



Developer Guide, Version 1

AWS IoT Greengrass



AWS IoT Greengrass: Developer Guide, Version 1

Copyright © 2025 Amazon Web Services, Inc. and/or its affiliates. All rights reserved.

Amazon's trademarks and trade dress may not be used in connection with any product or service that is not Amazon's, in any manner that is likely to cause confusion among customers, or in any manner that disparages or discredits Amazon. All other trademarks not owned by Amazon are the property of their respective owners, who may or may not be affiliated with, connected to, or sponsored by Amazon.

Table of Contents

.....	xxi
What is AWS IoT Greengrass?	1
AWS IoT Greengrass Core software	3
AWS IoT Greengrass Core software versions	4
AWS IoT Greengrass groups	14
Devices in AWS IoT Greengrass	16
SDKs	18
Supported platforms and requirements	19
AWS IoT Greengrass downloads	32
AWS IoT Greengrass Core software	32
AWS IoT Greengrass snap software	40
AWS IoT Greengrass Docker software	41
AWS IoT Greengrass Core SDK	43
Supported machine learning runtimes and libraries	44
AWS IoT Greengrass ML SDK software	45
We want to hear from you	45
Install the AWS IoT Greengrass Core software	45
Download and extract a tar.gz file	46
Run the Greengrass device setup script	46
Install from an APT repository	46
Run AWS IoT Greengrass in a Docker container	48
Run AWS IoT Greengrass in a snap	49
Archive a core software installation	60
Configure the AWS IoT Greengrass core	62
AWS IoT Greengrass core configuration file	63
Service endpoints must match the certificate type	123
Connect on port 443 or through a network proxy	124
Configure a write directory	134
Configure MQTT settings	137
Activate automatic IP detection	155
Start Greengrass on system boot	159
See also	160
AWS IoT Greengrass V1 maintenance policy	161
AWS IoT Greengrass versioning scheme	161

Lifecycle phases for the AWS IoT Greengrass Core software	162
Maintenance policy for AWS IoT Greengrass Core software	162
Maintenance phase schedule	163
Deprecation schedule	163
Support policy for Lambda functions	163
Support policy for AWS IoT Device Tester for AWS IoT Greengrass V1	164
End of maintenance schedule	164
End of maintenance for AWS IoT Greengrass Core software v1.x Docker images	41
End of maintenance for AWS IoT Greengrass Core software v1.x APT repository	165
End of maintenance for AWS IoT Greengrass Core software v1.11.x Snap	165
Getting started with AWS IoT Greengrass	167
Choose how to get started	167
Requirements	170
Create an AWS account	171
Sign up for an AWS account	172
Create a user with administrative access	172
Quick start: Greengrass device setup	174
Requirements	174
Run Greengrass device setup	175
Troubleshooting issues	179
Greengrass device setup configuration options	180
Module 1: Environment setup for Greengrass	190
Setting up a Raspberry Pi	190
Setting up an Amazon EC2 instance	198
Setting up other devices	203
Module 2: Installing the AWS IoT Greengrass Core software	206
Provision an AWS IoT thing to use as a Greengrass core	207
Create an Greengrass group	211
Install and run AWS IoT Greengrass on the core device	212
Module 3 (part 1): Lambda functions on AWS IoT Greengrass	218
Create and package a Lambda function	219
Configure the Lambda function for AWS IoT Greengrass	224
Deploy cloud configurations to a core device	227
Verify the Lambda function is running on the core device	229
Module 3 (part 2): Lambda functions on AWS IoT Greengrass	230
Create and package the Lambda function	230

Configure long-lived Lambda functions for AWS IoT Greengrass	234
Test long-lived Lambda functions	235
Test on-demand Lambda functions	238
Module 4: Interacting with client devices in an AWS IoT Greengrass group	242
Create client devices in an AWS IoT Greengrass group	244
Configure subscriptions	247
Install the AWS IoT Device SDK for Python	248
Test communications	254
Module 5: Interacting with device shadows	258
Configure devices and subscriptions	260
Download required files	262
Test communications (device syncs disabled)	263
Test communications (device syncs enabled)	266
Module 6: Accessing other AWS services	267
Configure the group role	269
Create and configure the Lambda function	271
Configure subscriptions	274
Test communications	275
Module 7: Simulating hardware security integration	277
Install SoftHSM	278
Configure SoftHSM	278
Import the private key	279
Configure the Greengrass core	281
Test the configuration	284
See also	285
OTA updates of AWS IoT Greengrass Core software	286
Requirements	286
IAM permissions for OTA updates	287
Considerations	290
Greengrass OTA update agent	291
Integration with init systems	292
Managed respawn with OTA updates	292
Create an OTA update	294
CreateSoftwareUpdateJob API	297
Deploy AWS IoT Greengrass groups	300
Deploying groups (console)	301

Deploying groups (API)	302
Getting the group ID	304
Overview of the group object model	305
Groups	305
Group versions	306
Group components	307
Updating groups	308
See also	309
Get deployment notifications	310
Group deployment status change event	311
Prerequisites for creating EventBridge rules	312
Configure deployment notifications (console)	313
Configure deployment notifications (CLI)	314
Configure deployment notifications (AWS CloudFormation)	315
See also	315
Reset deployments	315
Reset deployments from the AWS IoT console	316
Reset deployments with the AWS IoT Greengrass API	316
See also	318
Create bulk deployments	318
Prerequisites	318
Create and upload the bulk deployment input file	319
Create and configure an IAM execution role for bulk deployments	321
Allow your execution role access to your S3 Bucket	324
Deploy the groups	325
Test the deployment	328
Troubleshooting bulk deployments	330
See also	332
Run local Lambda functions	333
SDKs	334
Migrating cloud-based Lambda functions	337
Reference functions by alias or version	338
Controlling Greengrass Lambda function execution	338
Group-specific configuration settings	339
Running a Lambda function as root	343
Considerations when choosing Lambda function containerization	344

Setting the default access identity for Lambda functions in a group	348
Setting default containerization for Lambda functions in a group	349
Communication flows	350
Communication using MQTT messages	351
Other communication flows	351
Retrieve the input topic (or subject)	352
Lifecycle configuration	354
Lambda executables	356
Create a Lambda executable	357
Run AWS IoT Greengrass in a Docker container	358
Prerequisites	360
Get the AWS IoT Greengrass container image from Amazon ECR	361
Create and configure the Greengrass group and core	365
Run AWS IoT Greengrass locally	365
Configure "No container" containerization for the group	368
Deploy Lambda functions to the Docker container	369
(Optional) Deploy client devices that interact with Greengrass in the Docker container	369
Stopping the AWS IoT Greengrass Docker container	370
Troubleshooting AWS IoT Greengrass in a Docker container	370
Access local resources	374
Supported resource types	374
Requirements	375
Volume resources under the /proc directory	376
Group owner file access permission	376
See also	377
Using the CLI	377
Create local resources	378
Create the Greengrass function	380
Add the Lambda function to the group	381
Troubleshooting	383
Using the console	384
Prerequisites	385
Create a Lambda function deployment package	385
Create and publish a Lambda function	387
Add the Lambda function to the group	389
Add a local resource to the group	390

Add subscriptions to the group	391
Deploy the group	392
Test local resource access	393
Perform machine learning inference	396
How AWS IoT Greengrass ML inference works	396
Machine learning resources	397
Supported model sources	397
Requirements	400
Runtimes and libraries for ML inference	400
SageMaker AI Neo deep learning runtime	400
MXNet versioning	401
MXNet on Raspberry Pi	401
TensorFlow model-serving limitations on Raspberry Pi	401
Access machine learning resources	402
Access permissions for machine learning resources	402
Defining access permissions for Lambda functions (console)	405
Defining access permissions for Lambda functions (API)	406
Accessing machine learning resources from Lambda function code	409
Troubleshooting	410
See also	412
How to configure machine learning inference	412
Prerequisites	413
Configure the Raspberry Pi	414
Install the MXNet framework	416
Create a model package	416
Create and publish a Lambda function	417
Add the Lambda function to the group	420
Add resources to the group	422
Add a subscription to the group	424
Deploy the group	425
Test the app	426
Next steps	430
Configuring an Intel Atom	430
Configuring an NVIDIA Jetson TX2	434
How to configure optimized machine learning inference	438
Prerequisites	413

Configure the Raspberry Pi	440
Install the Neo deep learning runtime	442
Create an inference Lambda function	443
Add the Lambda function to the group	446
Add a Neo-optimized model resource to the group	448
Add your camera device resource to the group	450
Add subscriptions to the group	452
Deploy the group	452
Test the example	453
Configuring an Intel Atom	454
Configuring an NVIDIA Jetson TX2	457
Troubleshooting AWS IoT Greengrass ML inference	427
Next steps	464
Manage data streams	465
Stream management workflow	466
Requirements	468
Data security	469
Local data security	469
Client authentication	470
See also	470
Configure stream manager	471
Stream manager parameters	471
Configure settings (console)	474
Configure settings (CLI)	477
See also	487
Use StreamManagerClient to work with streams	487
Create message stream	488
Append message	493
Read messages	499
List streams	502
Describe message stream	503
Update message stream	506
Delete message stream	510
See also	511
Export configurations for supported AWS Cloud destinations	512
Export data streams (console)	528

Prerequisites	529
Create a Lambda function deployment package	531
Create a Lambda function	535
Add a function to the group	537
Enable stream manager	538
Configure local logging	538
Deploy the group	539
Test the application	540
See also	541
Export data streams (CLI)	542
Prerequisites	543
Create a Lambda function deployment package	545
Create a Lambda function	549
Create a function definition and version	551
Create a logger definition and version	553
Get the ARN of your core definition version	554
Create a group version	555
Create a deployment	556
Test the application	557
See also	559
Deploy secrets to the core	560
Secrets encryption	561
Requirements	562
Specify the private key for secret encryption	563
Allow AWS IoT Greengrass to get secret values	564
See also	566
Work with secret resources	566
Creating and managing secrets	566
Using local secrets	571
How to create a secret resource (console)	574
Prerequisites	576
Create a Secrets Manager secret	576
Add a secret resource to a group	577
Create a Lambda function deployment package	578
Create a Lambda function	580
Add the function to the group	582

Attach the secret resource to the function	583
Add subscriptions to the group	584
Deploy the group	584
Test the Lambda function	586
See also	586
Integrate with services and protocols using connectors	587
Requirements	588
Using Greengrass connectors	589
Configuration parameters	591
Parameters used to access group resources	591
Updating connector parameters	592
Inputs and outputs	592
Input topics	593
Containerization support	594
Upgrading connector versions	594
Logging	595
AWS-provided Greengrass connectors	596
CloudWatch Metrics	599
Device Defender	615
Docker application deployment	621
IoT Analytics	664
IoT Ethernet IP Protocol Adapter	680
IoT SiteWise	685
Kinesis Firehose	700
ML Feedback	717
ML Image Classification	735
ML Object Detection	760
Modbus-RTU Protocol Adapter	777
Modbus-TCP Protocol Adapter	796
Raspberry Pi GPIO	801
Serial Stream	811
ServiceNow MetricBase Integration	825
SNS	840
Splunk Integration	851
Twilio Notifications	865
Get started with connectors (console)	882

Prerequisites	883
Create a Secrets Manager secret	884
Add a secret resource to a group	885
Add a connector to the group	886
Create a Lambda function deployment package	886
Create a Lambda function	888
Add a function to the group	890
Add subscriptions to the group	890
Deploy the group	891
Test the solution	893
See also	894
Get started with connectors (CLI)	894
Prerequisites	896
Create a Secrets Manager secret	897
Create a resource definition and version	898
Create a connector definition and version	899
Create a Lambda function deployment package	900
Create a Lambda function	902
Create a function definition and version	904
Create a subscription definition and version	905
Create a group version	906
Create a deployment	908
Test the solution	909
See also	910
Greengrass Discovery RESTful API	911
Request	911
Response	912
Discovery authorization	912
Example discover response documents	913
Security	916
Overview of AWS IoT Greengrass security	917
Device connection workflow	918
Configuring AWS IoT Greengrass security	919
Security principals	920
Managed subscriptions in the MQTT messaging workflow	923
TLS cipher suites support	923

Data protection	926
Data encryption	927
Hardware security integration	930
Device authentication and authorization	947
X.509 certificates	948
AWS IoT policies	950
Minimal AWS IoT policy for the core device	953
Identity and access management	957
Audience	957
Authenticating with identities	958
Managing access using policies	961
See also	963
How AWS IoT Greengrass works with IAM	963
Greengrass service role	972
Greengrass group role	980
Cross-service confused deputy prevention	990
Identity-based policy examples	991
Troubleshooting identity and access issues	994
Compliance validation	997
Resilience	998
Infrastructure security	999
Configuration and vulnerability analysis	999
VPC endpoints (AWS PrivateLink)	1000
Considerations for AWS IoT Greengrass VPC endpoints	1001
Create an interface VPC endpoint for AWS IoT Greengrass control plane operations	1002
Creating a VPC endpoint policy for AWS IoT Greengrass	1002
Security best practices	1003
Grant minimum possible permissions	1003
Don't hardcode credentials in Lambda functions	1003
Don't log sensitive information	1004
Create targeted subscriptions	1004
Keep your device clock in sync	1004
Manage device authentication with the Greengrass core	1005
See also	1006
Logging and monitoring	1007
Monitoring tools	1007

See also	1008
Monitoring with AWS IoT Greengrass logs	1008
Accessing CloudWatch Logs	1008
Accessing file system logs	1010
Default logging configuration	1011
Configure logging for AWS IoT Greengrass	1012
Logging limitations	1015
CloudTrail logs	1016
Logging AWS IoT Greengrass API calls with AWS CloudTrail	1016
AWS IoT Greengrass information in CloudTrail	1017
Understanding AWS IoT Greengrass log file entries	1018
See also	1021
Gathering system health telemetry data	1021
Configuring telemetry settings	1024
Subscribing to receive telemetry data	1028
Troubleshooting AWS IoT Greengrass telemetry	1035
Calling the local health check API	1035
Get health information for all workers	1036
Get health information about specified workers	1037
Worker health information	1039
Tagging your Greengrass resources	1043
Tag basics	1043
Tagging support (console)	1043
Tagging support (API)	1044
Using tags with IAM policies	1045
Example IAM policies	1046
See also	1048
AWS CloudFormation support for AWS IoT Greengrass	1049
Create resources	1049
Deploy resources	1050
Example template	1051
Supported AWS Regions	1064
Using AWS IoT Device Tester for AWS IoT Greengrass V1	1065
AWS IoT Greengrass qualification suite	1065
Custom test suites	1066
Supported versions of AWS IoT Device Tester for AWS IoT Greengrass V1	1066

Unsupported IDT versions for for AWS IoT Greengrass	1067
Use IDT to run the AWS IoT Greengrass qualification suite	1072
Test suite versions	1073
Test group descriptions	1074
Prerequisites	1078
Configure your device to run IDT tests	1088
Configure IDT settings	1109
Run the AWS IoT Greengrass qualification suite	1124
Understanding results and logs	1129
Use IDT to develop and run your own test suites	1133
Download the latest version of IDT for AWS IoT Greengrass	1078
Test suite creation workflow	1134
Tutorial: Build and run the sample IDT test suite	1134
Tutorial: Develop a simple IDT test suite	1139
Create IDT test suite configuration files	1149
Configure the IDT state machine	1156
Create IDT test case executables	1180
Use the IDT context	1187
Configure settings for test runners	1191
Debug and run custom test suites	1202
Review IDT test results and logs	1205
IDT usage metrics	1211
IDT for AWS IoT Greengrass troubleshooting	1218
Error codes	1218
Resolving IDT for AWS IoT Greengrass errors	1238
Support policy for AWS IoT Device Tester for AWS IoT Greengrass V1	1243
Troubleshooting	1244
AWS IoT Greengrass Core issues	1244
Error: The configuration file is missing the CaPath, CertPath or KeyPath. The Greengrass daemon process with [pid = <pid>] died.	1246
Error: Failed to parse /<greengrass-root>/config/config.json.	1247
Error: Error occurred while generating TLS config: ErrUnknownURIScheme	1247
Error: Runtime failed to start: unable to start workers: container test timed out.	1247
Error: Failed to invoke PutLogEvents on local Cloudwatch, logGroup: /GreengrassSystem/connection_manager, error: RequestError: send request failed caused by: Post http://	

<path>/cloudwatch/logs/: dial tcp <address>: getsockopt: connection refused, response: { }. 1248

Error: Unable to create server due to: failed to load group: chmod /<greengrass-root>/ggc/deployment/lambda/arn:aws:lambda:<region>:<account-id>:function:<function-name>:<version>/<file-name>: no such file or directory. 1248

The AWS IoT Greengrass Core software doesn't start after you changed from running with no containerization to running in a Greengrass container. 1249

Error: Spool size should be at least 262144 bytes. 1249

Error: [ERROR]-Cloud messaging error: Error occurred while trying to publish a message. {"errorString": "operation timed out"} 1249

Error: container_linux.go:344: starting container process caused "process_linux.go:424: container init caused \"rootfs_linux.go:64: mounting \\\"/greengrass/ggc/socket/greengrass_ipc.sock\\\" to rootfs \\\"/greengrass/ggc/packages/<version>/rootfs/merged\\\" at \\\"/greengrass_ipc.sock\\\" caused \\\"stat /greengrass/ggc/socket/greengrass_ipc.sock: permission denied\\\"\". 1250

Error: Greengrass daemon running with PID: <process-id>. Some system components failed to start. Check 'runtime.log' for errors. 1250

Device shadow does not sync with the cloud. 995

ERROR: unable to accept TCP connection. accept tcp [::]:8000: accept4: too many open files. 1251

Error: Runtime execution error: unable to start lambda container. container_linux.go:259: starting container process caused "process_linux.go:345: container init caused \"rootfs_linux.go:50: preparing rootfs caused \\\"permission denied\\\"\". 1251

Warning: [WARN]-[5]GK Remote: Error retrieving public key data: ErrPrincipalNotConfigured: private key for MqttCertificate is not set. 1252

Error: Permission denied when attempting to use role arn:aws:iam::<account-id>:role/<role-name> to access s3 url https://<region>-greengrass-updates.s3.<region>.amazonaws.com/core/<architecture>/greengrass-core-<distribution-version>.tar.gz. 995

The AWS IoT Greengrass core is configured to use a network proxy and your Lambda function can't make outgoing connections. 1252

The core is in an infinite connect-disconnect loop. The runtime.log file contains a continuous series of connect and disconnect entries. 1253

Error: unable to start lambda container. container_linux.go:259: starting container process caused "process_linux.go:345: container init caused \"rootfs_linux.go:62: mounting \\\"proc\\\" to rootfs \\\"\" 1254

[ERROR]-runtime execution error: unable to start lambda container. {"errorString": "failed to initialize container mounts: failed to mask greengrass root in overlay upper dir: failed to create mask device at directory <ggc-path>: file exists"}	1254
[ERROR]-Deployment failed. {"deploymentId": "<deployment-id>", "errorString": "container test process with pid <pid> failed: container process state: exit status 1"}	1255
Error: [ERROR]-runtime execution error: unable to start lambda container. {"errorString": "failed to initialize container mounts: failed to create overlay fs for container: mounting overlay at /greengrass/ggc/packages/<ggc- version>/rootfs/merged failed: failed to mount with args source=\"no_source \" dest=\"/greengrass/ggc/packages/<ggc-version>/rootfs/merged\" fstype=\"overlay\" flags=\"0\" data=\"lowerdir=/greengrass/ggc/packages/<ggc-version>/dns:/,upperdir=/ greengrass/ggc/packages/<ggc-version>/rootfs/upper,workdir=/greengrass/ggc/ packages/<ggc-version>/rootfs/work\": too many levels of symbolic links"}	1256
Error: [DEBUG]-Failed to get routes. Discarding message.	1257
Error: [Errno 24] Too many open <lambda-function>,[Errno 24] Too many open files	1257
Error: ds server failed to start listening to socket: listen unix <ggc-path>/ggc/socket/ greengrass_ipc.sock: bind: invalid argument	1257
[INFO] (Copier) aws.greengrass.StreamManager: stdout. Caused by: com.fasterxml.jackson.databind.JsonMappingException: Instant exceeds minimum or maximum instant	1257
GPG error: https://dnw9lb6lzp2d8.cloudfront.net stable InRelease: The following signatures were invalid: EXPKEYSIG 68D644ABD2327D47 AWS Greengrass Master Key ..	1258
Deployment issues	1258
Your current deployment does not work and you want to revert to a previous working deployment.	1260
You see a 403 Forbidden error on deployment in the logs.	1262
A ConcurrentDeployment error occurs when you run the create-deployment command for the first time.	1262
Error: Greengrass is not authorized to assume the Service Role associated with this account, or the error: Failed: TES service role is not associated with this account.	995
Error: unable to execute download step in deployment. error while downloading: error while downloading the Group definition file: ... x509: certificate has expired or is not yet valid	1263
The deployment doesn't finish.	1263

Error: Unable to find java or java8 executables, or the error: Deployment <deployment-id> of type NewDeployment for group <group-id> failed error: worker with <worker-id> failed to initialize with reason Installed Java version must be greater than or equal to 8	1264
The deployment doesn't finish, and runtime.log contains multiple "wait 1s for container to stop" entries.	1264
The deployment doesn't finish, and runtime.log contains "[ERROR]-Greengrass deployment error: failed to report deployment status back to cloud {"deploymentId": "<deployment-id>", "errorString": "Failed to initiate PUT, endpoint: https://<deployment-status>, error: Put https://<deployment-status>: proxyconnect tcp: x509: certificate signed by unknown authority"}"	1265
Error: Deployment <deployment-id> of type NewDeployment for group <group-id> failed error: Error while processing. group config is invalid: 112 or [119 0] don't have rw permission on the file: <path>.	1266
Error: <list-of-function-arns> are configured to run as root but Greengrass is not configured to run Lambda functions with root permissions.	1266
Error: Deployment <deployment-id> of type NewDeployment for group <group-id> failed error: Greengrass deployment error: unable to execute download step in deployment. error while processing: unable to load the group file downloaded: could not find UID based on user name, userName: ggc_user: user: unknown user ggc_user.	1266
Error: [ERROR]-runtime execution error: unable to start lambda container. {"errorString": "failed to initialize container mounts: failed to mask greengrass root in overlay upper dir: failed to create mask device at directory <ggc-path>: file exists"}	1267
Error: Deployment <deployment-id> of type NewDeployment for group <group-id> failed error: process start failed: container_linux.go:259: starting container process caused "process_linux.go:250: running exec setns process for init caused \"wait: no child processes \\\"".	1267
Error: [WARN]-MQTT[client] dial tcp: lookup <host-prefix>-ats.iot.<region>.amazonaws.com: no such host ... [ERROR]-Greengrass deployment error: failed to report deployment status back to cloud ... net/http: request canceled while waiting for connection (Client.Timeout exceeded while awaiting headers)	1268
Create group and create function issues	1268
Error: Your 'IsolationMode' configuration for the group is invalid.	1269
Error: Your 'IsolationMode' configuration for function with arn <function-arn> is invalid.	1269
Error: MemorySize configuration for function with arn <function-arn> is not allowed in IsolationMode=NoContainer.	1269

Error: Access Sysfs configuration for function with arn <function-arn> is not allowed in IsolationMode=NoContainer.	1270
Error: MemorySize configuration for function with arn <function-arn> is required in IsolationMode=GreengrassContainer.	1270
Error: Function <function-arn> refers to resource of type <resource-type> that is not allowed in IsolationMode=NoContainer.	1270
Error: Execution configuration for function with arn <function-arn> is not allowed.	1271
Discovery issues	1271
Error: Device is a member of too many groups, devices may not be in more than 10 groups	1271
Machine learning resource issues	1271
InvalidMLModelOwner - GroupOwnerSetting is provided in ML model resource, but GroupOwner or GroupPermission is not present	410
NoContainer function cannot configure permission when attaching Machine Learning resources. <function-arn> refers to Machine Learning resource <resource-id> with permission <ro/rw> in resource access policy.	411
Function <function-arn> refers to Machine Learning resource <resource-id> with missing permission in both ResourceAccessPolicy and resource OwnerSetting.	411
Function <function-arn> refers to Machine Learning resource <resource-id> with permission \"rw\", while resource owner setting GroupPermission only allows \"ro\".	411
NoContainer Function <function-arn> refers to resources of nested destination path.	411
Lambda <function-arn> gains access to resource <resource-id> by sharing the same group owner id	412
AWS IoT Greengrass core in Docker issues	1274
Error: Unknown options: -no-include-email.	370
Warning: IPv4 is disabled. Networking will not work.	370
Error: A firewall is blocking file Sharing between windows and the containers.	371
Error: An error occurred (AccessDeniedException) when calling the GetAuthorizationToken operation: User: arn:aws:iam::<account-id>:user/<user-name> is not authorized to perform: ecr:GetAuthorizationToken on resource: *	371
Error: Cannot create container for the service greengrass: Conflict. The container name \"/aws-iot-greengrass\" is already in use.	1275
Error: [FATAL]-Failed to reset thread's mount namespace due to an unexpected error: \"operation not permitted\". To maintain consistency, GGC will crash and need to be manually restarted.	1276
Troubleshooting with logs	1276


Troubleshooting storage issues	1277
Troubleshooting messages	1278
Troubleshooting shadow synchronization timeout issues	1278
Check AWS re:Post	1279
Document history	1280
Earlier updates	1301

AWS IoT Greengrass Version 1 entered the extended life phase on June 30, 2023. For more information, see the [AWS IoT Greengrass V1 maintenance policy](#). After this date, AWS IoT Greengrass V1 won't release updates that provide features, enhancements, bug fixes, or security patches. Devices that run on AWS IoT Greengrass V1 won't be disrupted and will continue to operate and to connect to the cloud. We strongly recommend that you [migrate to AWS IoT Greengrass Version 2](#), which adds [significant new features](#) and [support for additional platforms](#).

What is AWS IoT Greengrass?

AWS IoT Greengrass is software that extends cloud capabilities to local devices. This enables devices to collect and analyze data closer to the source of information, react autonomously to local events, and communicate securely with each other on local networks. Local devices can also communicate securely with AWS IoT Core and export IoT data to the AWS Cloud. AWS IoT Greengrass developers can use AWS Lambda functions and prebuilt [connectors](#) to create serverless applications that are deployed to devices for local execution.

The following diagram shows the basic architecture of AWS IoT Greengrass.



AWS IoT Greengrass makes it possible for customers to build IoT devices and application logic. Specifically, AWS IoT Greengrass provides cloud-based management of application logic that runs on devices. Locally deployed Lambda functions and connectors are triggered by local events, messages from the cloud, or other sources.

In AWS IoT Greengrass, devices securely communicate on a local network and exchange messages with each other without having to connect to the cloud. AWS IoT Greengrass provides a local pub/sub message manager that can intelligently buffer messages if connectivity is lost so that inbound and outbound messages to the cloud are preserved.

AWS IoT Greengrass protects user data:

- Through the secure authentication and authorization of devices.
- Through secure connectivity in the local network.

- Between local devices and the cloud.

Device security credentials function in a group until they are revoked, even if connectivity to the cloud is disrupted, so that the devices can continue to securely communicate locally.

AWS IoT Greengrass provides secure, over-the-air updates of Lambda functions.

AWS IoT Greengrass consists of:

- Software distributions
 - AWS IoT Greengrass Core software
 - AWS IoT Greengrass Core SDK
- Cloud service
 - AWS IoT Greengrass API
- Features
 - Lambda runtime
 - Shadows implementation
 - Message manager
 - Group management
 - Discovery service
 - Over-the-air update agent
 - Stream manager
 - Local resource access
 - Local machine learning inference
 - Local secrets manager
 - Connectors with built-in integration with services, protocols, and software

Topics

- [AWS IoT Greengrass Core software](#)
- [AWS IoT Greengrass groups](#)
- [Devices in AWS IoT Greengrass](#)
- [SDKs](#)
- [Supported platforms and requirements](#)

- [AWS IoT Greengrass downloads](#)
- [We want to hear from you](#)
- [Install the AWS IoT Greengrass Core software](#)
- [Configure the AWS IoT Greengrass core](#)

AWS IoT Greengrass Core software

The AWS IoT Greengrass Core software provides the following functionality:

- Deployment and the local running of connectors and Lambda functions.
- Process data streams locally with automatic exports to the AWS Cloud.
- MQTT messaging over the local network between devices, connectors, and Lambda functions using managed subscriptions.
- MQTT messaging between AWS IoT and devices, connectors, and Lambda functions using managed subscriptions.
- Secure connections between devices and the AWS Cloud using device authentication and authorization.
- Local shadow synchronization of devices. Shadows can be configured to sync with the AWS Cloud.
- Controlled access to local device and volume resources.
- Deployment of cloud-trained machine learning models for running local inference.
- Automatic IP address detection that enables devices to discover the Greengrass core device.
- Central deployment of new or updated group configuration. After the configuration data is downloaded, the core device is restarted automatically.
- Secure, over-the-air (OTA) software updates of user-defined Lambda functions.
- Secure, encrypted storage of local secrets and controlled access by connectors and Lambda functions.

AWS IoT Greengrass core instances are configured through AWS IoT Greengrass APIs that create and update AWS IoT Greengrass group definitions stored in the cloud.

AWS IoT Greengrass Core software versions

AWS IoT Greengrass provides several options for installing the AWS IoT Greengrass Core software, including tar.gz download files, a quick start script, and apt installations on supported Debian platforms. For more information, see [the section called “Install the AWS IoT Greengrass Core software”](#).

The following tabs describe what's new and changed in AWS IoT Greengrass Core software versions.

GGC v1.11

1.11.6

Bug fixes and improvements:

- Improved resilience if sudden power loss occurs during a deployment.
- Fixed an issue where stream manager data corruption could prevent the AWS IoT Greengrass Core software from starting.
- Fixed an issue where new client devices couldn't connect to the core in certain scenarios.
- Fixed an issue where stream manager stream names couldn't contain `.log`.

1.11.5

Bug fixes and improvements:

- General performance improvements and bug fixes.

1.11.4

Bug fixes and improvements:

- Fixed an issue with stream manager that prevented upgrades to AWS IoT Greengrass Core software v1.11.3. If you are using stream manager to export data to the cloud, you can now use an OTA update to upgrade an earlier v1.x version of the AWS IoT Greengrass Core software to v1.11.4.
- General performance improvements and bug fixes.

1.11.3

Bug fixes and improvements:

- Fixed an issue that caused AWS IoT Greengrass Core software running in a snap on an Ubuntu device to stop responding after a sudden power loss to the device.

- Fixed an issue that caused delayed delivery of MQTT messages to long-lived Lambda functions.
- Fixed an issue that caused MQTT messages to not be sent correctly when the `maxWorkItemCount` value was set to a value greater than 1024.
- Fixed an issue that caused the OTA update agent to ignore the MQTT KeepAlive period specified in the `keepAlive` property in [config.json](#).
- General performance improvements and bug fixes.

Important

If you are using stream manager to export data to the cloud, do *not* upgrade to AWS IoT Greengrass Core software v1.11.3 from an earlier v1.x version. If you are enabling stream manager for the first time, we strongly recommend that you first install the latest version of the AWS IoT Greengrass Core software.

1.11.1

Bug fixes and improvements:

- Fixed an issue that caused increased memory use for stream manager.
- Fixed an issue that caused stream manager to reset the sequence number of the stream to 0 if the Greengrass core device was turned off for longer than the specified time-to-live (TTL) period of the stream data.
- Fixed an issue that prevented stream manager from correctly stopping retry attempts to export data to the AWS Cloud.

1.11.0

New features:

- A telemetry agent on the Greengrass core collects local telemetry data and publishes it to AWS Cloud. To retrieve the telemetry data for further processing, customers can create an Amazon EventBridge rule and subscribe to a target. For more information, see [Gathering system health telemetry data from AWS IoT Greengrass core devices](#).
- A local HTTP API returns a snapshot of the current state of local worker processes started by AWS IoT Greengrass. For more information, see [Calling the local health check API](#).
- A [stream manager](#) automatically exports data to Amazon S3 and AWS IoT SiteWise.

New [stream manager parameters](#) let you update existing streams and pause or resume data export.

- Support for running Python 3.8.x Lambda functions on the core.
- A new `ggDaemonPort` property in [config.json](#) that use to configure the Greengrass core IPC port number. The default port number is 8000.

A new `systemComponentAuthTimeout` property in [config.json](#) that you use to configure the timeout for Greengrass core IPC authentication. The default timeout is 5000 milliseconds.

- Increased the maximum number of AWS IoT devices per AWS IoT Greengrass group from 200 to 2500.

Increased the maximum number of subscriptions per group from 1000 to 10000.

For more information, see [AWS IoT Greengrass endpoints and quotas](#).

Bug fixes and improvements:

- General optimization that can reduce the memory utilization of the Greengrass service processes.
- A new runtime configuration parameter (`mountAllBlockDevices`) lets Greengrass use bind mounts to mount all block devices into a container after setting up the OverlayFS. This feature resolved an issue that caused Greengrass deployment failure if `/usr` isn't under the `/` hierarchy.
- Fixed an issue that caused AWS IoT Greengrass core failure if `/tmp` is a symlink.
- Fixed an issue to let the Greengrass deployment agent remove unused machine learning model artifacts from the `mlmodel_public` folder.
- General performance improvements and bug fixes.

Extended life versions

1.10.5

Bug fixes and improvements:

- General performance improvements and bug fixes.

1.10.4

Bug fixes and improvements:

- Fixed an issue that caused AWS IoT Greengrass Core software running in a snap on an Ubuntu device to stop responding after a sudden power loss to the device.
- Fixed an issue that caused delayed delivery of MQTT messages to long-lived Lambda functions.
- Fixed an issue that caused MQTT messages to not be sent correctly when the `maxWorkItemCount` value was set to a value greater than 1024.
- Fixed an issue that caused the OTA update agent to ignore the MQTT KeepAlive period specified in the `keepAlive` property in [config.json](#).
- General performance improvements and bug fixes.

1.10.3

Bug fixes and improvements:

- A new `systemComponentAuthTimeout` property in [config.json](#) that you use to configure the timeout for Greengrass core IPC authentication. The default timeout is 5000 milliseconds.
- Fixed an issue that caused increased memory use for stream manager.

1.10.2

Bug fixes and improvements:

- A new `mqttOperationTimeout` property in [config.json](#) that you use to set the timeout for publish, subscribe, and unsubscribe operations in MQTT connections with AWS IoT Core.
- General performance improvements and bug fixes.

1.10.1

Bug fixes and improvements:

- [Stream manager](#) is more resilient to file data corruption.
- Fixed an issue that causes a `sysfs` mount failure on devices using Linux kernel 5.1 and later.
- General performance improvements and bug fixes.

1.10.0

New features:

- A stream manager that processes data streams locally and exports them to the AWS Cloud automatically. This feature requires Java 8 on the Greengrass core device. For more information, see [Manage data streams](#).
- A new Greengrass Docker application deployment connector that runs a Docker application on a core device. For more information, see [the section called “Docker application deployment”](#).
- A new IoT SiteWise connector that sends industrial device data from OPC-UA servers to asset properties in AWS IoT SiteWise. For more information, see [the section called “IoT SiteWise”](#).
- Lambda functions that run without containerization can access machine learning resources in the Greengrass group. For more information, see [the section called “Access machine learning resources”](#).
- Support for MQTT persistent sessions with AWS IoT. For more information, see [the section called “MQTT persistent sessions with AWS IoT Core”](#).
- Local MQTT traffic can travel over a port other than the default port 8883. For more information, see [the section called “MQTT port for local messaging”](#).
- New queueFullPolicy options in the [AWS IoT Greengrass Core SDK](#) for reliable message publishing from Lambda functions.
- Support for running Node.js 12.x Lambda functions on the core.
- Over-the-air (OTA) updates with hardware security integration can be configured with OpenSSL 1.1.
- General performance improvements and bug fixes.

1.9.4

Bug fixes and improvements:

- General performance improvements and bug fixes.

1.9.3

New features:

- Support for Armv6L. AWS IoT Greengrass Core software v1.9.3 or later can be installed on Raspbian distributions on Armv6L architectures (for example, on Raspberry Pi Zero devices).

- OTA updates on port 443 with ALPN. Greengrass cores that use port 443 for MQTT traffic now support over-the-air (OTA) software updates. AWS IoT Greengrass uses the Application Layer Protocol Network (ALPN) TLS extension to enable these connections. For more information, see [OTA updates of AWS IoT Greengrass Core software](#) and [the section called “Connect on port 443 or through a network proxy”](#).

Bug fixes and improvements:

- Fixes a bug introduced in v1.9.0 that prevented Python 2.7 Lambda functions from sending binary payloads to other Lambda functions.
- General performance improvements and bug fixes.

1.9.2

New features:

- Support for [OpenWrt](#). AWS IoT Greengrass Core software v1.9.2 or later can be installed on OpenWrt distributions with Armv8 (AArch64) and Armv7l architectures. Currently, OpenWrt does not support ML inference.

1.9.1

Bug fixes and improvements:

- Fixes a bug introduced in v1.9.0 that drops messages from the cloud that contain wildcard characters in the topic.

1.9.0

New features:

- Support for Python 3.7 and Node.js 8.10 Lambda runtimes. Lambda functions that use Python 3.7 and Node.js 8.10 runtimes can now run on an AWS IoT Greengrass core. (AWS IoT Greengrass continues to support the Python 2.7 and Node.js 6.10 runtimes.)
- Optimized MQTT connections. The Greengrass core establishes fewer connections with the AWS IoT Core. This change can reduce operational costs for charges that are based on the number of connections.
- Elliptic Curve (EC) key for the local MQTT server. The local MQTT server supports EC keys in addition to RSA keys. (The MQTT server certificate has an SHA-256 RSA signature, regardless of the key type.) For more information, see [the section called “Security principals”](#).

Bug fixes and improvements:

- General performance improvements and bug fixes.

1.8.4

Fixed an issue with shadow synchronization and device certificate manager reconnection.

General performance improvements and bug fixes.

1.8.3

General performance improvements and bug fixes.

1.8.2

General performance improvements and bug fixes.

1.8.1

General performance improvements and bug fixes.

1.8.0

New features:

- Configurable default access identity for Lambda functions in the group. This group-level setting determines the default permissions that are used to run Lambda functions. You can set the user ID, group ID, or both. Individual Lambda functions can override the default access identity of their group. For more information, see [the section called “Setting the default access identity for Lambda functions in a group”](#).
- HTTPS traffic over port 443. HTTPS communication can be configured to travel over port 443 instead of the default port 8443. This complements AWS IoT Greengrass support for the Application Layer Protocol Network (ALPN) TLS extension and allows all Greengrass messaging traffic—both MQTT and HTTPS—to use port 443. For more information, see [the section called “Connect on port 443 or through a network proxy”](#).
- Predictably named client IDs for AWS IoT connections. This change enables support for AWS IoT Device Defender and [AWS IoT lifecycle events](#), so you can receive notifications for connect, disconnect, subscribe, and unsubscribe events. Predictable naming also makes it easier to create logic around connection IDs (for example, to create [subscribe policy templates](#) based on certificate attributes). For more information, see [the section called “Client IDs for MQTT connections with AWS IoT”](#).

Bug fixes and improvements:

- Fixed an issue with shadow synchronization and device certificate manager reconnection.

- General performance improvements and bug fixes.

1.7.1

New features:

- Greengrass connectors provide built-in integration with local infrastructure, device protocols, AWS, and other cloud services. For more information, see [Integrate with services and protocols using connectors](#).
- AWS IoT Greengrass extends AWS Secrets Manager to core devices, which makes your passwords, tokens, and other secrets available to connectors and Lambda functions. Secrets are encrypted in transit and at rest. For more information, see [Deploy secrets to the core](#).
- Support for a hardware root of trust security option. For more information, see [the section called "Hardware security integration"](#).
- Isolation and permission settings that allow Lambda functions to run without Greengrass containers and to use the permissions of a specified user and group. For more information, see [the section called "Controlling Greengrass Lambda function execution"](#).
- You can run AWS IoT Greengrass in a Docker container (on Windows, macOS, or Linux) by configuring your Greengrass group to run with no containerization. For more information, see [the section called "Run AWS IoT Greengrass in a Docker container"](#).
- MQTT messaging on port 443 with Application Layer Protocol Negotiation (ALPN) or connection through a network proxy. For more information, see [the section called "Connect on port 443 or through a network proxy"](#).
- The SageMaker AI Neo deep learning runtime, which supports machine learning models that have been optimized by the SageMaker AI Neo deep learning compiler. For information about the Neo deep learning runtime, see [the section called "Runtimes and libraries for ML inference"](#).
- Support for Raspbian Stretch (2018-06-27) on Raspberry Pi core devices.

Bug fixes and improvements:

- General performance improvements and bug fixes.

In addition, the following features are available with this release:

- The AWS IoT Device Tester for AWS IoT Greengrass, which you can use to verify that your CPU architecture, kernel configuration, and drivers work with AWS IoT Greengrass. For more information, see [Using AWS IoT Device Tester for AWS IoT Greengrass V1](#).

- The AWS IoT Greengrass Core software, AWS IoT Greengrass Core SDK, and AWS IoT Greengrass Machine Learning SDK packages are available for download through Amazon CloudFront. For more information, see [the section called “AWS IoT Greengrass downloads”](#).

1.6.1

New features:

- Lambda executables that run binary code on the Greengrass core. Use the new AWS IoT Greengrass Core SDK for C to write Lambda executables in C and C++. For more information, see [the section called “Lambda executables”](#).
- Optional local storage message cache that can persist across restarts. You can configure the storage settings for MQTT messages that are queued for processing. For more information, see [the section called “MQTT message queue”](#).
- Configurable maximum reconnect retry interval for when the core device is disconnected. For more information, see the `mqttMaxConnectionRetryInterval` property in [the section called “AWS IoT Greengrass core configuration file”](#).
- Local resource access to the host `/proc` directory. For more information, see [Access local resources](#).
- Configurable write directory. The AWS IoT Greengrass Core software can be deployed to read-only and read-write locations. For more information, see [the section called “Configure a write directory”](#).

Bug fixes and improvements:

- Performance improvement for publishing messages in the Greengrass core and between devices and the core.
- Reduced the compute resources required to process logs generated by user-defined Lambda functions.

1.5.0

New features:

- AWS IoT Greengrass Machine Learning (ML) Inference is generally available. You can perform ML inference locally on AWS IoT Greengrass devices using models that are built and trained in the cloud. For more information, see [Perform machine learning inference](#).
- Greengrass Lambda functions now support binary data as input payload, in addition to JSON. To use this feature, you must upgrade to AWS IoT Greengrass Core SDK version 1.1.0, which you can download from the [AWS IoT Greengrass Core SDK](#) downloads page.

Bug fixes and improvements:

- Reduced the overall memory footprint.
- Performance improvements for sending messages to the cloud.
- Performance and stability improvements for the download agent, Device Certificate Manager, and OTA update agent.
- Minor bug fixes.

1.3.0

New features:

- Over-the-air (OTA) update agent capable of handling cloud-deployed, Greengrass update jobs. The agent is found under the new `/greengrass/ota` directory. For more information, see [OTA updates of AWS IoT Greengrass Core software](#).
- Local resource access feature allows Greengrass Lambda functions to access local resources, such as peripheral devices and volumes. For more information, see [Access local resources with Lambda functions and connectors](#).

1.1.0

New features:

- Deployed AWS IoT Greengrass groups can be reset by deleting Lambda functions, subscriptions, and configurations. For more information, see [the section called “Reset deployments”](#).
- Support for Node.js 6.10 and Java 8 Lambda runtimes, in addition to Python 2.7.

To migrate from the previous version of the AWS IoT Greengrass core:


- Copy certificates from the `/greengrass/configuration/certs` folder to `/greengrass/certs`.
- Copy `/greengrass/configuration/config.json` to `/greengrass/config/config.json`.
- Run `/greengrass/ggc/core/greengrassd` instead of `/greengrass/greengrassd`.
- Deploy the group to the new core.

1.0.0

Initial version

AWS IoT Greengrass groups

A Greengrass group is a collection of settings and components, such as a Greengrass core, devices, and subscriptions. Groups are used to define a scope of interaction. For example, a group might represent one floor of a building, one truck, or an entire mining site. The following diagram shows the components that can make up a Greengrass group.



In the preceding diagram:

A: Greengrass group definition

Information about group settings and components.

B: Greengrass group settings

These include:

- Greengrass group role.
- Certificate authority and local connection configuration.
- Greengrass core connectivity information.
- Default Lambda runtime environment. For more information, see [the section called "Setting default containerization for Lambda functions in a group"](#).

- CloudWatch and local logs configuration. For more information, see [the section called “Monitoring with AWS IoT Greengrass logs”](#).

C: Greengrass core

The AWS IoT thing (device) that represents the Greengrass core. For more information, see [the section called “Configure the AWS IoT Greengrass core”](#).

D: Lambda function definition

A list of Lambda functions that run locally on the core, with associated configuration data. For more information, see [Run local Lambda functions](#).

E: Subscription definition

A list of subscriptions that enable communication using MQTT messages. A subscription defines:

- A message source and message target. These can be client devices, Lambda functions, connectors, AWS IoT Core, and the local shadow service.
- A topic or subject that's used to filter messages.

For more information, see [the section called “Managed subscriptions in the MQTT messaging workflow”](#).

F: Connector definition

A list of connectors that run locally on the core, with associated configuration data. For more information, see [Integrate with services and protocols using connectors](#).

G: Device definition

A list of AWS IoT things (known as client devices or devices) that are members of the Greengrass group, with associated configuration data. For more information, see [the section called “Devices in AWS IoT Greengrass”](#).

H: Resource definition

A list of local resources, machine learning resources, and secret resources on the Greengrass core, with associated configuration data. For more information, see [Access local resources](#), [Perform machine learning inference](#), and [Deploy secrets to the core](#).

When deployed, the Greengrass group definition, Lambda functions, connectors, resources, and subscription table are copied to the core device. For more information, see [Deploy AWS IoT Greengrass groups](#).

Devices in AWS IoT Greengrass

A Greengrass group can contain two types of AWS IoT device:

Greengrass core

A Greengrass core is a device that runs the AWS IoT Greengrass Core software, which allows it to communicate directly with AWS IoT Core and the AWS IoT Greengrass service. A core has its own device certificate used for authenticating with AWS IoT Core. It has a device shadow and an entry in the AWS IoT Core registry. Greengrass cores run a local Lambda runtime, deployment agent, and IP address tracker that sends IP address information to the AWS IoT Greengrass service to allow client devices to automatically discover their group and core connection information. For more information, see [the section called “Configure the AWS IoT Greengrass core”](#).

Note

A Greengrass group must contain exactly one core.

Client device

Client devices (also called *connected devices*, *Greengrass devices*, or *devices*) are devices that connect to a Greengrass core over MQTT. They have their own device certificate for AWS IoT Core authentication, a device shadow, and an entry in the AWS IoT Core registry. Client devices can run [FreeRTOS](#) or use the [AWS IoT Device SDK](#) or [AWS IoT Greengrass Discovery API](#) to get discovery information used to connect and authenticate with the core in the same Greengrass group. To learn how to use the AWS IoT console to create and configure a client device for AWS IoT Greengrass, see [the section called “Module 4: Interacting with client devices in an AWS IoT Greengrass group”](#). Or, for examples that show you how to use the AWS CLI to create and configure a client device for AWS IoT Greengrass, see [create-device-definition](#) in the *AWS CLI Command Reference*.

In a Greengrass group, you can create subscriptions that allow client devices to communicate over MQTT with Lambda functions, connectors, and other client devices in the group, and with AWS IoT Core or the local shadow service. MQTT messages are routed through the core. If the core device loses connectivity to the cloud, client devices can continue to communicate over the local network. Client devices can vary in size, from smaller microcontroller-based devices to

large appliances. Currently, a Greengrass group can contain up to 2,500 client devices. A client device can be a member of up to 10 groups.

Note

OPC-UA is an information exchange standard for industrial communication. To implement support for OPC-UA on the Greengrass core, you can use the [IoT SiteWise connector](#). The connector sends industrial device data from OPC-UA servers to asset properties in AWS IoT SiteWise.

The following table shows how these device types are related.

	Core	Device
Certificate	✓	✓
IoT Policy	✓	✓
IoT Thing	✓	✓
Device use	Gateway	Sensor and/or Actuator
Software	AWS IoT Greengrass Core Software	Amazon FreeRTOS / AWS IoT Device SDK
Group membership	✓	✓
Functions outside a Greengrass Group	✗	✓

The AWS IoT Greengrass core device stores certificates in two locations:

- Core device certificate in `/greengrass-root/certs`. Typically, the core device certificate is named `hash.cert.pem` (for example, `86c84488a5.cert.pem`). This certificate is used by the AWS IoT client for mutual authentication when the core connects to the AWS IoT Core and AWS IoT Greengrass services.
- MQTT server certificate in `/greengrass-root/ggc/var/state/server`. The MQTT server certificate is named `server.crt`. This certificate is used for mutual authentication between the local MQTT server (on the Greengrass core) and Greengrass devices.

Note

`greengrass-root` represents the path where the AWS IoT Greengrass Core software is installed on your device. Typically, this is the `/greengrass` directory.

SDKs

The following AWS-provided SDKs are used to work with AWS IoT Greengrass:

AWS SDK

Use the AWS SDK to build applications that interact with any AWS service, including Amazon S3, Amazon DynamoDB, AWS IoT, AWS IoT Greengrass, and more. In the context of AWS IoT Greengrass, you can use the AWS SDK in deployed Lambda functions to make direct calls to any AWS service. For more information, see [AWS SDKs](#).

Note

The operations specific to Greengrass that are available in the AWS SDKs are also available in the [AWS IoT Greengrass API](#) and [AWS CLI](#).

AWS IoT Device SDK

The AWS IoT Device SDK helps devices connect to AWS IoT Core and AWS IoT Greengrass. For more information, see [AWS IoT Device SDKs](#) in the *AWS IoT Developer Guide*.

Client devices can use any of the AWS IoT Device SDK v2 platforms to discover connectivity information for a Greengrass core. Connectivity information includes:

- The IDs of the Greengrass groups that the client device belongs to.

- The IP addresses of the Greengrass core in each group. These are also called *core endpoints*.
- The group CA certificate, which devices use for mutual authentication with the core. For more information, see [the section called “Device connection workflow”](#).

Note

In v1 of the AWS IoT Device SDKs, only the C++ and Python platforms provide built-in discovery support.

AWS IoT Greengrass Core SDK

The AWS IoT Greengrass Core SDK enables Lambda functions to interact with the Greengrass core, publish messages to AWS IoT, interact with the local shadow service, invoke other deployed Lambda functions, and access secret resources. This SDK is used by Lambda functions that run on an AWS IoT Greengrass core. For more information, see [AWS IoT Greengrass Core SDK](#).

AWS IoT Greengrass Machine Learning SDK

The AWS IoT Greengrass Machine Learning SDK enables Lambda functions to consume machine learning models that are deployed to the Greengrass core as machine learning resources. This SDK is used by Lambda functions that run on an AWS IoT Greengrass core and interact with a local inference service. For more information, see [AWS IoT Greengrass Machine Learning SDK](#).

Supported platforms and requirements

The following tabs list supported platforms and requirements for the AWS IoT Greengrass Core software.

Note

You can download the AWS IoT Greengrass Core software from the [AWS IoT Greengrass Core Software](#) downloads.

GGC v1.11

Supported platforms:

- Architecture: Armv7l
 - OS: Linux
 - OS: Linux ([OpenWrt](#))
- Architecture: Armv8 (AArch64)
 - OS: Linux
 - OS: Linux ([OpenWrt](#))
- Architecture: Armv6l
 - OS: Linux
- Architecture: x86_64
 - OS: Linux
- Windows, macOS, and Linux platforms can run AWS IoT Greengrass in a Docker container. For more information, see [the section called “Run AWS IoT Greengrass in a Docker container”](#).

Requirements:

- Minimum 128 MB disk space available for the AWS IoT Greengrass Core software. If you use the [OTA update agent](#), the minimum is 400 MB.
- Minimum 128 MB RAM allocated to the AWS IoT Greengrass Core software. With [stream manager](#) enabled, the minimum is 198 MB RAM.

Note

Stream manager is enabled by default if you use the **Default Group creation** option on the AWS IoT console to create your Greengrass group.

- Linux kernel version:
 - Linux kernel version 4.4 or later is required to support running AWS IoT Greengrass with [containers](#).
 - Linux kernel version 3.17 or later is required to support running AWS IoT Greengrass without containers. In this configuration, the default Lambda function containerization for the Greengrass group must be set to **No container**. For instructions, see [the section called “Setting default containerization for Lambda functions in a group”](#).
- [GNU C Library](#) (glibc) version 2.14 or later. OpenWrt distributions require [musl C Library](#) version 1.1.16 or later.

- The `/var/run` directory must be present on the device.
- The `/dev/stdin`, `/dev/stdout`, and `/dev/stderr` files must be available.
- Hardlink and softlink protection must be enabled on the device. Otherwise, AWS IoT Greengrass can only be run in insecure mode, using the `-i` flag.
- The following Linux kernel configurations must be enabled on the device:
 - Namespace:
 - `CONFIG_IPC_NS`
 - `CONFIG_UTS_NS`
 - `CONFIG_USER_NS`
 - `CONFIG_PID_NS`
 - Cgroups:
 - `CONFIG_CGROUP_DEVICE`
 - `CONFIG_CGROUPS`
 - `CONFIG_MEMCG`

The kernel must support [cgroups](#). The following requirements apply when running AWS IoT Greengrass with [containers](#):

- The *memory* cgroup must be enabled and mounted to allow AWS IoT Greengrass to set the memory limit for Lambda functions.
- The *devices* cgroup must be enabled and mounted if Lambda functions with [local resource access](#) are used to open files on the AWS IoT Greengrass core device.
- Others:
 - `CONFIG_POSIX_MQUEUE`
 - `CONFIG_OVERLAY_FS`
 - `CONFIG_HAVE_ARCH_SECCOMP_FILTER`
 - `CONFIG_SECCOMP_FILTER`
 - `CONFIG_KEYS`
 - `CONFIG_SECCOMP`
 - `CONFIG_SHMEM`
- The root certificate for Amazon S3 and AWS IoT must be present in the system trust store.
- [Stream manager](#) requires the Java 8 runtime and a minimum of 70 MB RAM in addition to the base AWS IoT Greengrass Core software memory requirement. Stream manager is enabled

by default when you use the **Default Group creation** option on the AWS IoT console. Stream manager is not supported on OpenWrt distributions.

- Libraries that support the [AWS Lambda runtime](#) required by the Lambda functions you want to run locally. Required libraries must be installed on the core and added to the PATH environment variable. Multiple libraries can be installed on the same core.
 - [Python](#) version 3.8 for functions that use the Python 3.8 runtime.
 - [Python](#) version 3.7 for functions that use the Python 3.7 runtime.
 - [Python](#) version 2.7 for functions that use the Python 2.7 runtime.
 - [Node.js](#) version 12.x for functions that use the Node.js 12.x runtime.
 - [Java](#) version 8 or later for functions that use the Java 8 runtime.

 Note

Running Java on an OpenWrt distribution isn't officially supported. However, if your OpenWrt build has Java support, you might be able to run Lambda functions authored in Java on your OpenWrt devices.

For more information about AWS IoT Greengrass support for Lambda runtimes, see [Run local Lambda functions](#).

- The following shell commands (not the BusyBox variants) are required by the [over-the-air \(OTA\) update agent](#):
 - `wget`
 - `realpath`
 - `tar`
 - `readlink`
 - `basename`
 - `dirname`
 - `pidof`
 - `df`
 - `grep`
 - `umount`

- `gzip`
- `mkdir`
- `rm`
- `ln`
- `cut`
- `cat`
- `/bin/bash`

GGC v1.10

Supported platforms:

- Architecture: Armv7l
 - OS: Linux
 - OS: Linux ([OpenWrt](#))
- Architecture: Armv8 (AArch64)
 - OS: Linux
 - OS: Linux ([OpenWrt](#))
- Architecture: Armv6l
 - OS: Linux
- Architecture: x86_64
 - OS: Linux
- Windows, macOS, and Linux platforms can run AWS IoT Greengrass in a Docker container. For more information, see [the section called “Run AWS IoT Greengrass in a Docker container”](#).

Requirements:

- Minimum 128 MB disk space available for the AWS IoT Greengrass Core software. If you use the [OTA update agent](#), the minimum is 400 MB.
- Minimum 128 MB RAM allocated to the AWS IoT Greengrass Core software. With [stream manager](#) enabled, the minimum is 198 MB RAM.

 Note

Stream manager is enabled by default if you use the **Default Group creation** option on the AWS IoT console to create your Greengrass group.

- Linux kernel version:
 - Linux kernel version 4.4 or later is required to support running AWS IoT Greengrass with [containers](#).
 - Linux kernel version 3.17 or later is required to support running AWS IoT Greengrass without containers. In this configuration, the default Lambda function containerization for the Greengrass group must be set to **No container**. For instructions, see [the section called "Setting default containerization for Lambda functions in a group"](#).
- [GNU C Library](#) (glibc) version 2.14 or later. OpenWrt distributions require [musl C Library](#) version 1.1.16 or later.
- The `/var/run` directory must be present on the device.
- The `/dev/stdin`, `/dev/stdout`, and `/dev/stderr` files must be available.
- Hardlink and softlink protection must be enabled on the device. Otherwise, AWS IoT Greengrass can only be run in insecure mode, using the `-i` flag.
- The following Linux kernel configurations must be enabled on the device:
 - Namespace:
 - `CONFIG_IPC_NS`
 - `CONFIG_UTS_NS`
 - `CONFIG_USER_NS`
 - `CONFIG_PID_NS`
 - Cgroups:
 - `CONFIG_CGROUP_DEVICE`
 - `CONFIG_CGROUPS`
 - `CONFIG_MEMCG`

The kernel must support [cgroups](#). The following requirements apply when running AWS IoT Greengrass with [containers](#):

- The *memory* cgroup must be enabled and mounted to allow AWS IoT Greengrass to set the memory limit for Lambda functions.

- The *devices* cgroup must be enabled and mounted if Lambda functions with [local resource access](#) are used to open files on the AWS IoT Greengrass core device.
- Others:
 - CONFIG_POSIX_MQUEUE
 - CONFIG_OVERLAY_FS
 - CONFIG_HAVE_ARCH_SECCOMP_FILTER
 - CONFIG_SECCOMP_FILTER
 - CONFIG_KEYS
 - CONFIG_SECCOMP
 - CONFIG_SHMEM
- The root certificate for Amazon S3 and AWS IoT must be present in the system trust store.
- [Stream manager](#) requires the Java 8 runtime and a minimum of 70 MB RAM in addition to the base AWS IoT Greengrass Core software memory requirement. Stream manager is enabled by default when you use the **Default Group creation** option on the AWS IoT console. Stream manager is not supported on OpenWrt distributions.
- Libraries that support the [AWS Lambda runtime](#) required by the Lambda functions you want to run locally. Required libraries must be installed on the core and added to the PATH environment variable. Multiple libraries can be installed on the same core.
 - [Python](#) version 3.7 for functions that use the Python 3.7 runtime.
 - [Python](#) version 2.7 for functions that use the Python 2.7 runtime.
 - [Node.js](#) version 12.x for functions that use the Node.js 12.x runtime.
 - [Java](#) version 8 or later for functions that use the Java 8 runtime.

 Note

Running Java on an OpenWrt distribution isn't officially supported. However, if your OpenWrt build has Java support, you might be able to run Lambda functions authored in Java on your OpenWrt devices.

For more information about AWS IoT Greengrass support for Lambda runtimes, see [Run local Lambda functions](#).

- The following shell commands (not the BusyBox variants) are required by the [over-the-air \(OTA\) update agent](#):
 - `wget`
 - `realpath`
 - `tar`
 - `readlink`
 - `basename`
 - `dirname`
 - `pidof`
 - `df`
 - `grep`
 - `umount`
 - `mv`
 - `gzip`
 - `mkdir`
 - `rm`
 - `ln`
 - `cut`
 - `cat`
 - `/bin/bash`

GGC v1.9

Supported platforms:

- Architecture: Armv7l
 - OS: Linux
 - OS: Linux ([OpenWrt](#))
- Architecture: Armv8 (AArch64)
 - OS: Linux
 - OS: Linux ([OpenWrt](#))

- OS: Linux
- Architecture: x86_64
- OS: Linux
- Windows, macOS, and Linux platforms can run AWS IoT Greengrass in a Docker container. For more information, see [the section called “Run AWS IoT Greengrass in a Docker container”](#).

Requirements:

- Minimum 128 MB disk space available for the AWS IoT Greengrass Core software. If you use the [OTA update agent](#), the minimum is 400 MB.
- Minimum 128 MB RAM allocated to the AWS IoT Greengrass Core software.
- Linux kernel version:
 - Linux kernel version 4.4 or later is required to support running AWS IoT Greengrass with [containers](#).
 - Linux kernel version 3.17 or later is required to support running AWS IoT Greengrass without containers. In this configuration, the default Lambda function containerization for the Greengrass group must be set to **No container**. For instructions, see [the section called “Setting default containerization for Lambda functions in a group”](#).
- [GNU C Library](#) (glibc) version 2.14 or later. OpenWrt distributions require [musl C Library](#) version 1.1.16 or later.
- The `/var/run` directory must be present on the device.
- The `/dev/stdin`, `/dev/stdout`, and `/dev/stderr` files must be available.
- Hardlink and softlink protection must be enabled on the device. Otherwise, AWS IoT Greengrass can only be run in insecure mode, using the `-i` flag.
- The following Linux kernel configurations must be enabled on the device:
 - Namespace:
 - `CONFIG_IPC_NS`
 - `CONFIG_UTS_NS`
 - `CONFIG_USER_NS`
 - `CONFIG_PID_NS`
 - Cgroups:
 - `CONFIG_CGROUP_DEVICE`

- CONFIG_CGROUPS
- CONFIG_MEMCG

The kernel must support [cgroups](#). The following requirements apply when running AWS IoT Greengrass with [containers](#):

- The *memory* cgroup must be enabled and mounted to allow AWS IoT Greengrass to set the memory limit for Lambda functions.
- The *devices* cgroup must be enabled and mounted if Lambda functions with [local resource access](#) are used to open files on the AWS IoT Greengrass core device.
- Others:
 - CONFIG_POSIX_MQUEUE
 - CONFIG_OVERLAY_FS
 - CONFIG_HAVE_ARCH_SECCOMP_FILTER
 - CONFIG_SECCOMP_FILTER
 - CONFIG_KEYS
 - CONFIG_SECCOMP
 - CONFIG_SHMEM
- The root certificate for Amazon S3 and AWS IoT must be present in the system trust store.
- Libraries that support the [AWS Lambda runtime](#) required by the Lambda functions you want to run locally. Required libraries must be installed on the core and added to the PATH environment variable. Multiple libraries can be installed on the same core.
 - [Python](#) version 2.7 for functions that use the Python 2.7 runtime.
 - [Python](#) version 3.7 for functions that use the Python 3.7 runtime.
 - [Node.js](#) version 6.10 or later for functions that use the Node.js 6.10 runtime.
 - [Node.js](#) version 8.10 or later for functions that use the Node.js 8.10 runtime.
 - [Java](#) version 8 or later for functions that use the Java 8 runtime.

 Note

Running Java on an OpenWrt distribution isn't officially supported. However, if your OpenWrt build has Java support, you might be able to run Lambda functions authored in Java on your OpenWrt devices.

For more information about AWS IoT Greengrass support for Lambda runtimes, see [Run local Lambda functions](#).

- The following shell commands (not the BusyBox variants) are required by the [over-the-air \(OTA\) update agent](#):
 - `wget`
 - `realpath`
 - `tar`
 - `readlink`
 - `basename`
 - `dirname`
 - `pidof`
 - `df`
 - `grep`
 - `umount`
 - `mv`
 - `gzip`
 - `mkdir`
 - `rm`
 - `ln`
 - `cut`
 - `cat`

GGC v1.8

- Supported platforms:
 - Architecture: Armv7l; OS: Linux
 - Architecture: x86_64; OS: Linux
 - Architecture: Armv8 (AArch64); OS: Linux
 - Windows, macOS, and Linux platforms can run AWS IoT Greengrass in a Docker container. For more information, see [the section called “Run AWS IoT Greengrass in a Docker container”](#).

- Linux platforms can run a version of AWS IoT Greengrass with limited functionality using the Greengrass snap, which is available through [Snapcraft](#). For more information, see [the section called “AWS IoT Greengrass snap software”](#).
- The following items are required:
 - Minimum 128 MB disk space available for the AWS IoT Greengrass Core software. If you use the [OTA update agent](#), the minimum is 400 MB.
 - Minimum 128 MB RAM allocated to the AWS IoT Greengrass Core software.
 - Linux kernel version:
 - Linux kernel version 4.4 or later is required to support running AWS IoT Greengrass with [containers](#).
 - Linux kernel version 3.17 or later is required to support running AWS IoT Greengrass without containers. In this configuration, the default Lambda function containerization for the Greengrass group must be set to **No container**. For instructions, see [the section called “Setting default containerization for Lambda functions in a group”](#).
 - [GNU C Library](#) (glibc) version 2.14 or later.
 - The `/var/run` directory must be present on the device.
 - The `/dev/stdin`, `/dev/stdout`, and `/dev/stderr` files must be available.
 - Hardlink and softlink protection must be enabled on the device. Otherwise, AWS IoT Greengrass can only be run in insecure mode, using the `-i` flag.
- The following Linux kernel configurations must be enabled on the device:
 - Namespace:
 - `CONFIG_IPC_NS`
 - `CONFIG_UTS_NS`
 - `CONFIG_USER_NS`
 - `CONFIG_PID_NS`
 - Cgroups:
 - `CONFIG_CGROUP_DEVICE`
 - `CONFIG_CGROUPS`
 - `CONFIG_MEMCG`

The kernel must support [cgroups](#). The following requirements apply when running AWS IoT Greengrass with [containers](#):

- The *memory* cgroup must be enabled and mounted to allow AWS IoT Greengrass to set the memory limit for Lambda functions.
- The *devices* cgroup must be enabled and mounted if Lambda functions with [local resource access](#) are used to open files on the AWS IoT Greengrass core device.
- Others:
 - CONFIG_POSIX_MQUEUE
 - CONFIG_OVERLAY_FS
 - CONFIG_HAVE_ARCH_SECCOMP_FILTER
 - CONFIG_SECCOMP_FILTER
 - CONFIG_KEYS
 - CONFIG_SECCOMP
 - CONFIG_SHMEM
- The root certificate for Amazon S3 and AWS IoT must be present in the system trust store.
- The following items are conditionally required:
 - Libraries that support the [AWS Lambda runtime](#) required by the Lambda functions you want to run locally. Required libraries must be installed on the core and added to the PATH environment variable. Multiple libraries can be installed on the same core.
 - [Python](#) version 2.7 for functions that use the Python 2.7 runtime.
 - [Node.js](#) version 6.10 or later for functions that use the Node.js 6.10 runtime.
 - [Java](#) version 8 or later for functions that use the Java 8 runtime.
 - The following shell commands (not the BusyBox variants) are required by the [over-the-air \(OTA\) update agent](#):
 - wget
 - realpath
 - tar
 - readlink
 - basename
 - dirname
 - pidof
 - df
 - grep

- `umount`
- `mv`
- `gzip`
- `mkdir`
- `rm`
- `ln`
- `cut`
- `cat`

For information about AWS IoT Greengrass quotas (limits), see [Service Quotas](#) in the *Amazon Web Services General Reference*.

For pricing information, see [AWS IoT Greengrass pricing](#) and [AWS IoT Core pricing](#).

AWS IoT Greengrass downloads

You can use the following information to find and download software for use with AWS IoT Greengrass.

Topics

- [AWS IoT Greengrass Core software](#)
- [AWS IoT Greengrass snap software](#)
- [AWS IoT Greengrass Docker software](#)
- [AWS IoT Greengrass Core SDK](#)
- [Supported machine learning runtimes and libraries](#)
- [AWS IoT Greengrass ML SDK software](#)

AWS IoT Greengrass Core software

The AWS IoT Greengrass Core software extends AWS functionality onto an AWS IoT Greengrass core device, making it possible for local devices to act locally on the data they generate.

v1.11

1.11.6

Bug fixes and improvements:

- Improved resilience if sudden power loss occurs during a deployment.
- Fixed an issue where stream manager data corruption could prevent the AWS IoT Greengrass Core software from starting.
- Fixed an issue where new client devices couldn't connect to the core in certain scenarios.
- Fixed an issue where stream manager stream names couldn't contain `.log`.

1.11.5

Bug fixes and improvements:

- General performance improvements and bug fixes.

1.11.4

Bug fixes and improvements:

- Fixed an issue with stream manager that prevented upgrades to AWS IoT Greengrass Core software v1.11.3. If you are using stream manager to export data to the cloud, you can now use an OTA update to upgrade an earlier v1.x version of the AWS IoT Greengrass Core software to v1.11.4.
- General performance improvements and bug fixes.

1.11.3

Bug fixes and improvements:

- Fixed an issue that caused AWS IoT Greengrass Core software running in a snap on an Ubuntu device to stop responding after a sudden power loss to the device.
- Fixed an issue that caused delayed delivery of MQTT messages to long-lived Lambda functions.
- Fixed an issue that caused MQTT messages to not be sent correctly when the `maxWorkItemCount` value was set to a value greater than 1024.
- Fixed an issue that caused the OTA update agent to ignore the MQTT KeepAlive period specified in the `keepAlive` property in [config.json](#).
- General performance improvements and bug fixes.

⚠ Important

If you are using stream manager to export data to the cloud, do *not* upgrade to AWS IoT Greengrass Core software v1.11.3 from an earlier v1.x version. If you are enabling stream manager for the first time, we strongly recommend that you first install the latest version of the AWS IoT Greengrass Core software.

1.11.1

Bug fixes and improvements:

- Fixed an issue that caused increased memory use for stream manager.
- Fixed an issue that caused stream manager to reset the sequence number of the stream to 0 if the Greengrass core device was turned off for longer than the specified time-to-live (TTL) period of the stream data.
- Fixed an issue that prevented stream manager from correctly stopping retry attempts to export data to the AWS Cloud.

1.11.0

New features:

- A telemetry agent on the Greengrass core collects local telemetry data and publishes it to AWS Cloud. To retrieve the telemetry data for further processing, customers can create an Amazon EventBridge rule and subscribe to a target. For more information, see [Gathering system health telemetry data from AWS IoT Greengrass core devices](#).
- A local HTTP API returns a snapshot of the current state of local worker processes started by AWS IoT Greengrass. For more information, see [Calling the local health check API](#).
- A [stream manager](#) automatically exports data to Amazon S3 and AWS IoT SiteWise.

New [stream manager parameters](#) let you update existing streams and pause or resume data export.

- Support for running Python 3.8.x Lambda functions on the core.
- A new `ggDaemonPort` property in [config.json](#) that use to configure the Greengrass core IPC port number. The default port number is 8000.

A new system `ComponentAuthTimeout` property in [config.json](#) that you use to configure the timeout for Greengrass core IPC authentication. The default timeout is 5000 milliseconds.

- Increased the maximum number of AWS IoT devices per AWS IoT Greengrass group from 200 to 2500.

Increased the maximum number of subscriptions per group from 1000 to 10000.

For more information, see [AWS IoT Greengrass endpoints and quotas](#).

Bug fixes and improvements:

- General optimization that can reduce the memory utilization of the Greengrass service processes.
- A new runtime configuration parameter (`mountAllBlockDevices`) lets Greengrass use bind mounts to mount all block devices into a container after setting up the OverlayFS. This feature resolved an issue that caused Greengrass deployment failure if `/usr` isn't under the `/` hierarchy.
- Fixed an issue that caused AWS IoT Greengrass core failure if `/tmp` is a symlink.
- Fixed an issue to let the Greengrass deployment agent remove unused machine learning model artifacts from the `mlmodel_public` folder.
- General performance improvements and bug fixes.

To install the AWS IoT Greengrass Core software on your core device, download the package for your architecture and operating system (OS), and then follow the steps in the [Getting Started Guide](#).

 Tip

AWS IoT Greengrass also provides other options for installing the AWS IoT Greengrass Core software. For example, you can use [Greengrass device setup](#) to configure your environment and install the latest version of the AWS IoT Greengrass Core software. Or, on supported Debian platforms, you can use the [APT package manager](#) to install or upgrade the AWS IoT Greengrass Core software. For more information, see [the section called “Install the AWS IoT Greengrass Core software”](#).

Architecture	Operating system	Link
Armv8 (AArch64)	Linux	Download
Armv8 (AArch64)	Linux (OpenWrt)	Download
Armv7l	Linux	Download
Armv7l	Linux (OpenWrt)	Download
Armv6l	Linux	Download
x86_64	Linux	Download

Extended life versions

1.10.5

New features in v1.10:

- A stream manager that processes data streams locally and exports them to the AWS Cloud automatically. This feature requires Java 8 on the Greengrass core device. For more information, see [Manage data streams](#).
- A new Greengrass Docker application deployment connector that runs a Docker application on a core device. For more information, see [the section called “Docker application deployment”](#).
- A new IoT SiteWise connector that sends industrial device data from OPC-UA servers to asset properties in AWS IoT SiteWise. For more information, see [the section called “IoT SiteWise”](#).
- Lambda functions that run without containerization can access machine learning resources in the Greengrass group. For more information, see [the section called “Access machine learning resources”](#).
- Support for MQTT persistent sessions with AWS IoT. For more information, see [the section called “MQTT persistent sessions with AWS IoT Core”](#).
- Local MQTT traffic can travel over a port other than the default port 8883. For more information, see [the section called “MQTT port for local messaging”](#).

- New `queueFullPolicy` options in the [AWS IoT Greengrass Core SDK](#) for reliable message publishing from Lambda functions.
- Support for running Node.js 12.x Lambda functions on the core.

Bug fixes and improvements:

- Over-the-air (OTA) updates with hardware security integration can be configured with OpenSSL 1.1.
- [Stream manager](#) is more resilient to file data corruption.
- Fixed an issue that causes a `sysfs` mount failure on devices using Linux kernel 5.1 and later.
- A new `mqttOperationTimeout` property in [config.json](#) that you use to set the timeout for publish, subscribe, and unsubscribe operations in MQTT connections with AWS IoT Core.
- Fixed an issue that caused increased memory use for stream manager.
- A new `systemComponentAuthTimeout` property in [config.json](#) that you use to configure the timeout for Greengrass core IPC authentication. The default timeout is 5000 milliseconds.
- Fixed an issue that caused the OTA update agent to ignore the MQTT `KeepAlive` period specified in the `keepAlive` property in [config.json](#).
- Fixed an issue that caused MQTT messages to not be sent correctly when the `maxWorkItemCount` value was set to a value greater than 1024.
- Fixed an issue that caused delayed delivery of MQTT messages to long-lived Lambda functions.
- Fixed an issue that caused AWS IoT Greengrass Core software running in a snap on an Ubuntu device to stop responding after a sudden power loss to the device.
- General performance improvements and bug fixes.

To install the AWS IoT Greengrass Core software on your core device, download the package for your architecture and operating system (OS), and then follow the steps in the [Getting Started Guide](#).

Architecture	Operating system	Link
Armv8 (AArch64)	Linux	Download

Architecture	Operating system	Link
Armv8 (AArch64)	Linux (OpenWrt)	Download
Armv7l	Linux	Download
Armv7l	Linux (OpenWrt)	Download
Armv6l	Linux	Download
x86_64	Linux	Download

1.9.4

New features in v1.9:

- Support for Python 3.7 and Node.js 8.10 Lambda runtimes. Lambda functions that use Python 3.7 and Node.js 8.10 runtimes can now run on an AWS IoT Greengrass core. (AWS IoT Greengrass continues to support the Python 2.7 and Node.js 6.10 runtimes.)
- Optimized MQTT connections. The Greengrass core establishes fewer connections with the AWS IoT Core. This change can reduce operational costs for charges that are based on the number of connections.
- Elliptic Curve (EC) key for the local MQTT server. The local MQTT server supports EC keys in addition to RSA keys. (The MQTT server certificate has an SHA-256 RSA signature, regardless of the key type.) For more information, see [the section called “Security principals”](#).
- Support for [OpenWrt](#). AWS IoT Greengrass Core software v1.9.2 or later can be installed on OpenWrt distributions with Armv8 (AArch64) and Armv7l architectures. Currently, OpenWrt does not support ML inference.
- Support for Armv6l. AWS IoT Greengrass Core software v1.9.3 or later can be installed on Raspbian distributions on Armv6l architectures (for example, on Raspberry Pi Zero devices).
- OTA updates on port 443 with ALPN. Greengrass cores that use port 443 for MQTT traffic now support over-the-air (OTA) software updates. AWS IoT Greengrass uses the Application Layer Protocol Network (ALPN) TLS extension to enable these connections. For more information, see [OTA updates of AWS IoT Greengrass Core software](#) and [the section called “Connect on port 443 or through a network proxy”](#).

To install the AWS IoT Greengrass Core software on your core device, download the package for your architecture and operating system (OS), and then follow the steps in the [Getting Started Guide](#).

Architecture	Operating system	Link
Armv8 (AArch64)	Linux	Download
Armv8 (AArch64)	Linux (OpenWrt)	Download
Armv7l	Linux	Download
Armv7l	Linux (OpenWrt)	Download
Armv6l	Linux	Download
x86_64	Linux	Download

1.8.4

- New features:
 - Configurable default access identity for Lambda functions in the group. This group-level setting determines the default permissions that are used to run Lambda functions. You can set the user ID, group ID, or both. Individual Lambda functions can override the default access identity of their group. For more information, see [the section called “Setting the default access identity for Lambda functions in a group”](#).
 - HTTPS traffic over port 443. HTTPS communication can be configured to travel over port 443 instead of the default port 8443. This complements AWS IoT Greengrass support for the Application Layer Protocol Network (ALPN) TLS extension and allows all Greengrass messaging traffic—both MQTT and HTTPS—to use port 443. For more information, see [the section called “Connect on port 443 or through a network proxy”](#).
 - Predictably named client IDs for AWS IoT connections. This change enables support for AWS IoT Device Defender and [AWS IoT lifecycle events](#), so you can receive notifications for connect, disconnect, subscribe, and unsubscribe events. Predictable naming also makes it easier to create logic around connection IDs (for example, to create [subscribe policy](#) templates based on certificate attributes). For more information, see [the section called “Client IDs for MQTT connections with AWS IoT”](#).

Bug fixes and improvements:

- Fixed an issue with shadow synchronization and device certificate manager reconnection.
- General performance improvements and bug fixes.

To install the AWS IoT Greengrass Core software on your core device, download the package for your architecture and operating system (OS), and then follow the steps in the [Getting Started Guide](#).

Architecture	Operating system	Link
Armv8 (AArch64)	Linux	Download
Armv7l	Linux	Download
x86_64	Linux	Download

By downloading this software, you agree to the [Greengrass Core Software License Agreement](#).

For information about other options for installing the AWS IoT Greengrass Core software on your device, see [the section called “Install the AWS IoT Greengrass Core software”](#).

AWS IoT Greengrass snap software

AWS IoT Greengrass snap 1.11.x enables you to run a limited version of AWS IoT Greengrass through convenient software packages, along with all necessary dependencies, in a containerized environment.

Note

The AWS IoT Greengrass snap is available for AWS IoT Greengrass Core software v1.11.x. AWS IoT Greengrass doesn't provide a snap for v1.10.x. Unsupported versions don't receive bug fixes or updates.

The AWS IoT Greengrass snap doesn't support connectors and machine learning (ML) inference.

For more information, see [the section called “Run AWS IoT Greengrass in a snap”](#).

AWS IoT Greengrass Docker software

AWS provides a Dockerfile and Docker images that make it easier for you to run AWS IoT Greengrass in a Docker container.

Dockerfile

Dockerfiles contain source code for building custom AWS IoT Greengrass container images. Images can be modified to run on different platform architectures or to reduce the image size. For instructions, see the README file.

Download your target AWS IoT Greengrass Core software version.

v1.11

- [Dockerfile for AWS IoT Greengrass v1.11.6](#).

Extended life versions

v1.10

[Dockerfile for AWS IoT Greengrass v1.10.5](#).

v1.9

[Dockerfile for AWS IoT Greengrass v1.9.4](#).

v1.8

[Dockerfile for AWS IoT Greengrass v1.8.1](#).

Docker image

Docker images have the AWS IoT Greengrass Core software and dependencies installed on Amazon Linux 2 (x86_64) and Alpine Linux (x86_64, Armv7l, or AArch64) base images. You can use prebuilt images to start experimenting with AWS IoT Greengrass.

Important

On June 30, 2022, AWS IoT Greengrass ended maintenance for AWS IoT Greengrass Core software v1.x Docker images that are published to Amazon Elastic Container

Registry (Amazon ECR) and Docker Hub. You can continue to download these Docker images from Amazon ECR and Docker Hub until June 30, 2023, which is 1 year after maintenance ended. However, the AWS IoT Greengrass Core software v1.x Docker images no longer receive security patches or bug fixes after maintenance ended on June 30, 2022. If you run a production workload that depends on these Docker images, we recommend that you build your own Docker images using the Dockerfiles that AWS IoT Greengrass provides. For more information, see [AWS IoT Greengrass Version 1 maintenance policy](#).

Download a prebuilt image from [Docker Hub](#) or Amazon Elastic Container Registry (Amazon ECR).

- For Docker Hub, use the *version* tag to download a specific version of the Greengrass Docker image. To find tags for all available images, check the **Tags** page on Docker Hub.
- For Amazon ECR, use the `latest` tag to download the latest available version of the Greengrass Docker image. For more information about listing available image versions and downloading images from Amazon ECR, see [Running AWS IoT Greengrass in a Docker container](#).

 Warning

Starting with v1.11.6 of the AWS IoT Greengrass Core software, the Greengrass Docker images no longer include Python 2.7, because Python 2.7 reached end-of-life in 2020 and no longer receives security updates. If you choose to update to these Docker images, we recommend that you validate that your applications work with the new Docker images before you deploy the updates to production devices. If you require Python 2.7 for your application that uses a Greengrass Docker image, you can modify the Greengrass Dockerfile to include Python 2.7 for your application.

AWS IoT Greengrass doesn't provide Docker images for AWS IoT Greengrass Core software v1.11.1.

Note

By default, `alpine-aarch64` and `alpine-armv7l` images can run only on Arm-based hosts. To run these images on an x86 host, you can install [QEMU](#) and mount the QEMU libraries on the host. For example:

```
docker run --rm --privileged multiarch/qemu-user-static --reset -p yes
```

AWS IoT Greengrass Core SDK

Lambda functions use the AWS IoT Greengrass Core SDK to interact with the AWS IoT Greengrass core locally. This allows deployed Lambda functions to:

- Exchange MQTT messages with AWS IoT Core.
- Exchange MQTT messages with connectors, client devices, and other Lambda functions in the Greengrass group.
- Interact with the local shadow service.
- Invoke other local Lambda functions.
- Access [secret resources](#).
- Interact with [stream manager](#).

Download the AWS IoT Greengrass Core SDK for your language or platform from GitHub.

- [AWS IoT Greengrass Core SDK for Java](#)
- [AWS IoT Greengrass Core SDK for Node.js](#)
- [AWS IoT Greengrass Core SDK for Python](#)
- [AWS IoT Greengrass Core SDK for C](#)

For more information, see [AWS IoT Greengrass Core SDK](#).

Supported machine learning runtimes and libraries

To [perform inference](#) on a Greengrass core, you must install the machine learning runtime or library for your ML model type.

AWS IoT Greengrass supports the following ML model types. Use these links to find information about how to install the runtime or library for your model type and device platform.

- [Deep Learning Runtime \(DLR\)](#)
- [MXNet](#)
- [TensorFlow](#)

Machine learning samples

AWS IoT Greengrass provides samples that you can use with supported ML runtimes and libraries. These samples are released under the [Greengrass Core Software License Agreement](#).

Deep learning runtime (DLR)

Download the sample for your device platform:

- DLR sample for [Raspberry Pi](#)
- DLR sample for [NVIDIA Jetson TX2](#)
- DLR sample for [Intel Atom](#)

For a tutorial that uses the DLR sample, see [the section called “How to configure optimized machine learning inference”](#).

MXNet

Download the sample for your device platform:

- MXNet sample for [Raspberry Pi](#)
- MXNet sample for [NVIDIA Jetson TX2](#)
- MXNet sample for [Intel Atom](#)

For a tutorial that uses the MXNet sample, see [the section called “How to configure machine learning inference”](#).

TensorFlow

Download the [Tensorflow sample](#) for your device platform. This sample works with Raspberry Pi, NVIDIA Jetson TX2, and Intel Atom.

AWS IoT Greengrass ML SDK software

The [AWS IoT Greengrass Machine Learning SDK](#) enables the Lambda functions you author to consume a local machine learning model and send data to the [ML Feedback](#) connector for uploading and publishing.

v1.1.0

- [Python 3.7.](#)

v1.0.0

- [Python 2.7.](#)

We want to hear from you

We welcome your feedback. To contact us, visit [AWS re:Post](#) and use the [AWS IoT Greengrass tag](#).

Install the AWS IoT Greengrass Core software

The AWS IoT Greengrass Core software extends AWS functionality onto an AWS IoT Greengrass core device, making it possible for local devices to act locally on the data they generate.

AWS IoT Greengrass provides several options for installing the AWS IoT Greengrass Core software:

- [Download and extract a tar.gz file.](#)
- [Run the Greengrass Device Setup script.](#)
- [Install from an APT repository.](#)

AWS IoT Greengrass also provides containerized environments that run the AWS IoT Greengrass Core software.

- [Run AWS IoT Greengrass in a Docker container.](#)
- [Run AWS IoT Greengrass in a snap.](#)

Download and extract the AWS IoT Greengrass Core software package

Choose the AWS IoT Greengrass Core software for your platform to download as a tar.gz file and extract on your device. You can download recent versions of the software. For more information, see [the section called “AWS IoT Greengrass Core software”](#).

Run the Greengrass device setup script

Run Greengrass device setup to configure your device, install the latest AWS IoT Greengrass Core software version, and deploy a Hello World Lambda function in minutes. For more information, see [the section called “Quick start: Greengrass device setup”](#).

Install the AWS IoT Greengrass Core software from an APT repository

Important

As of February 11, 2022, you can no longer install or update the AWS IoT Greengrass Core software from an APT repository. On devices where you added the AWS IoT Greengrass repository, you must [remove the repository from the sources list](#). Devices that run the software from the APT repository will continue to operate normally. We recommend that you update the AWS IoT Greengrass Core software using [tar files](#).

The APT repository provided by AWS IoT Greengrass includes the following packages:

- `aws-iot-greengrass-core`. Installs the AWS IoT Greengrass Core software.
- `aws-iot-greengrass-keyring`. Installs the GnuPG (GPG) keys used to sign the AWS IoT Greengrass package repository.

By downloading this software, you agree to the [Greengrass Core Software License Agreement](#).

Topics

- [Use systemd scripts to manage the Greengrass daemon lifecycle](#)
- [Uninstall the AWS IoT Greengrass core software using the APT repository](#)
- [Remove the AWS IoT Greengrass core software repository sources](#)

Use systemd scripts to manage the Greengrass daemon lifecycle

The `aws-iot-greengrass-core` package also installs systemd scripts that you can use to manage the AWS IoT Greengrass Core software (daemon) lifecycle.

- To start the Greengrass daemon during boot:

```
systemctl enable greengrass.service
```

- To start the Greengrass daemon:

```
systemctl start greengrass.service
```

- To stop the Greengrass daemon:

```
systemctl stop greengrass.service
```

- To check the status of the Greengrass daemon:

```
systemctl status greengrass.service
```

Uninstall the AWS IoT Greengrass core software using the APT repository

When you uninstall the AWS IoT Greengrass core software, you can choose whether to preserve or remove the AWS IoT Greengrass core software's configuration information, such as device certificates, group information, and log files.

To uninstall the AWS IoT Greengrass core software and preserve configuration information

- Run the following command to remove the AWS IoT Greengrass core software packages and preserve configuration information in the `/greengrass` folder.

```
sudo apt remove aws-iot-greengrass-core aws-iot-greengrass-keyring
```


To uninstall the AWS IoT Greengrass core software and remove configuration information

1. Run the following command to remove the AWS IoT Greengrass core software packages and remove configuration information from the `/greengrass` folder.

```
sudo apt purge aws-iot-greengrass-core aws-iot-greengrass-keyring
```

2. Remove the AWS IoT Greengrass core software repository from your sources list. For more information, see [Remove the AWS IoT Greengrass core software repository sources](#).

Remove the AWS IoT Greengrass core software repository sources

You can remove the AWS IoT Greengrass core software repository sources when you no longer need to install or update the AWS IoT Greengrass core software from the APT repository. After February 11, 2022, you must remove the repository from your sources list to avoid an error when you run `apt update`.

To remove the APT repository from the sources list

- Run the following commands to remove the AWS IoT Greengrass core software repository from the sources list.

```
sudo rm /etc/apt/sources.list.d/greengrass.list  
sudo apt update
```

Run AWS IoT Greengrass in a Docker container

AWS IoT Greengrass provides a Dockerfile and Docker images that make it easier for you to run the AWS IoT Greengrass Core software in a Docker container. For more information, see [the section called "AWS IoT Greengrass Docker software"](#).

Note

You can also run a Docker application on a Greengrass core device. To do so, use the [Greengrass Docker application deployment connector](#).

Run AWS IoT Greengrass in a snap

AWS IoT Greengrass snap 1.11.x enables you to run a limited version of AWS IoT Greengrass through convenient software packages, along with all necessary dependencies, in a containerized environment.

On December 31, 2023, AWS IoT Greengrass will end maintenance for the AWS IoT Greengrass core software version 1.11.x Snap that is published on snapcraft.io. Devices currently running the Snap will continue to work until further notice. However, the AWS IoT Greengrass core Snap will no longer receive security patches or bug fixes after maintenance ends.

Snap concepts

The following are essential snap concepts to help you understand how to use the AWS IoT Greengrass snap:

Channel

A snap component that defines which version of a snap is installed and tracked for updates. Snaps are automatically updated to the latest version of the current channel.

Interface

A snap component that grants access to resources, such as networks and user files.

To run the AWS IoT Greengrass snap, the following interfaces must be connected. Note that `greengrass-support-no-container` must be connected first and never disconnected.

- **greengrass-support-no-container**
- hardware-observe
- home-for-hooks
- hugepages-control
- log-observe
- mount-observe
- network
- network-bind
- network-control
- process-control
- system-observe

The other interfaces are optional. If your Lambda functions require access to specific resources, you might need to connect to the appropriate interfaces.

[Refresh](#)

Snapd are automatically updated. The snapd daemon is the snap package manager that checks for updates four times a day by default. Each update check is called a refresh. When a refresh occurs, the daemon stops, the snap gets updated, and then the daemon restarts.

For more information, see the [Snapcraft](#) website.

What's new with AWS IoT Greengrass snap v1.11.x

The following describes what's new and changed with the version 1.11.x of the AWS IoT Greengrass snap.

- This version supports only the snap_daemon user, exposed as user ID (UID) and group (GID) 584788.
- This version supports only noncontainerized Lambda functions.

Important

Because noncontainerized Lambda functions must share the same user (snap_daemon), the Lambda functions have no isolation from each other. For more information, see [Controlling execution of Greengrass Lambda functions by using group-specific configuration](#).

- This version supports C, C++, Java 8, Node.js 12.x, Python 2.7, Python 3.7, and Python 3.8 runtimes.

Note

To avoid redundant Python runtimes, Python 3.7 Lambda functions actually run the Python 3.8 runtime.

Getting started with AWS IoT Greengrass snap

The following procedure helps you install and configure the AWS IoT Greengrass snap on your device.

Requirements

To run the AWS IoT Greengrass snap, you must do the following:

- Run the AWS IoT Greengrass snap on a supported Linux distribution, such as Ubuntu, Linux Mint, Debian, and Fedora.
- Install the snapd daemon on your device. The snapd daemon including the snap tool manages the snap environment on your device.

For the list of supported Linux distributions and installation instructions, see [Installing snapd](#) in the *Snap documentation*.

Install and configure the AWS IoT Greengrass snap

The following tutorial shows you how to install and configure the AWS IoT Greengrass snap on your device.

Note

- Although this tutorial uses an Amazon EC2 instance (x86 t2.micro Ubuntu 20.04), you can run the AWS IoT Greengrass snap with physical hardware, such as a Raspberry Pi.
- The snapd daemon is preinstalled on Ubuntu.

1. Install the core18 snap by running the following command in your device's terminal:

```
sudo snap install core18
```

The core18 snap is a [base snap](#) that provides a runtime environment with commonly used libraries. This snap is built from [Ubuntu 18.04 LTS](#).

2. Upgrade snapd by running the following command:

```
sudo snap install --channel=edge snapd; sudo snap refresh --channel=edge snapd
```

3. Run the `snap list` command to check if you have the AWS IoT Greengrass snap installed.

The following example response shows that snapd is installed, but `aws-iot-greengrass` isn't.

Name	Version	Rev	Tracking	Publisher	Notes
amazon-ssm-agent	3.0.161.0	2996	latest/stable/...	aws#	classic
core	16-2.48	10444	latest/stable	canonical#	core
core18	20200929	1932	latest/stable	canonical#	base
lxd	4.0.4	18150	4.0/stable/...	canonical#	-
snapd	2.48+git548.g929ccfb	10526	latest/edge	canonical#	snapd

4. Choose one of the following options to install AWS IoT Greengrass snap 1.11.x.

- To install the AWS IoT Greengrass snap, run the following command:

```
sudo snap install aws-iot-greengrass
```

Example response:

```
aws-iot-greengrass 1.11.5 from Amazon Web Services (aws) installed
```

- To migrate from an earlier version to v1.11.x or update to the latest available patch version, run the following command:

```
sudo snap refresh --channel=1.11.x aws-iot-greengrass
```

Like other snaps, the AWS IoT Greengrass snap uses channels to manage minor versions. Snaps are automatically updated to the latest available version of the current channel. For examples, if you specify `--channel=1.11.x`, your AWS IoT Greengrass snap is updated to v1.11.5.

You can run the `snap info aws-iot-greengrass` command to get the list of available channels for AWS IoT Greengrass.

Example response:

```
name:      aws-iot-greengrass
summary:   AWS supported software that extends cloud capabilities to local devices.
publisher: Amazon Web Services (aws#)
store-url: https://snapcraft.io/aws-iot-greengrass
contact:   https://repost.aws/tags/TA4ckIed1sR4enZBey29rKTg/aws-io-t-greengrass
license:   Proprietary
description: |
```

AWS IoT Greengrass seamlessly extends AWS onto edge devices so they can act locally on the data they generate, while still using the cloud for management, analytics, and durable storage.

AWS IoT Greengrass snap v1.11.0 enables you to run a limited version of AWS IoT Greengrass with all necessary dependencies in a containerized environment.

The AWS IoT Greengrass snap doesn't support connectors and machine learning (ML) inference.

By downloading this software you agree to the Greengrass Core Software License Agreement

(<https://s3-us-west-2.amazonaws.com/greengrass-release-license/greengrass-license-v1.pdf>).

For more information, see *Run AWS IoT Greengrass in a snap*

(<https://docs.aws.amazon.com/greengrass/latest/developerguide/install-ggc.html#gg-snap-support>) in

the AWS IoT Greengrass Developer.

If you need help, try the AWS IoT Greengrass tag on AWS re:Post

(<https://repost.aws/tags/TA4ckIed1sR4enZBey29rKTg/aws-io-t-greengrass>) or connect with an AWS IQ expert

(<https://iq.aws.amazon.com/services/aws/greengrass>).

snap-id: SRDuhPJGj4XPxFNNZQK0TvURAp0wxKnd

channels:

```
latest/stable: 1.11.3 2021-06-15 (59) 111MB -
latest/candidate: 1.11.3 2021-06-14 (59) 111MB -
latest/beta: 1.11.3 2021-06-14 (59) 111MB -
latest/edge: 1.11.3 2021-06-14 (59) 111MB -
1.11.x/stable: 1.11.3 2021-06-15 (59) 111MB -
1.11.x/candidate: 1.11.3 2021-06-15 (59) 111MB -
1.11.x/beta: 1.11.3 2021-06-15 (59) 111MB -
1.11.x/edge: 1.11.3 2021-06-15 (59) 111MB -
```

5. To access specific resources that your Lambda functions need, you can connect to additional interfaces.

Run the following command to get the list of AWS IoT Greengrass snap supported interfaces:

```
snap connections aws-iot-greengrass
```

Example response:

Interface	Plug	Slot
Notes		

camera		aws-iot-greengrass:camera	-
-			
dvb		aws-iot-greengrass:dvb	-
-			
gpio		aws-iot-greengrass:gpio	-
-			
gpio-memory-control		aws-iot-greengrass:gpio-memory-control	-
-			
greengrass-support		aws-iot-greengrass:greengrass-support-no-container	
:greengrass-support	-		
hardware-observe		aws-iot-greengrass:hardware-observe	
:hardware-observe	manual		
hardware-random-control		aws-iot-greengrass:hardware-random-control	-
-			
home		aws-iot-greengrass:home-for-greengrassd	-
-			
home		aws-iot-greengrass:home-for-hooks	:home
manual			
hugepages-control		aws-iot-greengrass:hugepages-control	
:hugepages-control	manual		
i2c		aws-iot-greengrass:i2c	-
-			
iio		aws-iot-greengrass:iio	-
-			
joystick		aws-iot-greengrass:joystick	-
-			
log-observe		aws-iot-greengrass:log-observe	:log-
observe	manual		
mount-observe		aws-iot-greengrass:mount-observe	
:mount-observe	manual		
network		aws-iot-greengrass:network	
:network	-		
network-bind		aws-iot-greengrass:network-bind	
:network-bind	-		
network-control		aws-iot-greengrass:network-control	
:network-control	-		
opengl		aws-iot-greengrass:opengl	
:opengl	-		
optical-drive		aws-iot-greengrass:optical-drive	
:optical-drive	-		
process-control		aws-iot-greengrass:process-control	
:process-control	-		
raw-usb		aws-iot-greengrass:raw-usb	-
-			

removable-media	aws-iot-greengrass:removable-media	-
-		
serial-port	aws-iot-greengrass:serial-port	-
-		
spi	aws-iot-greengrass:spi	-
-		
system-observe	aws-iot-greengrass:system-observe	
:system-observe	-	

If you see a hyphen (-) in the Slot column, the corresponding interface isn't connected.

- Follow [Installing the AWS IoT Greengrass Core software](#) to create an AWS IoT thing, a Greengrass group, security resources that enable secure communications with AWS IoT, and the AWS IoT Greengrass Core software configuration file. The configuration file, `config.json`, contains configuration specific to your Greengrass core, such as the location of certificate files and the AWS IoT device data endpoint.

Note

If you downloaded the file to a different device, follow this [step](#) to transfer the files to the AWS IoT Greengrass core device.

- For the AWS IoT Greengrass snap, make sure that you update the [config.json](#) file, as shown in the following:
 - Replace each instance of *certificateId* with the certificate ID in the name of the certificate and key files.
 - If you downloaded a different Amazon root CA certificate than Amazon Root CA 1, replace each instance of *AmazonRootCA1.pem* with the name of the Amazon root CA file.

```
{
  ...
  "crypto" : {
    "principals" : {
      "SecretsManager" : {
        "privateKeyPath" : "file:///snap/aws-iot-greengrass/current/greengrass/
certs/certificateId-private.pem.keyy"
      },
      "IoTCertificate" : {
```



```
    "privateKeyPath" : "file:///snap/aws-iot-greengrass/current/greengrass/
certs/certificateId-private.pem.key",
    "certificatePath" : "file:///snap/aws-iot-greengrass/current/greengrass/
certs/certificateId-certificate.pem.crt"
  },
  "caPath" : "file:///snap/aws-iot-greengrass/current/greengrass/
certs/AmazonRootCA1.pem"
},
"writeDirectory": "/var/snap/aws-iot-greengrass/current/ggc-write-directory",
"pidFileDirectory": "/var/snap/aws-iot-greengrass/current/pidFileDirectory"
}
```

8. Run the following command to add your AWS IoT Greengrass certificate and configuration files:

```
sudo snap set aws-iot-greengrass gg-certs=/home/ubuntu/my-certs
```

Deploying a Lambda function

This section shows you how to deploy a customer managed Lambda function on the AWS IoT Greengrass snap.

Important

AWS IoT Greengrass snap v1.11 only supports noncontainerized Lambda functions.

1. Run the following command to start the AWS IoT Greengrass daemon:

```
sudo snap start aws-iot-greengrass
```

Example response:

```
Started.
```

Note

If you get an error, you can use the `snap run` command for a detailed error message. For more troubleshooting information, see [error: cannot perform the following tasks: - Run service command "start" for services \["greengrassd"\] of snap "aws-iot-greengrass" \(\[start snap.aws-iot-greengrass.greengrassd.service\] failed with exit status 1: Job for snap.aws-iot-greengrass.greengrassd.service failed because the control process exited with error code. See "systemctl status snap.aws-iot-greengrass.greengrassd.service" and "journalctl -xe" for details.\)](#).

2. Run the following command to confirm that the daemon is running:

```
snap services aws-iot-greengrass.greengrassd
```

Example response:

Service	Startup	Current	Notes
aws-iot-greengrass.greengrassd	disabled	active	-

3. Follow [Module 3 \(part 1\): Lambda functions on AWS IoT Greengrass](#) to create and deploy a Hello World Lambda function. However, before you deploy the Lambda function, complete the next step.
4. Make sure that your Lambda function run as the `snap_daemon` user and in the no container mode. To update the settings of your Greengrass group, do the following in the AWS IoT Greengrass console:
 - a. Sign in to the AWS IoT Greengrass console.
 - b. In the AWS IoT console navigation pane, under **Manage**, expand **Greengrass devices**, and then choose **Groups (V1)**.
 - c. Under **Greengrass groups**, choose the target group.
 - d. On the group configuration page, in the navigation pane, choose the **Lambda functions** tab.
 - e. Under **Default Lambda function runtime environment**, choose **Edit**, and do the following:

- i. For **Default system user and group**, choose **Another user ID/group ID**, and then enter **584788** for both **System user ID (number)** and **System group ID (number)**.
- ii. For **Default Lambda function containerization**, choose **No container**.
- iii. Choose **Save**.

Stopping the AWS IoT Greengrass daemon

You can use the `snap stop` command to stop a service.

To stop the AWS IoT Greengrass daemon, run the following command:

```
sudo snap stop aws-iot-greengrass
```

The command should return `Stopped..`

To check if you successfully stopped the snap, run the following command:

```
snap services aws-iot-greengrass.greengrassd
```

Example response:

Service	Startup	Current	Notes
aws-iot-greengrass.greengrassd	disabled	inactive	-

Uninstalling the AWS IoT Greengrass snap

To uninstall the AWS IoT Greengrass snap, run the following command:

```
sudo snap remove aws-iot-greengrass
```

Example response:

```
aws-iot-greengrass removed
```

Troubleshooting the AWS IoT Greengrass snap

Use the following information to help troubleshoot issues with the AWS IoT Greengrass snap.

Got permission denied errors.

Solution: Permission denied errors are often because of missing interfaces. For the list of missing interfaces and detailed troubleshooting information, you can use the `snappy-debug` tool.

Run the following command to install the tool.

```
sudo snap install snappy-debug
```

Example response:

```
snappy-debug 0.36-snapd2.45.1 from Canonical# installed
```

Run the `sudo snappy-debug` command in a separate terminal session. The operation continues until a permission denied error occurs.

For example, if your Lambda function tries to read a file in the `$HOME` directory, you may get the following response:

```
INFO: Following '/var/log/syslog'. If have dropped messages, use:
INFO: $ sudo journalctl --output=short --follow --all | sudo snappy-debug
kernel.printk_ratelimit = 0
= AppArmor =
Time: Dec  6 04:48:26
Log: apparmor="DENIED" operation="mknod" profile="snap.aws-iot-greengrass.greengrassd"
     name="/home/ubuntu/my-file.txt" pid=12345 comm="touch" requested_mask="c"
     denied_mask="c" fsuid=0 ouid=0
File: /home/ubuntu/my-file.txt (write)
Suggestion:
* add 'home' to 'plugins'
```

This example shows that creating the `/home/ubuntu/my-file.txt` file caused the permission error. It also suggests that you add `home` to `plugins`. However, this suggestion is not applicable. The `home-for-greengrassd` and `home-for-hooks` plugins are only given the read-only access.

For more information, see [The snappy-debug snap](#) in the *Snap documentation*.

error: cannot perform the following tasks: - Run service command "start" for services ["greengrassd"] of snap "aws-iot-greengrass" ([start snap.aws-iot-greengrass.greengrassd.service] failed with exit status 1: Job for snap.aws-iot-greengrass.greengrassd.service failed because the control process exited with error code. See "systemctl status snap.aws-iot-greengrass.greengrassd.service" and "journalctl -xe" for details.)

Solution: You might see this error when the `snap start aws-iot-greengrass` command fails to start the AWS IoT Greengrass Core software.

For more troubleshooting information, run the following command:

```
sudo snap run aws-iot-greengrass.greengrassd
```

Example response:

```
Couldn't find /snap/aws-iot-greengrass/44/greengrass/config/config.json.
```

This examples shows that AWS IoT Greengrass couldn't find the `config.json` file. You might check the configuration and certificate files.

/var/snap/aws-iot-greengrass/current/ggc-write-directory/packages/1.11.5/rootfs/merged is not an absolute path or is a symlink.

Solution: The AWS IoT Greengrass snap supports only noncontainerized Lambda functions. Make sure that you run your Lambda functions in the no container mode. For more information, see [Considerations when choosing Lambda function containerization](#) in the *AWS IoT Greengrass Version 1 Developer Guide*.

The snapd daemon failed to restart after you ran the `sudo snap refresh snapd` command.

Solution: Follow steps 6 through 8 in [Install and configure the AWS IoT Greengrass snap](#) to add the AWS IoT Greengrass certificate and configuration files to the AWS IoT Greengrass snap.

Archive an AWS IoT Greengrass Core software installation

When you upgrade to a new version of the AWS IoT Greengrass Core software, you can archive the currently installed version. This preserves your current installation environment so you can test a new software version on the same hardware. This also makes it easy to roll back to your archived version for any reason.

To archive the current installation and install a new version

1. Download the [AWS IoT Greengrass Core software](#) installation package that you want to upgrade to.
2. Copy the package to the destination core device. For instructions that show how to transfer files, see this [step](#).

Note

You copy your current certificates, keys, and configuration file to the new installation later.

Run the commands in the following steps in your core device terminal.

3. Make sure that the Greengrass daemon is stopped on the core device.
 - a. To check whether the daemon is running:

```
ps aux | grep -E 'greengrass.*daemon'
```

If the output contains a root entry for `/greengrass/ggc/packages/ggc-version/bin/daemon`, then the daemon is running.

Note

This procedure is written with the assumption that the AWS IoT Greengrass Core software is installed in the `/greengrass` directory.

- b. To stop the daemon:

```
cd /greengrass/ggc/core/  
sudo ./greengrassd stop
```

4. Move the current Greengrass root directory to a different directory.

```
sudo mv /greengrass /greengrass_backup
```

5. Untar the new software on the core device. Replace the *os-architecture* and *version* placeholders in the command.

```
sudo tar -zxvf greengrass-os-architecture-version.tar.gz -C /
```

6. Copy the archived certificates, keys, and configuration file to the new installation.

```
sudo cp /greengrass_backup/certs/* /greengrass/certs  
sudo cp /greengrass_backup/config/* /greengrass/config
```

7. Start the daemon:

```
cd /greengrass/ggc/core/  
sudo ./greengrassd start
```

Now, you can make a group deployment to test the new installation. If something fails, you can restore the archived installation.

To restore the archived installation

1. Stop the daemon.
2. Delete the new `/greengrass` directory.
3. Move the `/greengrass_backup` directory back to `/greengrass`.
4. Start the daemon.

Configure the AWS IoT Greengrass core

An AWS IoT Greengrass core is an AWS IoT thing (device) that acts as a hub or gateway in edge environments. Like other AWS IoT devices, a core exists in the registry, has a device shadow, and uses a device certificate to authenticate with AWS IoT Core and AWS IoT Greengrass. The core device runs the AWS IoT Greengrass Core software, which enables it to manage local processes for Greengrass groups, such as communication, shadow sync, and token exchange.

The AWS IoT Greengrass Core software provides the following functionality:

- Deployment and the local running of connectors and Lambda functions.
- Process data streams locally with automatic exports to the AWS Cloud.
- MQTT messaging over the local network between devices, connectors, and Lambda functions using managed subscriptions.

- MQTT messaging between AWS IoT and devices, connectors, and Lambda functions using managed subscriptions.
- Secure connections between devices and the AWS Cloud using device authentication and authorization.
- Local shadow synchronization of devices. Shadows can be configured to sync with the AWS Cloud.
- Controlled access to local device and volume resources.
- Deployment of cloud-trained machine learning models for running local inference.
- Automatic IP address detection that enables devices to discover the Greengrass core device.
- Central deployment of new or updated group configuration. After the configuration data is downloaded, the core device is restarted automatically.
- Secure, over-the-air (OTA) software updates of user-defined Lambda functions.
- Secure, encrypted storage of local secrets and controlled access by connectors and Lambda functions.

AWS IoT Greengrass core configuration file

The configuration file for the AWS IoT Greengrass Core software is `config.json`. It is located in the `/greengrass-root/config` directory.

Note

`greengrass-root` represents the path where the AWS IoT Greengrass Core software is installed on your device. Typically, this is the `/greengrass` directory.

If you use the **Default Group creation** option from the AWS IoT Greengrass console, then the `config.json` file is deployed to the core device in a working state.

You can review the contents of this file by running the following command:

```
cat /greengrass-root/config/config.json
```

The following is an example `config.json` file. This is the version that's generated when you create the core from the AWS IoT Greengrass console.

GGC v1.11

```

{
  "coreThing": {
    "caPath": "root.ca.pem",
    "certPath": "hash.cert.pem",
    "keyPath": "hash.private.key",
    "thingArn": "arn:partition:iot:region:account-id:thing/core-thing-name",
    "iotHost": "host-prefix-ats.iot.region.amazonaws.com",
    "ggHost": "greengrass-ats.iot.region.amazonaws.com",
    "keepAlive": 600,
    "ggDaemonPort": 8000,
    "systemComponentAuthTimeout": 5000
  },
  "runtime": {
    "maxWorkItemCount": 1024,
    "maxConcurrentLimit": 25,
    "lruSize": 25,
    "mountAllBlockDevices": "no",
    "cgroup": {
      "useSystemd": "yes"
    }
  },
  "managedRespawn": false,
  "crypto": {
    "principals": {
      "SecretsManager": {
        "privateKeyPath": "file:///greengrass/certs/hash.private.key"
      },
      "IoTCertificate": {
        "privateKeyPath": "file:///greengrass/certs/hash.private.key",
        "certificatePath": "file:///greengrass/certs/hash.cert.pem"
      }
    },
    "caPath": "file:///greengrass/certs/root.ca.pem"
  },
  "writeDirectory": "/var/snap/aws-iot-greengrass/current/ggc-write-directory",
  "pidFileDirectory": "/var/snap/aws-iot-greengrass/current/pidFileDirectory"
}

```

The `config.json` file supports the following properties:

coreThing

Field	Description	Notes
caPath	The path to the AWS IoT root CA relative to the <i>/greengrass-root</i> /certs directory.	For backward compatibility with versions earlier than 1.7.0. This property is ignored when the crypto object is present.  Note Make sure that your endpoints correspond to your certificate type.
certPath	The path to the core device certificate relative to the <i>/greengrass-root</i> /certs directory.	For backward compatibility with versions earlier than 1.7.0. This property is ignored when the crypto object is present.
keyPath	The path to the core private key relative to <i>/greengrass-root</i> /certs directory.	For backward compatibility with versions earlier than 1.7.0. This property is ignored when the crypto object is present.
thingArn	The Amazon Resource Name (ARN) of the AWS IoT thing that represents the AWS IoT Greengrass core device.	Find the ARN for your core in the AWS IoT Greengrass console under Cores , or by running the aws greengrass get-core-definition-version CLI command.
iotHost	Your AWS IoT endpoint.	Find the endpoint in the AWS IoT console under

Field	Description	Notes
		Settings, or by running the aws iot describe-endpoint --endpoint-type iot:Data-ATS CLI command. This command returns the Amazon Trust Services (ATS) endpoint. For more information, see the Server authentication documentation.  Note Make sure that your endpoints correspond to your certificate type. Make sure that your endpoints correspond to your AWS Region.

Field	Description	Notes
ggHost	Your AWS IoT Greengrass endpoint.	This is your <code>iotHost</code> endpoint with the host prefix replaced by <i>greengrass</i> (for example, <code>greengrass-ats.iot.<i>region</i>.amazonaws.com</code>). Use the same AWS Region as <code>iotHost</code>.  Note Make sure that your endpoints correspond to your certificate type. Make sure that your endpoints correspond to your AWS Region.
iotMqttPort	Optional. The port number to use for MQTT communication with AWS IoT.	Valid values are 8883 or 443. The default value is 8883. For more information, see Connect on port 443 or through a network proxy .
iotHttpPort	Optional. The port number used to create HTTPS connections to AWS IoT.	Valid values are 8443 or 443. The default value is 8443. For more information, see Connect on port 443 or through a network proxy .

Field	Description	Notes
ggMqttPort	Optional. The port number to use for MQTT communication over the local network.	Valid values are 1024 through 65535. The default value is 8883. For more information, see the section called “MQTT port for local messaging” .
ggHttpPort	Optional. The port number used to create HTTPS connections to the AWS IoT Greengrass service.	Valid values are 8443 or 443. The default value is 8443. For more information, see Connect on port 443 or through a network proxy .
keepAlive	Optional. The MQTT KeepAlive period, in seconds.	Valid range is between 30 and 1200 seconds. The default value is 600.
networkProxy	Optional. An object that defines a proxy server to connect to.	The proxy server can be HTTP or HTTPS. For more information, see Connect on port 443 or through a network proxy .
mqttOperationTimeout	Optional. The amount of time (in seconds) to allow the Greengrass core to complete a publish, subscribe, or unsubscribe operation in MQTT connections to AWS IoT Core.	The default value is 5. The minimum value is 5.

Field	Description	Notes
<code>ggDaemonPort</code>	Optional. The Greengrass core IPC port number.	This property is available in AWS IoT Greengrass v1.11.0 or later. Valid values are between 1024 and 65535. The default value is 8000.
<code>systemComponentAuthTimeout</code>	Optional. The time (in milliseconds) to allow the Greengrass core IPC to complete authentication.	This property is available in AWS IoT Greengrass v1.11.0 or later. Valid values are between 500 and 5000. The default value is 5000.

runtime

Field	Description	Notes
<code>maxWorkItemCount</code>	Optional. The maximum number of work items that the Greengrass daemon can process at a time. Work items that exceed this limit are ignored. The work item queue is shared by system components, user-defined Lambda functions, and connectors.	The default value is 1024. The maximum value is limited by your device hardware. Increasing this value increases the memory that AWS IoT Greengrass uses. You can increase this value if you expect your core to receive heavy MQTT message traffic.
<code>maxConcurrentLimit</code>	Optional. The maximum number of concurrent unpinning Lambda workers	The default value is 25. The minimum value is defined by <code>lruSize</code> .

Field	Description	Notes
	that the Greengrass daemon can have. You can specify a different integer to override this parameter.	
<code>lruSize</code>	Optional. Defines the minimum value for <code>maxConcurrentLimit</code> .	The default value is 25.
<code>mountAllBlockDevices</code>	Optional. Enables AWS IoT Greengrass to use bind mounts to mount all block devices into a container after setting up the OverlayFS.	This property is available in AWS IoT Greengrass v1.11.0 or later. Valid values are yes and no. The default value is no. Set this value to yes if your <code>/usr</code> directory isn't under the <code>/</code> hierarchy.
<code>postStartHealthCheckTimeout</code>	Optional. The time (in milliseconds) after starting that the Greengrass daemon waits for the health check to finish.	The default timeout is 30 seconds (30000 ms).
<code>cgroup</code>		
<code>useSystemd</code>	Indicates whether your device uses systemd .	Valid values are yes or no. Run the <code>check_ggc_dependencies</code> script in Module 1 to see if your device uses <code>systemd</code> .
crypto		

The crypto contains properties that support private key storage on a hardware security module (HSM) through PKCS#11 and local secret storage. For more information, see [the section called “Security principals”](#), [the section called “Hardware security integration”](#), and [Deploy secrets to the core](#). Configurations for private key storage on HSMs or in the file system are supported.

Field	Description	Notes
caPath	The absolute path to the AWS IoT root CA.	Must be a file URI of the form: <code>file:///absolute/path/to/file</code> .
 Note Make sure that your endpoints correspond to your certificate type.		
PKCS11		
OpenSSLEngine	Optional. The absolute path to the OpenSSL engine .so file to enable PKCS#11 support on OpenSSL.	Must be a path to a file on the file system. This property is required if you're using the Greengrass OTA update agent with hardware security. For more information, see the section called “Configure OTA updates”.
P11Provider	The absolute path to the PKCS#11 implementation's libdl-loadable library.	Must be a path to a file on the file system.

Field	Description	Notes
slotLabel	The slot label that's used to identify the hardware module.	Must conform to PKCS#11 label specifications.
slotUserPin	The user PIN that's used to authenticate the Greengrass core to the module.	Must have sufficient permissions to perform C_Sign with the configured private keys.
principals		
IoTCertificate	The certificate and private key that the core uses to make requests to AWS IoT.	
IoTCertificate. privateKeyPath	The path to the core private key.	For file system storage, must be a file URI of the form: file:///absolute/path/to/file . For HSM storage, must be an RFC 7512 PKCS#11 path that specifies the object label.
IoTCertificate. certificatePath	The absolute path to the core device certificate.	Must be a file URI of the form: file:///absolute/path/to/file .
MQTTServerCertificate	Optional. The private key that the core uses in combination with the certificate to act as an MQTT server or gateway.	

Field	Description	Notes
MQTTServerCertificate.privateKeyPath	The path to the local MQTT server private key.	Use this value to specify your own private key for the local MQTT server. For file system storage, must be a file URI of the form: <code>file:///absolute/path/to/file</code>. For HSM storage, must be an RFC 7512 PKCS#11 path that specifies the object label. If this property is omitted, AWS IoT Greengrass rotates the key based your rotation settings. If specified, the customer is responsible for rotating the key.
SecretsManager	The private key that secures the data key used for encryption. For more information, see Deploy secrets to the core .	

Field	Description	Notes
SecretsManager.privateKeyPath	The path to the local secrets manager private key.	Only an RSA key is supported. For file system storage, must be a file URI of the form: <code>file:///absolute/path/to/file</code>. For HSM storage, must be an RFC 7512 PKCS#11 path that specifies the object label. The private key must be generated using the PKCS#1 v1.5 padding mechanism.

The following configuration properties are also supported:

Field	Description	Notes
mqttMaxConnectionRetryInterval	Optional. The maximum interval (in seconds) between MQTT connection retries if the connection is dropped.	Specify this value as an unsigned integer. The default is 60.
managedRespawn	Optional. Indicates that the OTA agent needs to run custom code before an update.	Valid values are true or false. For more information, see OTA updates of AWS IoT Greengrass Core software .
writeDirectory	Optional. The write directory where AWS IoT Greengrass creates all read/write resources.	For more information, see Configure a write directory for AWS IoT Greengrass .

Field	Description	Notes
pidFileDirectory	Optional. AWS IoT Greengrass stores its process ID (PID) under this directory.	The default value is <code>/var/run</code> .

Extended life versions

The following versions of the AWS IoT Greengrass Core software are in the [extended life phase](#). This information is included for reference purposes only.

GGC v1.10

```
{
  "coreThing" : {
    "caPath" : "root.ca.pem",
    "certPath" : "hash.cert.pem",
    "keyPath" : "hash.private.key",
    "thingArn" : "arn:partition:iot:region:account-id:thing/core-thing-name",
    "iotHost" : "host-prefix-ats.iot.region.amazonaws.com",
    "ggHost" : "greengrass-ats.iot.region.amazonaws.com",
    "keepAlive" : 600,
    "systemComponentAuthTimeout": 5000
  },
  "runtime" : {
    "maxWorkItemCount" : 1024,
    "maxConcurrentLimit" : 25,
    "lruSize": 25,
    "cgroup" : {
      "useSystemd" : "yes"
    }
  },
  "managedRespawn" : false,
  "crypto" : {
    "principals" : {
      "SecretsManager" : {
        "privateKeyPath" : "file:///greengrass/certs/hash.private.key"
      },
      "IoTCertificate" : {
        "privateKeyPath" : "file:///greengrass/certs/hash.private.key",
        "certificatePath" : "file:///greengrass/certs/hash.cert.pem"
      }
    }
  }
}
```

```

    }
  },
  "caPath" : "file:///greengrass/certs/root.ca.pem"
}

```

The `config.json` file supports the following properties:

coreThing

Field	Description	Notes
caPath	The path to the AWS IoT root CA relative to the <code>/greengrass-root /</code> certs directory.	For backward compatibility with versions earlier than 1.7.0. This property is ignored when the crypto object is present.
		 Note Make sure that your endpoints correspond to your certificate type.
certPath	The path to the core device certificate relative to the <code>/greengrass-root /</code> certs directory.	For backward compatibility with versions earlier than 1.7.0. This property is ignored when the crypto object is present.
keyPath	The path to the core private key relative to <code>/greengrass-root /</code> certs directory.	For backward compatibility with versions earlier than 1.7.0. This property is ignored when the crypto object is present.

Field	Description	Notes
thingArn	The Amazon Resource Name (ARN) of the AWS IoT thing that represents the AWS IoT Greengrass core device.	Find the ARN for your core in the AWS IoT Greengrass console under Cores , or by running the aws greengrass get-core-definition-version CLI command.
iotHost	Your AWS IoT endpoint.	Find the endpoint in the AWS IoT console under Settings, or by running the aws iot describe-endpoint --endpoint-type iot:Data-ATS CLI command. This command returns the Amazon Trust Services (ATS) endpoint. For more information, see the Server authentication documentation.  Note Make sure that your endpoints correspond to your certificate type. Make sure that your endpoints correspond to your AWS Region.

Field	Description	Notes
ggHost	Your AWS IoT Greengrass endpoint.	This is your <code>iotHost</code> endpoint with the host prefix replaced by <i>greengrass</i> (for example, <code>greengrass-ats.iot.<i>region</i>.amazonaws.com</code>). Use the same AWS Region as <code>iotHost</code>.  Note Make sure that your endpoints correspond to your certificate type. Make sure that your endpoints correspond to your AWS Region.
iotMqttPort	Optional. The port number to use for MQTT communication with AWS IoT.	Valid values are 8883 or 443. The default value is 8883. For more information, see Connect on port 443 or through a network proxy .
iotHttpPort	Optional. The port number used to create HTTPS connections to AWS IoT.	Valid values are 8443 or 443. The default value is 8443. For more information, see Connect on port 443 or through a network proxy .

Field	Description	Notes
<code>ggMqttPort</code>	Optional. The port number to use for MQTT communication over the local network.	Valid values are 1024 through 65535. The default value is 8883. For more information, see the section called “MQTT port for local messaging” .
<code>ggHttpPort</code>	Optional. The port number used to create HTTPS connections to the AWS IoT Greengrass service.	Valid values are 8443 or 443. The default value is 8443. For more information, see Connect on port 443 or through a network proxy .
<code>keepAlive</code>	Optional. The MQTT KeepAlive period, in seconds.	Valid range is between 30 and 1200 seconds. The default value is 600.
<code>networkProxy</code>	Optional. An object that defines a proxy server to connect to.	The proxy server can be HTTP or HTTPS. For more information, see Connect on port 443 or through a network proxy .
<code>mqttOperationTimeout</code>	Optional. The amount of time (in seconds) to allow the Greengrass core to complete a publish, subscribe, or unsubscribe operation in MQTT connections to AWS IoT Core.	This property is available starting in AWS IoT Greengrass v1.10.2. The default value is 5. The minimum value is 5.

runtime

Field	Description	Notes
<code>maxWorkItemCount</code>	Optional. The maximum number of work items that the Greengrass daemon can process at a time. Work items that exceed this limit are ignored. The work item queue is shared by system components, user-defined Lambda functions, and connectors.	The default value is 1024. The maximum value is limited by your device hardware. Increasing this value increases the memory that AWS IoT Greengrass uses. You can increase this value if you expect your core to receive heavy MQTT message traffic.
<code>maxConcurrentLimit</code>	Optional. The maximum number of concurrent unpinned Lambda workers that the Greengrass daemon can have. You can specify a different integer to override this parameter.	The default value is 25. The minimum value is defined by <code>lruSize</code>.
<code>lruSize</code>	Optional. Defines the minimum value for <code>maxConcurrentLimit</code>.	The default value is 25.
<code>postStartHealthCheckTimeout</code>	Optional. The time (in milliseconds) after starting that the Greengrass daemon waits for the health check to finish.	The default timeout is 30 seconds (30000 ms).
<code>cgroup</code>		
<code>useSystemd</code>	Indicates whether your device uses systemd.	Valid values are yes or no. Run the <code>check_ggc_dependencies</code> script

Field**Description****Notes**

in [Module 1](#) to see if your device uses systemd.

crypto

The `crypto` contains properties that support private key storage on a hardware security module (HSM) through PKCS#11 and local secret storage. For more information, see [the section called “Security principals”](#), [the section called “Hardware security integration”](#), and [Deploy secrets to the core](#). Configurations for private key storage on HSMs or in the file system are supported.

Field**Description****Notes**`caPath`

The absolute path to the AWS IoT root CA.

Must be a file URI of the form:
`file:///absolute/path/to/file` .

Note

Make sure that your [endpoints correspond to your certificate type](#).

PKCS11`OpenSSLEngine`

Optional. The absolute path to the OpenSSL engine .so file to enable PKCS#11 support on OpenSSL.

Must be a path to a file on the file system.

This property is required if you're using the Greengrass OTA update agent with hardware security. For more information, see [the section](#)

Field	Description	Notes
		called "Configure OTA updates" .
P11Provider	The absolute path to the PKCS#11 implementation's libdl-loadable library.	Must be a path to a file on the file system.
slotLabel	The slot label that's used to identify the hardware module.	Must conform to PKCS#11 label specifications.
slotUserPin	The user PIN that's used to authenticate the Greengrass core to the module.	Must have sufficient permissions to perform C_Sign with the configured private keys.
principals		
IoTCertificate	The certificate and private key that the core uses to make requests to AWS IoT.	
IoTCertificate.privateKeyPath	The path to the core private key.	For file system storage, must be a file URI of the form: <code>file:///absolute/path/to/file</code>. For HSM storage, must be an RFC 7512 PKCS#11 path that specifies the object label.
IoTCertificate.certificatePath	The absolute path to the core device certificate.	Must be a file URI of the form: <code>file:///absolute/path/to/file</code>.

Field	Description	Notes
MQTTServerCertificate	Optional. The private key that the core uses in combination with the certificate to act as an MQTT server or gateway.	
MQTTServerCertificate.privateKeyPath	The path to the local MQTT server private key.	Use this value to specify your own private key for the local MQTT server. For file system storage, must be a file URI of the form: <code>file:///absolute/path/to/file</code>. For HSM storage, must be an RFC 7512 PKCS#11 path that specifies the object label. If this property is omitted, AWS IoT Greengrass rotates the key based your rotation settings. If specified, the customer is responsible for rotating the key.
SecretsManager	The private key that secures the data key used for encryption. For more information, see Deploy secrets to the core .	

Field	Description	Notes
SecretsManager.privateKeyPath	The path to the local secrets manager private key.	Only an RSA key is supported. For file system storage, must be a file URI of the form: <code>file:///absolute/path/to/file</code>. For HSM storage, must be an RFC 7512 PKCS#11 path that specifies the object label. The private key must be generated using the PKCS#1 v1.5 padding mechanism.

The following configuration properties are also supported:

Field	Description	Notes
mqttMaxConnectionRetryInterval	Optional. The maximum interval (in seconds) between MQTT connection retries if the connection is dropped.	Specify this value as an unsigned integer. The default is 60.
managedRespawn	Optional. Indicates that the OTA agent needs to run custom code before an update.	Valid values are true or false. For more information, see OTA updates of AWS IoT Greengrass Core software .

Field	Description	Notes
writeDirectory	Optional. The write directory where AWS IoT Greengrass creates all read/write resources.	For more information, see Configure a write directory for AWS IoT Greengrass .

GGC v1.9

```
{
  "coreThing" : {
    "caPath" : "root.ca.pem",
    "certPath" : "hash.cert.pem",
    "keyPath" : "hash.private.key",
    "thingArn" : "arn:partition:iot:region:account-id:thing/core-thing-name",
    "iotHost" : "host-prefix-ats.iot.region.amazonaws.com",
    "ggHost" : "greengrass-ats.iot.region.amazonaws.com",
    "keepAlive" : 600
  },
  "runtime" : {
    "cgroup" : {
      "useSystemd" : "yes"
    }
  },
  "managedRespawn" : false,
  "crypto" : {
    "principals" : {
      "SecretsManager" : {
        "privateKeyPath" : "file:///greengrass/certs/hash.private.key"
      },
      "IoTCertificate" : {
        "privateKeyPath" : "file:///greengrass/certs/hash.private.key",
        "certificatePath" : "file:///greengrass/certs/hash.cert.pem"
      }
    },
    "caPath" : "file:///greengrass/certs/root.ca.pem"
  }
}
```

The `config.json` file supports the following properties:

coreThing

Field	Description	Notes
caPath	The path to the AWS IoT root CA relative to the <i>/greengrass-root /</i> certs directory.	For backward compatibility with versions earlier than 1.7.0. This property is ignored when the crypto object is present.  Note Make sure that your endpoints correspond to your certificate type.
certPath	The path to the core device certificate relative to the <i>/greengrass-root /</i> certs directory.	For backward compatibility with versions earlier than 1.7.0. This property is ignored when the crypto object is present.
keyPath	The path to the core private key relative to <i>/greengrass-root /</i> certs directory.	For backward compatibility with versions earlier than 1.7.0. This property is ignored when the crypto object is present.
thingArn	The Amazon Resource Name (ARN) of the AWS IoT thing that represents the AWS IoT Greengrass core device.	Find the ARN for your core in the AWS IoT Greengrass console under Cores , or by running the aws greengrass get-core-definition-version CLI command.

Field	Description	Notes
iotHost	Your AWS IoT endpoint.	Find the endpoint in the AWS IoT console under Settings, or by running the aws iot describe-endpoint --endpoint-t-type iot:Data-ATS CLI command. This command returns the Amazon Trust Services (ATS) endpoint. For more information, see the Server authentication documentation.  Note Make sure that your endpoints correspond to your certificate type. Make sure that your endpoints correspond to your AWS Region.

Field	Description	Notes
ggHost	Your AWS IoT Greengrass endpoint.	This is your <code>iotHost</code> endpoint with the host prefix replaced by <i>greengrass</i> (for example, <code>greengrass-ats.iot.<i>region</i>.amazonaws.com</code>). Use the same AWS Region as <code>iotHost</code>.  Note Make sure that your endpoints correspond to your certificate type. Make sure that your endpoints correspond to your AWS Region.
iotMqttPort	Optional. The port number to use for MQTT communication with AWS IoT.	Valid values are 8883 or 443. The default value is 8883. For more information, see Connect on port 443 or through a network proxy .
iotHttpPort	Optional. The port number used to create HTTPS connections to AWS IoT.	Valid values are 8443 or 443. The default value is 8443. For more information, see Connect on port 443 or through a network proxy .

Field	Description	Notes
<code>ggHttpPort</code>	Optional. The port number used to create HTTPS connections to the AWS IoT Greengrass service.	Valid values are 8443 or 443. The default value is 8443. For more information, see Connect on port 443 or through a network proxy .
<code>keepAlive</code>	Optional. The MQTT KeepAlive period, in seconds.	Valid range is between 30 and 1200 seconds. The default value is 600.
<code>networkProxy</code>	Optional. An object that defines a proxy server to connect to.	The proxy server can be HTTP or HTTPS. For more information, see Connect on port 443 or through a network proxy .

runtime

Field	Description	Notes
<code>maxConcurrentLimit</code>	Optional. The maximum number of concurrent unpinning Lambda workers that the Greengrass daemon can have. You can specify a different integer to override this parameter.	The default value is 25. The minimum value is defined by <code>lruSize</code> .
<code>lruSize</code>	Optional. Defines the minimum value for <code>maxConcurrentLimit</code> .	The default value is 25.
<code>postStartHealthCheckTimeout</code>	Optional. The time (in milliseconds) after starting	The default timeout is 30 seconds (30000 ms).

Field	Description	Notes
	that the Greengrass daemon waits for the health check to finish.	
cgroup		
useSystemd	Indicates whether your device uses systemd .	Valid values are yes or no. Run the <code>check_ggc_dependencies</code> script in Module 1 to see if your device uses systemd.
crypto		
The <code>crypto</code> object is added in v1.7.0. It introduces properties that support private key storage on a hardware security module (HSM) through PKCS#11 and local secret storage. For more information, see the section called “Security principals”, the section called “Hardware security integration”, and Deploy secrets to the core. Configurations for private key storage on HSMs or in the file system are supported.		

Field	Description	Notes
caPath	The absolute path to the AWS IoT root CA.	Must be a file URI of the form: <code>file:///absolute/path/to/file</code> .
 Note Make sure that your endpoints correspond to your certificate type.		
PKCS11		

Field	Description	Notes
OpenSSL Engine	Optional. The absolute path to the OpenSSL engine .so file to enable PKCS#11 support on OpenSSL.	Must be a path to a file on the file system. This property is required if you're using the Greengrass OTA update agent with hardware security. For more information, see the section called "Configure OTA updates" .
P11Provider	The absolute path to the PKCS#11 implementation's libdl-loadable library.	Must be a path to a file on the file system.
slotLabel	The slot label that's used to identify the hardware module.	Must conform to PKCS#11 label specifications.
slotUserPin	The user PIN that's used to authenticate the Greengrass core to the module.	Must have sufficient permissions to perform C_Sign with the configured private keys.
principals		
IoT Certificate	The certificate and private key that the core uses to make requests to AWS IoT.	

Field	Description	Notes
IoTCertificate .privateKeyPath	The path to the core private key.	For file system storage, must be a file URI of the form: <i>file:///absolute/path/to/file</i> . For HSM storage, must be an RFC 7512 PKCS#11 path that specifies the object label.
IoTCertificate .certificatePath	The absolute path to the core device certificate.	Must be a file URI of the form: <i>file:///absolute/path/to/file</i> .
MQTTServerCertificate	Optional. The private key that the core uses in combination with the certificate to act as an MQTT server or gateway.	

Field	Description	Notes
MQTTServerCertificatePrivateKeyPath	The path to the local MQTT server private key.	Use this value to specify your own private key for the local MQTT server. For file system storage, must be a file URI of the form: <code>file:///absolute/path/to/file</code>. For HSM storage, must be an RFC 7512 PKCS#11 path that specifies the object label. If this property is omitted, AWS IoT Greengrass rotates the key based your rotation settings. If specified, the customer is responsible for rotating the key.
SecretsManager	The private key that secures the data key used for encryption. For more information, see Deploy secrets to the core .	

Field	Description	Notes
SecretsManager.privateKeyPath	The path to the local secrets manager private key.	Only an RSA key is supported. For file system storage, must be a file URI of the form: <code>file:///absolute/path/to/file</code>. For HSM storage, must be an RFC 7512 PKCS#11 path that specifies the object label. The private key must be generated using the PKCS#1 v1.5 padding mechanism.

The following configuration properties are also supported.

Field	Description	Notes
mqttMaxConnectionRetryInterval	Optional. The maximum interval (in seconds) between MQTT connection retries if the connection is dropped.	Specify this value as an unsigned integer. The default is 60.
managedRespawn	Optional. Indicates that the OTA agent needs to run custom code before an update.	Valid values are true or false. For more information, see OTA updates of AWS IoT Greengrass Core software .

Field	Description	Notes
writeDirectory	Optional. The write directory where AWS IoT Greengrass creates all read/write resources.	For more information, see Configure a write directory for AWS IoT Greengrass .

GGC v1.8

```
{
  "coreThing" : {
    "caPath" : "root.ca.pem",
    "certPath" : "hash.cert.pem",
    "keyPath" : "hash.private.key",
    "thingArn" : "arn:aws:iot:region:account-id:thing/core-thing-name",
    "iotHost" : "host-prefix-ats.iot.region.amazonaws.com",
    "ggHost" : "greengrass-ats.iot.region.amazonaws.com",
    "keepAlive" : 600
  },
  "runtime" : {
    "cgroup" : {
      "useSystemd" : "yes"
    }
  },
  "managedRespawn" : false,
  "crypto" : {
    "principals" : {
      "SecretsManager" : {
        "privateKeyPath" : "file:///greengrass/certs/hash.private.key"
      },
      "IoTCertificate" : {
        "privateKeyPath" : "file:///greengrass/certs/hash.private.key",
        "certificatePath" : "file:///greengrass/certs/hash.cert.pem"
      }
    },
    "caPath" : "file:///greengrass/certs/root.ca.pem"
  }
}
```

The `config.json` file supports the following properties.

coreThing

Field	Description	Notes
caPath	The path to the AWS IoT root CA relative to the <i>/greengrass-root</i> /certs directory.	For backward compatibility with versions earlier than 1.7.0. This property is ignored when the crypto object is present.  Note Make sure that your endpoints correspond to your certificate type.
certPath	The path to the core device certificate relative to the <i>/greengrass-root</i> /certs directory.	For backward compatibility with versions earlier than 1.7.0. This property is ignored when the crypto object is present.
keyPath	The path to the core private key relative to <i>/greengrass-root</i> /certs directory.	For backward compatibility with versions earlier than 1.7.0. This property is ignored when the crypto object is present.
thingArn	The Amazon Resource Name (ARN) of the AWS IoT thing that represents the AWS IoT Greengrass core device.	Find the ARN for your core in the AWS IoT Greengrass console under Cores , or by running the aws greengrass get-core-definition-version CLI command.

Field	Description	Notes
iotHost	Your AWS IoT endpoint.	Find the endpoint in the AWS IoT console under Settings, or by running the aws iot describe-endpoint --endpoint-t-type iot:Data-ATS CLI command. This command returns the Amazon Trust Services (ATS) endpoint. For more information, see the Server authentication documentation.  Note Make sure that your endpoints correspond to your certificate type. Make sure your endpoints correspond to your AWS Region.

Field	Description	Notes
ggHost	Your AWS IoT Greengrass endpoint.	This is your <code>iotHost</code> endpoint with the host prefix replaced by <i>greengrass</i> (for example, <code>greengrass-ats.iot.<i>region</i>.amazonaws.com</code>). Use the same AWS Region as <code>iotHost</code>.  Note Make sure that your endpoints correspond to your certificate type. Make sure your endpoints correspond to your AWS Region.
iotMqttPort	Optional. The port number to use for MQTT communication with AWS IoT.	Valid values are 8883 or 443. The default value is 8883. For more information, see Connect on port 443 or through a network proxy .
iotHttpPort	Optional. The port number used to create HTTPS connections to AWS IoT.	Valid values are 8443 or 443. The default value is 8443. For more information, see Connect on port 443 or through a network proxy .

Field	Description	Notes
ggHttpPort	Optional. The port number used to create HTTPS connections to the AWS IoT Greengrass service.	Valid values are 8443 or 443. The default value is 8443. For more information, see Connect on port 443 or through a network proxy .
keepAlive	Optional. The MQTT KeepAlive period, in seconds.	Valid range is between 30 and 1200 seconds. The default value is 600.
networkProxy	Optional. An object that defines a proxy server to connect to.	The proxy server can be HTTP or HTTPS. For more information, see Connect on port 443 or through a network proxy .

runtime

Field	Description	Notes
cgroup		
useSystemd	Indicates whether your device uses systemd .	Valid values are yes or no. Run the <code>check_ggc_dependencies</code> script in Module 1 to see if your device uses systemd.

crypto

The `crypto` object is added in v1.7.0. It introduces properties that support private key storage on a hardware security module (HSM) through PKCS#11 and local secret storage. For more information, see [the section called “Security principals”](#), [the section called “Hardware](#)

[security integration](#)", and [Deploy secrets to the core](#). Configurations for private key storage on HSMs or in the file system are supported.

Field	Description	Notes
caPath	The absolute path to the AWS IoT root CA.	Must be a file URI of the form: <code>file:///absolute/path/to/file</code> .
 Note Make sure that your endpoints correspond to your certificate type.		
PKCS11		
OpenSSLEngine	Optional. The absolute path to the OpenSSL engine .so file to enable PKCS#11 support on OpenSSL.	Must be a path to a file on the file system. This property is required if you're using the Greengrass OTA update agent with hardware security. For more information, see the section called "Configure OTA updates" .
P11Provider	The absolute path to the PKCS#11 implementation's libdl-loadable library.	Must be a path to a file on the file system.
slotLabel	The slot label that's used to identify the hardware module.	Must conform to PKCS#11 label specifications.

Field	Description	Notes
slotUserPin	The user PIN that's used to authenticate the Greengrass core to the module.	Must have sufficient permissions to perform C_Sign with the configured private keys.
principals		
IoTCertificate	The certificate and private key that the core uses to make requests to AWS IoT.	
IoTCertificate.privateKeyPath	The path to the core private key.	For file system storage, must be a file URI of the form: <i>file:///absolute/path/to/file</i> . For HSM storage, must be an RFC 7512 PKCS#11 path that specifies the object label.
IoTCertificate.certificatePath	The absolute path to the core device certificate.	Must be a file URI of the form: <i>file:///absolute/path/to/file</i> .
MQTTServerCertificate	Optional. The private key that the core uses in combination with the certificate to act as an MQTT server or gateway.	

Field	Description	Notes
MQTTServerCertificatePrivateKeyPath	The path to the local MQTT server private key.	Use this value to specify your own private key for the local MQTT server. For file system storage, must be a file URI of the form: <code>file:///absolute/path/to/file</code>. For HSM storage, must be an RFC 7512 PKCS#11 path that specifies the object label. If this property is omitted, AWS IoT Greengrass rotates the key based your rotation settings. If specified, the customer is responsible for rotating the key.
SecretsManager	The private key that secures the data key used for encryption. For more information, see Deploy secrets to the core .	

Field	Description	Notes
SecretsManager.privateKeyPath	The path to the local secrets manager private key.	Only an RSA key is supported. For file system storage, must be a file URI of the form: <code>file:///absolute/path/to/file</code> . For HSM storage, must be an RFC 7512 PKCS#11 path that specifies the object label. The private key must be generated using the PKCS#1 v1.5 padding mechanism.

The following configuration properties are also supported:

Field	Description	Notes
mqttMaxConnectionRetryInterval	Optional. The maximum interval (in seconds) between MQTT connection retries if the connection is dropped.	Specify this value as an unsigned integer. The default is 60.
managedRespawn	Optional. Indicates that the OTA agent needs to run custom code before an update.	Valid values are true or false. For more information, see OTA updates of AWS IoT Greengrass Core software .

Field	Description	Notes
writeDirectory	Optional. The write directory where AWS IoT Greengrass creates all read/write resources.	For more information, see Configure a write directory for AWS IoT Greengrass .

GGC v1.7

```
{
  "coreThing" : {
    "caPath" : "root.ca.pem",
    "certPath" : "hash.cert.pem",
    "keyPath" : "hash.private.key",
    "thingArn" : "arn:aws:iot:region:account-id:thing/core-thing-name",
    "iotHost" : "host-prefix-ats.iot.region.amazonaws.com",
    "ggHost" : "greengrass-ats.iot.region.amazonaws.com",
    "keepAlive" : 600
  },
  "runtime" : {
    "cgroup" : {
      "useSystemd" : "yes"
    }
  },
  "managedRespawn" : false,
  "crypto" : {
    "principals" : {
      "SecretsManager" : {
        "privateKeyPath" : "file:///greengrass/certs/hash.private.key"
      },
      "IoTCertificate" : {
        "privateKeyPath" : "file:///greengrass/certs/hash.private.key",
        "certificatePath" : "file:///greengrass/certs/hash.cert.pem"
      }
    },
    "caPath" : "file:///greengrass/certs/root.ca.pem"
  }
}
```

The `config.json` file supports the following properties:

coreThing

Field	Description	Notes
caPath	The path to the AWS IoT root CA relative to the <i>/greengrass-root /</i> certs directory.	For backward compatibility with versions earlier than 1.7.0. This property is ignored when the crypto object is present.  Note Make sure that your endpoints correspond to your certificate type.
certPath	The path to the core device certificate relative to the <i>/greengrass-root /</i> certs directory.	For backward compatibility with versions earlier than 1.7.0. This property is ignored when the crypto object is present.
keyPath	The path to the core private key relative to <i>/greengrass-root /</i> certs directory.	For backward compatibility with versions earlier than 1.7.0. This property is ignored when the crypto object is present.
thingArn	The Amazon Resource Name (ARN) of the AWS IoT thing that represents the AWS IoT Greengrass core device.	Find the ARN for your core in the AWS IoT Greengrass console under Cores , or by running the aws greengrass get-core-definition-version CLI command.

Field	Description	Notes
iotHost	Your AWS IoT endpoint.	Find the endpoint in the AWS IoT console under Settings, or by running the aws iot describe-endpoint --endpoint-t-type iot:Data-ATS CLI command. This command returns the Amazon Trust Services (ATS) endpoint. For more information, see the Server authentication documentation.  Note Make sure that your endpoints correspond to your certificate type. Make sure your endpoints correspond to your AWS Region.

Field	Description	Notes
ggHost	Your AWS IoT Greengrass endpoint.	This is your <code>iotHost</code> endpoint with the host prefix replaced by <i>greengrass</i> (for example, <code>greengrass-ats.iot.<i>region</i>.amazonaws.com</code>). Use the same AWS Region as <code>iotHost</code>.  Note Make sure that your endpoints correspond to your certificate type. Make sure your endpoints correspond to your AWS Region.
iotMqttPort	Optional. The port number to use for MQTT communication with AWS IoT.	Valid values are 8883 or 443. The default value is 8883. For more information, see Connect on port 443 or through a network proxy .
keepAlive	Optional. The MQTT KeepAlive period, in seconds.	Valid range is between 30 and 1200 seconds. The default value is 600.

Field	Description	Notes
networkProxy	Optional. An object that defines a proxy server to connect to.	The proxy server can be HTTP or HTTPS. For more information, see Connect on port 443 or through a network proxy .

runtime

Field	Description	Notes
cgroup		
useSystemd	Indicates whether your device uses systemd .	Valid values are yes or no. Run the <code>check_ggc_dependencies</code> script in Module 1 to see if your device uses systemd.

crypto

The `crypto` object, added in v1.7.0, introduces properties that support private key storage on a hardware security module (HSM) through PKCS#11 and local secret storage. For more information, see [the section called “Hardware security integration”](#) and [Deploy secrets to the core](#). Configurations for private key storage on HSMs or in the file system are supported.

Field	Description	Notes
caPath	The absolute path to the AWS IoT root CA.	Must be a file URI of the form: <code>file:///absolute/path/to/file</code> .

Field	Description	Notes
		 Note Make sure that your endpoints correspond to your certificate type.
PKCS11		
OpenSSLEngine	Optional. The absolute path to the OpenSSL engine .so file to enable PKCS#11 support on OpenSSL.	Must be a path to a file on the file system. This property is required if you're using the Greengrass OTA update agent with hardware security. For more information, see the section called "Configure OTA updates".
P11Provider	The absolute path to the PKCS#11 implementation's libdl-loadable library.	Must be a path to a file on the file system.
slotLabel	The slot label that's used to identify the hardware module.	Must conform to PKCS#11 label specifications.
slotUserPin	The user PIN that's used to authenticate the Greengrass core to the module.	Must have sufficient permissions to perform C_Sign with the configured private keys.
principals		

Field	Description	Notes
IoTCertificate	The certificate and private key that the core uses to make requests to AWS IoT.	
IoTCertificate .privateKeyPath	The path to the core private key.	For file system storage, must be a file URI of the form: <i>file:///absolute/path/to/file</i> . For HSM storage, must be an RFC 7512 PKCS#11 path that specifies the object label.
IoTCertificate .certificatePath	The absolute path to the core device certificate.	Must be a file URI of the form: <i>file:///absolute/path/to/file</i> .
MQTTServerCertificate	Optional. The private key that the core uses in combination with the certificate to act as an MQTT server or gateway.	

Field	Description	Notes
MQTTServerCertificatePrivateKeyPath	The path to the local MQTT server private key.	Use this value to specify your own private key for the local MQTT server. For file system storage, must be a file URI of the form: <code>file:///absolute/path/to/file</code>. For HSM storage, must be an RFC 7512 PKCS#11 path that specifies the object label. If this property is omitted, AWS IoT Greengrass rotates the key based your rotation settings. If specified, the customer is responsible for rotating the key.
SecretsManager	The private key that secures the data key used for encryption. For more information, see Deploy secrets to the core .	

Field	Description	Notes
SecretsManager.privateKeyPath	The path to the local secrets manager private key.	Only an RSA key is supported. For file system storage, must be a file URI of the form: <code>file:///absolute/path/to/file</code>. For HSM storage, must be an RFC 7512 PKCS#11 path that specifies the object label. The private key must be generated using the PKCS#1 v1.5 padding mechanism.

The following configuration properties are also supported:

Field	Description	Notes
mqttMaxConnectionRetryInterval	Optional. The maximum interval (in seconds) between MQTT connection retries if the connection is dropped.	Specify this value as an unsigned integer. The default is 60.
managedRespawn	Optional. Indicates that the OTA agent needs to run custom code before an update.	Valid values are <code>true</code> or <code>false</code> . For more information, see OTA updates of AWS IoT Greengrass Core software .

Field	Description	Notes
writeDirectory	Optional. The write directory where AWS IoT Greengrass creates all read/write resources.	For more information, see Configure a write directory for AWS IoT Greengrass .

GGC v1.6

```
{
  "coreThing": {
    "caPath": "root-ca-pem",
    "certPath": "cloud-pem-crt",
    "keyPath": "cloud-pem-key",
    "thingArn": "arn:aws:iot:region:account-id:thing/core-thing-name",
    "iotHost": "host-prefix.iot.region.amazonaws.com",
    "ggHost": "greengrass.iot.region.amazonaws.com",
    "keepAlive": 600,
    "mqttMaxConnectionRetryInterval": 60
  },
  "runtime": {
    "cgroup": {
      "useSystemd": "yes/no"
    }
  },
  "managedRespawn": true,
  "writeDirectory": "/write-directory"
}
```

Note

If you use the **Default Group creation** option from the AWS IoT Greengrass console, then the `config.json` file is deployed to the core device in a working state that specifies the default configuration.

The `config.json` file supports the following properties:

Field	Description	Notes
caPath	The path to the AWS IoT root CA relative to the <code>/greengrass-root /certs</code> directory.	Save the file under <code>/greengrass-root /certs</code> .
certPath	The path to the AWS IoT Greengrass core certificate relative to the <code>/greengrass-root /certs</code> directory.	Save the file under <code>/greengrass-root /certs</code> .
keyPath	The path to the AWS IoT Greengrass core private key relative to <code>/greengrass-root /certs</code> directory.	Save the file under <code>/greengrass-root /certs</code> .
thingArn	The Amazon Resource Name (ARN) of the AWS IoT thing that represents the AWS IoT Greengrass core device.	Find the ARN for your core in the AWS IoT Greengrass console under Cores , or by running the aws greengrass get-core-definition-version CLI command.
iotHost	Your AWS IoT endpoint.	Find this in the AWS IoT console under Settings , or by running the aws iot describe-endpoint CLI command.
ggHost	Your AWS IoT Greengrass endpoint.	This value uses the format <code>greengrass.iot. <i>region</i>.amazonaws.com</code> . Use the same region as <code>iotHost</code> .

Field	Description	Notes
keepAlive	The MQTT KeepAlive period, in seconds.	This is an optional value. The default is 600.
mqttMaxConnectionRetryInterval	The maximum interval (in seconds) between MQTT connection retries if the connection is dropped.	Specify this value as an unsigned integer. This is an optional value. The default is 60.
useSystemd	Indicates whether your device uses systemd .	Valid values are yes or no. Run the <code>check_ggc_dependencies</code> script in Module 1 to see if your device uses systemd.
managedRespawn	An optional over-the-air (OTA) updates feature, this indicates that the OTA agent needs to run custom code before an update.	Valid values are true or false. For more information, see OTA updates of AWS IoT Greengrass Core software .
writeDirectory	The write directory where AWS IoT Greengrass creates all read/write resources.	This is an optional value. For more information, see Configure a write directory for AWS IoT Greengrass .

GGC v1.5

```
{
  "coreThing": {
    "caPath": "root-ca-pem",
    "certPath": "cloud-pem-crt",
    "keyPath": "cloud-pem-key",
    "thingArn": "arn:aws:iot:region:account-id:thing/core-thing-name",
    "iotHost": "host-prefix.iot.region.amazonaws.com",
    "ggHost": "greengrass.iot.region.amazonaws.com",
    "keepAlive": 600
  },
}
```

```

"runtime": {
  "cgroup": {
    "useSystemd": "yes/no"
  },
  "managedRespawn": true
}

```

The `config.json` file exists in `/greengrass-root/config` and contains the following parameters:

Field	Description	Notes
caPath	The path to the AWS IoT root CA relative to the <code>/greengrass-root / certs</code> folder.	Save the file under the <code>/greengrass-root / certs</code> folder.
certPath	The path to the AWS IoT Greengrass core certificate relative to the <code>/greengrass-root /certs</code> folder.	Save the file under the <code>/greengrass-root / certs</code> folder.
keyPath	The path to the AWS IoT Greengrass core private key relative to <code>/greengrass-root/certs</code> folder.	Save the file under the <code>/greengrass-root / certs</code> folder.
thingArn	The Amazon Resource Name (ARN) of the AWS IoT thing that represents the AWS IoT Greengrass core device.	Find the ARN for your core in the AWS IoT Greengrass console under Cores , or by running the aws greengrass get-core-definition-version CLI command.
iotHost	Your AWS IoT endpoint.	Find this in the AWS IoT console under Settings , or by running the aws iot

Field	Description	Notes
		describe-endpoint command.
ggHost	Your AWS IoT Greengrass endpoint.	This value uses the format greengrass.iot. <i>region</i> .amazonaws.com. Use the same region as iotHost.
keepAlive	The MQTT KeepAlive period, in seconds.	This is an optional value. The default value is 600 seconds.
useSystemd	Indicates whether your device uses systemd .	Valid values are yes or no. Run the check_ggc_dependencies script in Module 1 to see if your device uses systemd.
managedRespawn	An optional over-the-air (OTA) updates feature, this indicates that the OTA agent needs to run custom code before an update.	For more information, see OTA updates of AWS IoT Greengrass Core software .

GGC v1.3

```
{
  "coreThing": {
    "caPath": "root-ca-pem",
    "certPath": "cloud-pem-crt",
    "keyPath": "cloud-pem-key",
    "thingArn": "arn:aws:iot:region:account-id:thing/core-thing-name",
    "iotHost": "host-prefix.iot.region.amazonaws.com",
    "ggHost": "greengrass.iot.region.amazonaws.com",
    "keepAlive": 600
  },
}
```

```

"runtime": {
  "cgroup": {
    "useSystemd": "yes/no"
  },
  "managedRespawn": true
}

```

The `config.json` file exists in `/greengrass-root/config` and contains the following parameters:

Field	Description	Notes
caPath	The path to the AWS IoT root CA relative to the <code>/greengrass-root / certs</code> folder.	Save the file under the <code>/greengrass-root / certs</code> folder.
certPath	The path to the AWS IoT Greengrass core certificate relative to the <code>/greengrass-root /certs</code> folder.	Save the file under the <code>/greengrass-root / certs</code> folder.
keyPath	The path to the AWS IoT Greengrass core private key relative to <code>/greengrass-root/certs</code> folder.	Save the file under the <code>/greengrass-root / certs</code> folder.
thingArn	The Amazon Resource Name (ARN) of the AWS IoT thing that represents the AWS IoT Greengrass core.	You can find this value in the AWS IoT Greengrass console under the definition for your AWS IoT thing.
iotHost	Your AWS IoT endpoint.	You can find this value in the AWS IoT console under Settings .
ggHost	Your AWS IoT Greengrass endpoint.	You can find this value in the AWS IoT console under

Field	Description	Notes
		Settings with greengrass. prepended.
keepAlive	The MQTT KeepAlive period, in seconds.	This is an optional value. The default value is 600 seconds.
useSystemd	A binary flag, if your device uses systemd .	Values are yes or no. Use the dependency script in Module 1 to see if your device uses systemd.
managedRespawn	An optional over-the-air (OTA) updates feature, this indicates that the OTA agent needs to run custom code before an update.	For more information, see OTA updates of AWS IoT Greengrass Core software .

GGC v1.1

```
{
  "coreThing": {
    "caPath": "root-ca-pem",
    "certPath": "cloud-pem-crt",
    "keyPath": "cloud-pem-key",
    "thingArn": "arn:aws:iot:region:account-id:thing/core-thing-name",
    "iotHost": "host-prefix.iot.region.amazonaws.com",
    "ggHost": "greengrass.iot.region.amazonaws.com",
    "keepAlive": 600
  },
  "runtime": {
    "cgroup": {
      "useSystemd": "yes/no"
    }
  }
}
```


The `config.json` file exists in `/greengrass-root/config` and contains the following parameters:

Field	Description	Notes
<code>caPath</code>	The path to the AWS IoT root CA relative to the <code>/greengrass-root /certs</code> folder.	Save the file under the <code>/greengrass-root /certs</code> folder.
<code>certPath</code>	The path to the AWS IoT Greengrass core certificate relative to the <code>/greengrass-root /certs</code> folder.	Save the file under the <code>/greengrass-root /certs</code> folder.
<code>keyPath</code>	The path to the AWS IoT Greengrass core private key relative to the <code>/greengrass-root /certs</code> folder.	Save the file under the <code>/greengrass-root /certs</code> folder.
<code>thingArn</code>	The Amazon Resource Name (ARN) of the AWS IoT thing that represents the AWS IoT Greengrass core.	You can find this value in the AWS IoT Greengrass console under the definition for your AWS IoT thing.
<code>iotHost</code>	Your AWS IoT endpoint.	You can find this value in the AWS IoT console under Settings .
<code>ggHost</code>	Your AWS IoT Greengrass endpoint.	You can find this value in the AWS IoT console under Settings with <code>greengrass.</code> prepended.
<code>keepAlive</code>	The MQTT KeepAlive period, in seconds.	This is an optional value. The default value is 600 seconds.

Field	Description	Notes
useSystemd	A binary flag, if your device uses systemd .	Values are yes or no. Use the dependency script in Module 1 to see if your device uses systemd.

GGC v1.0

In AWS IoT Greengrass Core v1.0, `config.json` is deployed to *greengrass-root*/configuration.

```
{
  "coreThing": {
    "caPath": "root-ca-pem",
    "certPath": "cloud-pem-crt",
    "keyPath": "cloud-pem-key",
    "thingArn": "arn:aws:iot:region:account-id:thing/core-thing-name",
    "iotHost": "host-prefix.iot.region.amazonaws.com",
    "ggHost": "greengrass.iot.region.amazonaws.com",
    "keepAlive": 600
  },
  "runtime": {
    "cgroup": {
      "useSystemd": "yes/no"
    }
  }
}
```

The `config.json` file exists in */greengrass-root*/configuration and contains the following parameters:

Field	Description	Notes
caPath	The path to the AWS IoT root CA relative to the <i>/greengrass-root</i> / configuration/certs folder.	Save the file under the <i>/greengrass-root</i> / configuration/certs folder.

Field	Description	Notes
certPath	The path to the AWS IoT Greengrass core certificate relative to the <code>/greengrass-root/configuration/certs</code> folder.	Save the file under the <code>/greengrass-root/configuration/certs</code> folder.
keyPath	The path to the AWS IoT Greengrass core private key relative to the <code>/greengrass-root/configuration/certs</code> folder.	Save the file under the <code>/greengrass-root/configuration/certs</code> folder.
thingArn	The Amazon Resource Name (ARN) of the AWS IoT thing that represents the AWS IoT Greengrass core.	You can find this value in the AWS IoT Greengrass console under the definition for your AWS IoT thing.
iotHost	Your AWS IoT endpoint.	You can find this value in the AWS IoT console under Settings .
ggHost	Your AWS IoT Greengrass endpoint.	You can find this value in the AWS IoT console under Settings with <code>greengrass.</code> prepended.
keepAlive	The MQTT KeepAlive period, in seconds.	This is an optional value. The default value is 600 seconds.
useSystemd	A binary flag if your device uses systemd .	Values are yes or no. Use the dependency script in Module 1 to see if your device uses systemd.

Service endpoints must match the root CA certificate type

Your AWS IoT Core and AWS IoT Greengrass endpoints must correspond to the certificate type of the root CA certificate on your device. If the endpoints and certificate type do not match, authentication attempts fail between the device and AWS IoT Core or AWS IoT Greengrass. For more information, see [Server authentication](#) in the *AWS IoT Developer Guide*.

If your device uses an Amazon Trust Services (ATS) root CA certificate, which is the preferred method, it must also use ATS endpoints for device management and discovery data plane operations. ATS endpoints include the `ats` segment, as shown in the following syntax for the AWS IoT Core endpoint.

```
prefix-ats.iot.region.amazonaws.com
```

Note

For backward compatibility, AWS IoT Greengrass currently supports legacy VeriSign root CA certificates and endpoints in some AWS Regions. If you're using a legacy VeriSign root CA certificate, we recommend that you create an ATS endpoint and use an ATS root CA certificate instead. Otherwise, make sure to use the corresponding legacy endpoints. For more information, see [Supported legacy endpoints](#) in the *Amazon Web Services General Reference*.

Endpoints in config.json

On a Greengrass core device, endpoints are specified in the `coreThing` object in the [config.json](#) file. The `iotHost` property represents the AWS IoT Core endpoint. The `ggHost` property represents the AWS IoT Greengrass endpoint. In the following example snippet, these properties specify ATS endpoints.

```
{
  "coreThing" : {
    ...
    "iotHost" : "abcde1234uvwxyz-ats.iot.us-west-2.amazonaws.com",
    "ggHost" : "greengrass-ats.iot.us-west-2.amazonaws.com",
    ...
  },
}
```

AWS IoT Core endpoint

You can get your AWS IoT Core endpoint by running the [aws iot describe-endpoint](#) CLI command with the appropriate `--endpoint-type` parameter.

- To return an ATS signed endpoint, run:

```
aws iot describe-endpoint --endpoint-type iot:Data-ATS
```

- To return a legacy VeriSign signed endpoint, run:

```
aws iot describe-endpoint --endpoint-type iot:Data
```

AWS IoT Greengrass endpoint

Your AWS IoT Greengrass endpoint is your `iotHost` endpoint with the host prefix replaced by *greengrass*. For example, the ATS signed endpoint is `greengrass-ats.iot.region.amazonaws.com`. This uses the same Region as your AWS IoT Core endpoint.

Connect on port 443 or through a network proxy

This feature is available for AWS IoT Greengrass Core v1.7 and later.

Greengrass cores communicate with AWS IoT Core using the MQTT messaging protocol with TLS client authentication. By convention, MQTT over TLS uses port 8883. However, as a security measure, restrictive environments might limit inbound and outbound traffic to a small range of TCP ports. For example, a corporate firewall might open port 443 for HTTPS traffic, but close other ports that are used for less common protocols, such as port 8883 for MQTT traffic. Other restrictive environments might require all traffic to go through an HTTP proxy before connecting to the internet.

To enable communication in these scenarios, AWS IoT Greengrass allows the following configurations:

- **MQTT with TLS client authentication over port 443.** If your network allows connections to port 443, you can configure the core to use port 443 for MQTT traffic instead of the default port 8883. This can be a direct connection to port 443 or a connection through a network proxy server.

AWS IoT Greengrass uses the [Application Layer Protocol Network](#) (ALPN) TLS extension to enable this connection. As with the default configuration, MQTT over TLS on port 443 uses certificate-based client authentication.

When configured to use a direct connection to port 443, the core supports [over-the-air \(OTA\) updates](#) for AWS IoT Greengrass software. This support requires AWS IoT Greengrass Core v1.9.3 or later.

- **HTTPS communication over port 443.** AWS IoT Greengrass sends HTTPS traffic over port 8443 by default, but you can configure it to use port 443.
- **Connection through a network proxy.** You can configure a network proxy server to act as an intermediary for connecting to the Greengrass core. Only basic authentication and HTTP and HTTPS proxies are supported.

The proxy configuration is passed to user-defined Lambda functions through the `http_proxy`, `https_proxy`, and `no_proxy` environment variables. User-defined Lambda functions must use these passed-in settings to connect through the proxy. Common libraries used by Lambda functions to make connections (such as `boto3` or `cURL` and `python requests` packages) typically use these environment variables by default. If a Lambda function also specifies these same environment variables, AWS IoT Greengrass doesn't override them.

Important

Greengrass cores that are configured to use a network proxy don't support [OTA updates](#).

To configure MQTT over port 443

This feature requires AWS IoT Greengrass Core v1.7 or later.

This procedure allows the Greengrass core to use port 443 for MQTT messaging with AWS IoT Core.

1. Run the following command to stop the Greengrass daemon:

```
cd /greengrass-root/ggc/core/  
sudo ./greengrassd stop
```

2. Open `greengrass-root/config/config.json` for editing as the `su` user.

3. In the `coreThing` object, add the `iotMqttPort` property and set the value to **443**, as shown in the following example.

```
{
  "coreThing" : {
    "caPath" : "root.ca.pem",
    "certPath" : "12345abcde.cert.pem",
    "keyPath" : "12345abcde.private.key",
    "thingArn" : "arn:aws:iot:us-west-2:123456789012:thing/core-thing-name",
    "iotHost" : "abcd123456wxyz-ats.iot.us-west-2.amazonaws.com",
    "iotMqttPort" : 443,
    "ggHost" : "greengrass-ats.iot.us-west-2.amazonaws.com",
    "keepAlive" : 600
  },
  ...
}
```

4. Start the daemon.

```
cd /greengrass-root/ggc/core/
sudo ./greengrassd start
```

To configure HTTPS over port 443

This feature requires AWS IoT Greengrass Core v1.8 or later.

This procedure configures the core to use port 443 for HTTPS communication.

1. Run the following command to stop the Greengrass daemon:

```
cd /greengrass-root/ggc/core/
sudo ./greengrassd stop
```

2. Open `greengrass-root/config/config.json` for editing as the `su` user.
3. In the `coreThing` object, add the `iotHttpPort` and `ggHttpPort` properties, as shown in the following example.

```
{
  "coreThing" : {
```

```

    "caPath" : "root.ca.pem",
    "certPath" : "12345abcde.cert.pem",
    "keyPath" : "12345abcde.private.key",
    "thingArn" : "arn:aws:iot:us-west-2:123456789012:thing/core-thing-name",
    "iotHost" : "abcd123456wxyz-ats.iot.us-west-2.amazonaws.com",
    "iotHttpPort" : 443,
    "ggHost" : "greengrass-ats.iot.us-west-2.amazonaws.com",
    "ggHttpPort" : 443,
    "keepAlive" : 600
  },
  ...
}

```

4. Start the daemon.

```

cd /greengrass-root/ggc/core/
sudo ./greengrassd start

```

To configure a network proxy

This feature requires AWS IoT Greengrass Core v1.7 or later.

This procedure allows AWS IoT Greengrass to connect to the internet through an HTTP or HTTPS network proxy.

1. Run the following command to stop the Greengrass daemon:

```

cd /greengrass-root/ggc/core/
sudo ./greengrassd stop

```

2. Open **greengrass-root**/config/config.json for editing as the su user.

3. In the coreThing object, add the [networkProxy](#) object, as shown in the following example.

```

{
  "coreThing" : {
    "caPath" : "root.ca.pem",
    "certPath" : "12345abcde.cert.pem",
    "keyPath" : "12345abcde.private.key",
    "thingArn" : "arn:aws:iot:us-west-2:123456789012:thing/core-thing-name",

```



```
"iotHost" : "abcd123456wxyz-ats.iot.us-west-2.amazonaws.com",
"ggHost" : "greengrass-ats.iot.us-west-2.amazonaws.com",
"keepAlive" : 600,
"networkProxy": {
  "noProxyAddresses" : "http://128.12.34.56,www.mywebsite.com",
  "proxy" : {
    "url" : "https://my-proxy-server:1100",
    "username" : "Mary_Major",
    "password" : "pass@word1357"
  }
},
...
```

4. Start the daemon.

```
cd /greengrass-root/ggc/core/
sudo ./greengrassd start
```

networkProxy object

Use the `networkProxy` object to specify information about the network proxy. This object has the following properties.

Field	Description
<code>noProxyAddresses</code>	Optional. A comma-separated list of IP addresses or host names that are exempt from the proxy.
<code>proxy</code>	The proxy to connect to. A proxy has the following properties. <code>url</code>. The URL of the proxy server, in the format <code>scheme://userinfo@host:port</code>. <code>scheme</code>. The scheme. Must be <code>http</code> or <code>https</code>.

Field	Description
	• <code>userinfo</code>. Optional. The user name and password information. If specified, the <code>username</code> and <code>password</code> fields are ignored. • <code>host</code>. The host name or IP address of the proxy server. • <code>port</code>. Optional. The port number. If not specified, the following default values are used: • <code>http</code>: 80 • <code>https</code>: 443 • <code>username</code>. Optional. The user name to use to authenticate to the proxy server. • <code>password</code>. Optional. The password to use to authenticate to the proxy server.

Allowing endpoints

Communication between Greengrass devices and AWS IoT Core or AWS IoT Greengrass must be authenticated. This authentication is based on registered X.509 device certificates and cryptographic keys. To allow authenticated requests to pass through proxies without additional encryption, allow the following endpoints.

Endpoint	Port	Description
<code>greengrass. <i>region</i>.amazonaws.com</code>	443	Used for control plane operations for group management.

Endpoint	Port	Description
<i>prefix</i>-ats.iot. <i>region</i>.amazonaws.com or <i>prefix</i>.iot.<i>region</i>.amazonaws.com	MQTT: 8883 or 443 HTTPS: 8443 or 443	Used for data plane operations for device management, such as shadow sync. Allow the use of one or both endpoints, depending on whether your core and client devices use Amazon Trust Services (preferred) root CA certificates, legacy root CA certificates, or both. For more information, see the section called "Service endpoints must match the certificate type".

Endpoint	Port	Description
greengrass-ats.iot . <i>region</i> .amazonaws.com or greengrass.iot. <i>region</i> .amazonaws.com	8443 or 443	Used for device discovery operations. Allow the use of one or both endpoints, depending on whether your core and client devices use Amazon Trust Services (preferred) root CA certificates, legacy root CA certificates, or both. For more information, see the section called “Service endpoints must match the certificate type” .

Endpoint	Port	Description
		 Note Clients that connect on port 443 must implement the Application Layer Protocol Negotiation (ALPN) TLS extension and pass x-amzn-http-ca as the ProtocolName in the ProtocolNameList . For more

Endpoint	Port	Description
		information, see Protocols in the <i>AWS IoT Developer Guide</i> .
*.s3.amazonaws.com	443	Used for deployment operations and over-the-air updates. This format includes the * character because endpoint prefixes are controlled internally and might change at any time.

Endpoint	Port	Description
logs. <i>region</i> .amazonaws.com	443	Required if the Greengrass group is configured to write logs to CloudWatch.

Configure a write directory for AWS IoT Greengrass

This feature is available for AWS IoT Greengrass Core v1.6 and later.

By default, the AWS IoT Greengrass Core software is deployed under a single root directory where AWS IoT Greengrass performs all read and write operations. However, you can configure AWS IoT Greengrass to use a separate directory for all write operations, including creating directories and files. In this case, AWS IoT Greengrass uses two top-level directories:

- The *greengrass-root* directory, which you can leave as read-write or optionally make read-only. This contains the AWS IoT Greengrass Core software and other critical components that should remain immutable during runtime, such as certificates and `config.json`.
- The specified write directory. This contains writable content, such as logs, state information, and deployed user-defined Lambda functions.

This configuration results in the following directory structure.

Greengrass root directory

```
greengrass-root/
|-- certs/
|   |-- root.ca.pem
|   |-- hash.cert.pem
|   |-- hash.private.key
|   |-- hash.public.key
|-- config/
|   |-- config.json
|-- ggc/
```

```
|  |-- packages/  
|      |-- package-version/  
|          |-- bin/  
|              |-- daemon  
|          |-- greengrassd  
|          |-- lambda/  
|          |-- LICENSE/  
|          |-- release_notes_package-version.html  
|              |-- runtime/  
|                  |-- java8/  
|                  |-- nodejs8.10/  
|                  |-- python3.8/  
|  |-- core/
```

Write Directory

```
write-directory/  
|-- packages/  
|   |-- package-version/  
|       |-- ggc_root/  
|       |-- rootfs_nosys/  
|       |-- rootfs_sys/  
|       |-- var/  
|-- deployment/  
|   |-- group/  
|       |-- group.json  
|   |-- lambda/  
|   |-- mlmodel/  
|-- var/  
|   |-- log/  
|   |-- state/
```

To configure a write directory

1. Run the following command to stop the AWS IoT Greengrass daemon:

```
cd /greengrass-root/ggc/core/  
sudo ./greengrassd stop
```

2. Open *greengrass-root*/config/config.json for editing as the su user.

3. Add `writeDirectory` as a parameter and specify the path to the target directory, as shown in the following example.

```
{
  "coreThing": {
    "caPath": "root-CA.pem",
    "certPath": "hash.pem.crt",
    ...
  },
  ...
  "writeDirectory" : "/write-directory"
}
```

Note

You can update the `writeDirectory` setting as often as you want. After the setting is updated, AWS IoT Greengrass uses the newly specified write directory at the next start, but doesn't migrate content from the previous write directory.

4. Now that your write directory is configured, you can optionally make the *greengrass-root* directory read-only. For instructions, see [To Make the Greengrass Root Directory Read-Only](#).

Otherwise, start the AWS IoT Greengrass daemon:

```
cd /greengrass-root/ggc/core/
sudo ./greengrassd start
```

To make the Greengrass root directory read-only

Follow these steps only if you want to make the Greengrass root directory read-only. The write directory must be configured before you begin.

1. Grant access permissions to required directories:
 - a. Give read and write permissions to the `config.json` owner.

```
sudo chmod 0600 /greengrass-root/config/config.json
```

- b. Make `ggc_user` the owner of the certs and system Lambda directories.

```
sudo chown -R ggc_user:ggc_group /greengrass-root/certs/  
sudo chown -R ggc_user:ggc_group /greengrass-root/ggc/packages/1.11.6/lambda/
```

 Note

The `ggc_user` and `ggc_group` accounts are used by default to run system Lambda functions. If you configured the group-level [default access identity](#) to use different accounts, you should give permissions to that user (UID) and group (GID) instead.

2. Make the `greengrass-root` directory read-only by using your preferred mechanism.

 Note

One way to make the `greengrass-root` directory read-only is to mount the directory as read-only. However, to apply over-the-air (OTA) updates to the AWS IoT Greengrass Core software in a mounted directory, the directory must first be unmounted, and then remounted after the update. You can add these `umount` and `mount` operations to the `ota_pre_update` and `ota_post_update` scripts. For more information about OTA updates, see [the section called “Greengrass OTA update agent”](#) and [the section called “Managed respawn with OTA updates”](#).

3. Start the daemon.

```
cd /greengrass-root/ggc/core/  
sudo ./greengrassd start
```

If the permissions from step 1 aren't set correctly, the daemon won't start.

Configure MQTT settings

In the AWS IoT Greengrass environment, local client devices, Lambda functions, connectors, and system components can communicate with each other and with AWS IoT Core. All communication goes through the core, which manages the [subscriptions](#) that authorize MQTT communication between entities.

For information about MQTT settings you can configure for AWS IoT Greengrass, see the following sections:


- [the section called “Message quality of service”](#)
- [the section called “MQTT message queue”](#)
- [the section called “MQTT persistent sessions with AWS IoT Core”](#)
- [the section called “Client IDs for MQTT connections with AWS IoT”](#)
- [MQTT port for local messaging](#)
- [the section called “Timeout for publish, subscribe, unsubscribe operations in MQTT connections with the AWS Cloud”](#)

Note

OPC-UA is an information exchange standard for industrial communication. To implement support for OPC-UA on the Greengrass core, you can use the [IoT SiteWise connector](#). The connector sends industrial device data from OPC-UA servers to asset properties in AWS IoT SiteWise.

Message quality of service

AWS IoT Greengrass supports quality of service (QoS) levels 0 or 1, depending on your configuration and the target and direction of the communication. The Greengrass core acts as a client for communication with AWS IoT Core and a message broker for communication on the local network.



For more information about MQTT and QoS, see [Getting Started](#) on the MQTT website.

Communication with the AWS Cloud

- **Outbound messages use QoS 1**

The core sends messages destined for AWS Cloud targets using QoS 1. AWS IoT Greengrass uses an MQTT message queue to process these messages. If message delivery isn't confirmed by AWS IoT, the message is spooled to be retried later. The message cannot be retried if the queue is full. The message delivery confirmation can help minimize data loss from intermittent connectivity.

Because outbound messages to AWS IoT use QoS 1, the maximum rate at which the Greengrass core can send messages depends on the latency between the core and AWS IoT. Each time the core sends a message, it waits until AWS IoT acknowledges the message before it sends the next message. For example, if the round-trip time between the core and its AWS Region is 50 milliseconds, the core can send up to 20 messages per second. Consider this behavior when you choose the AWS Region where your core connects. To ingest high-volume IoT data to the AWS Cloud, you can use [stream manager](#).

For more information about the MQTT message queue, including how to configure a local storage cache that can persist messages destined for AWS Cloud targets, see [the section called "MQTT message queue"](#).

- **Inbound messages use QoS 0 (default) or QoS 1**

By default, the core subscribes with QoS 0 to messages from AWS Cloud sources. If you enable persistent sessions, the core subscribes with QoS 1. This can help minimize data loss from intermittent connectivity. To manage the QoS for these subscriptions, you configure persistence settings on the local spooler system component.

For more information, including how to enable the core to establish a persistent session with AWS Cloud targets, see [the section called "MQTT persistent sessions with AWS IoT Core"](#).

Communication with local targets

All local communication uses QoS 0. The core makes one attempt to send a message to a local target, which can be a Greengrass Lambda function, connector, or [client device](#). The core doesn't store messages or confirm delivery. Messages can be dropped anywhere between components.

 Note

Although direct communication between Lambda functions doesn't use MQTT messaging, the behavior is the same.

MQTT message queue for cloud targets

MQTT messages that are destined for AWS Cloud targets are queued to await processing. Queued messages are processed in first in, first out (FIFO) order. After a message is processed and published to AWS IoT Core, the message is removed from the queue.

By default, the Greengrass core stores in memory unprocessed messages destined for AWS Cloud targets. You can configure the core to store unprocessed messages in a local storage cache instead. Unlike in-memory storage, the local storage cache has the ability to persist across core restarts (for example, after a group deployment or a device reboot), so AWS IoT Greengrass can continue to process the messages. You can also configure the storage size.

 Warning

The Greengrass core might queue duplicate MQTT messages when it loses connection, because it retries a publish operation before the MQTT client detects that it's offline. To avoid duplicate MQTT messages for cloud targets, configure the core's `keepAlive` value to less than half of its `mqttOperationTimeout` value. For more information, see [AWS IoT Greengrass core configuration file](#).

AWS IoT Greengrass uses the spooler system component (the `GGCloudSpooler` Lambda function) to manage the message queue. You can use the following `GGCloudSpooler` environment variables to configure storage settings.

- **GG_CONFIG_STORAGE_TYPE.** The location of the message queue. The following are valid values:
 - `FileSystem`. Store unprocessed messages in the local storage cache on the disk of the physical core device. When the core restarts, queued messages are retained for processing. Messages are removed after they are processed.
 - `Memory` (default). Store unprocessed messages in memory. When the core restarts, queued messages are lost.

This option is optimized for devices with restricted hardware capabilities. When using this configuration, we recommend that you deploy groups or restart the device when the service disruption is the lowest.

- **GG_CONFIG_MAX_SIZE_BYTES.** The storage size, in bytes. This value can be any non-negative integer **greater than or equal to 262144** (256 KB); a smaller size prevents the AWS IoT Greengrass Core software from starting. The default size is 2.5 MB. When the size limit is reached, the oldest queued messages are replaced by new messages.

Note

This feature is available for AWS IoT Greengrass Core v1.6 and later. Earlier versions use in-memory storage with a queue size of 2.5 MB. You cannot configure storage settings for earlier versions.

To cache messages in local storage

You can configure AWS IoT Greengrass to cache messages to the file system so they persist across core restarts. To do this, you deploy a function definition version where the `GGCloudSpooler` function sets the storage type to `FileSystem`. You must use the AWS IoT Greengrass API to configure the local storage cache. You can't do this in the console.

The following procedure uses the [create-function-definition-version](#) CLI command to configure the spooler to save queued messages to the file system. It also configures a 2.6 MB queue size.

1. Get the IDs of the target Greengrass group and group version. This procedure assumes that this is the latest group and group version. The following query returns the most recently created group.

```
aws greengrass list-groups --query "reverse(sort_by(Groups, &CreationTimestamp))
[0]"
```

Or, you can query by name. Group names are not required to be unique, so multiple groups might be returned.

```
aws greengrass list-groups --query "Groups[?Name=='MyGroup']"
```

 Note

You can also find these values in the AWS IoT console. The group ID is displayed on the group's **Settings** page. Group version IDs are displayed on the group's **Deployments** tab.

2. Copy the Id and LatestVersion values from the target group in the output.
3. Get the latest group version.
 - Replace *group-id* with the Id that you copied.
 - Replace *latest-group-version-id* with the LatestVersion that you copied.

```
aws greengrass get-group-version \  
--group-id group-id \  
--group-version-id latest-group-version-id
```

4. From the Definition object in the output, copy the CoreDefinitionVersionArn and the ARNs of all other group components except FunctionDefinitionVersionArn. You use these values when you create a new group version.
5. From the FunctionDefinitionVersionArn in the output, copy the ID of the function definition. The ID is the GUID that follows the functions segment in the ARN, as shown in the following example.

```
arn:aws:greengrass:us-west-2:123456789012:/greengrass/  
definition/functions/bcfc6b49-beb0-4396-b703-6dEXAMPLEcu5/  
versions/0f7337b4-922b-45c5-856f-1aEXAMPLEsf6
```

 Note

Or, you can create a function definition by running the [create-function-definition](#) command, and then copying the ID from the output.

6. Add a function definition version to the function definition.

- Replace *function-definition-id* with the Id that you copied for the function definition.
- Replace *arbitrary-function-id* with a name for the function, such as **spooler-function**.
- Add any Lambda functions that you want to include in this version to the functions array. You can use the [get-function-definition-version](#) command to get the Greengrass Lambda functions from an existing function definition version.

Warning

Make sure that you specify a value for GG_CONFIG_MAX_SIZE_BYTES that's **greater than or equal to 262144**. A smaller size prevents the AWS IoT Greengrass Core software from starting.

```
aws greengrass create-function-definition-version \
--function-definition-id function-definition-id \
--functions '[{"FunctionArn":
  "arn:aws:lambda:::function:GGCloudSpooler:1", "FunctionConfiguration":
  {"Environment": {"Variables":
    {"GG_CONFIG_MAX_SIZE_BYTES": "2621440", "GG_CONFIG_STORAGE_TYPE": "FileSystem"}}, "Executable":
    "spooler", "MemorySize": 32768, "Pinned": true, "Timeout": 3}, "Id": "arbitrary-
    function-id"}]'
```

Note

If you previously set the GG_CONFIG_SUBSCRIPTION_QUALITY environment variable to [support persistent sessions with AWS IoT Core](#), include it in this function instance.

7. Copy the Arn of the function definition version from the output.
8. Create a group version that contains the system Lambda function.
 - Replace *group-id* with the Id for the group.
 - Replace *core-definition-version-arn* with the CoreDefinitionVersionArn that you copied from the latest group version.

- Replace *function-definition-version-arn* with the Arn that you copied for the new function definition version.
- Replace the ARNs for other group components (for example, SubscriptionDefinitionVersionArn or DeviceDefinitionVersionArn) that you copied from the latest group version.
- Remove any unused parameters. For example, remove the `--resource-definition-version-arn` if your group version doesn't contain any resources.

```
aws greengrass create-group-version \  
--group-id group-id \  
--core-definition-version-arn core-definition-version-arn \  
--function-definition-version-arn function-definition-version-arn \  
--device-definition-version-arn device-definition-version-arn \  
--logger-definition-version-arn logger-definition-version-arn \  
--resource-definition-version-arn resource-definition-version-arn \  
--subscription-definition-version-arn subscription-definition-version-arn
```

9. Copy the Version from the output. This is the ID of the new group version.

10. Deploy the group with the new group version.

- Replace *group-id* with the Id that you copied for the group.
- Replace *group-version-id* with the Version that you copied for the new group version.

```
aws greengrass create-deployment \  
--group-id group-id \  
--group-version-id group-version-id \  
--deployment-type NewDeployment
```

To update the storage settings, you use the AWS IoT Greengrass API to create a new function definition version that contains the GGCloudSpooler function with the updated configuration. Then add the function definition version to a new group version (along with your other group components) and deploy the group version. If you want to restore the default configuration, you can deploy a function definition version that doesn't include the GGCloudSpooler function.

This system Lambda function isn't visible in the console. However, after the function is added to the latest group version, it's included in deployments that you make from the console, unless you use the API to replace or remove it.

MQTT persistent sessions with AWS IoT Core

This feature is available for AWS IoT Greengrass Core v1.10 and later.

A Greengrass core can establish a persistent session with the AWS IoT message broker. A persistent session is an ongoing connection that allows the core to receive messages sent while the core is offline. The core is the client in the connection.

In a persistent session, the AWS IoT message broker saves all subscriptions the core makes during the connection. If the core disconnects, the AWS IoT message broker stores unacknowledged and new messages published as QoS 1 and destined for local targets, such as Lambda functions and [client devices](#). When the core reconnects, the persistent session is resumed and the AWS IoT message broker sends stored messages to the core at a maximum rate of 10 messages per second. Persistent sessions have a default expiry period of 1 hour, which begins when the message broker detects that the core disconnects. For more information, see [MQTT persistent sessions](#) in the *AWS IoT Developer Guide*.

AWS IoT Greengrass uses the spooler system component (the `GGCloudSpooler` Lambda function) to create subscriptions that have AWS IoT as the source. You can use the following `GGCloudSpooler` environment variable to configure persistent sessions.

- **GG_CONFIG_SUBSCRIPTION_QUALITY.** The quality of subscriptions that have AWS IoT as the source. The following are valid values:
 - `AtMostOnce` (default). Disables persistent sessions. Subscriptions use QoS 0.
 - `AtLeastOncePersistent`. Enables persistent sessions. Sets the `cleanSession` flag to `0` in `CONNECT` messages and subscribes with QoS 1.

Messages published with QoS 1 that the core receives are guaranteed to reach the Greengrass daemon's in-memory work queue. The core acknowledges the message after it's added to the queue. Subsequent communication from the queue to the local target (for example, Greengrass Lambda function, connector, or device) is sent as QoS 0. AWS IoT Greengrass doesn't guarantee delivery to local targets.

Note

You can use the [maxWorkItemCount](#) configuration property to control the size of the work item queue. For example, you can increase the queue size if your workload requires heavy MQTT traffic.

When persistent sessions are enabled, the core opens at least one additional connection for MQTT message exchange with AWS IoT. For more information, see [the section called “Client IDs for MQTT connections with AWS IoT”](#).

To configure MQTT persistent sessions

You can configure AWS IoT Greengrass to use persistent sessions with AWS IoT Core. To do this, you deploy a function definition version where the `GGCloudSpooler` function sets the subscription quality to `AtLeastOncePersistent`. This setting applies to all your subscriptions that have AWS IoT Core (`cloud`) as the source. You must use the AWS IoT Greengrass API to configure persistent sessions. You can't do this in the console.

The following procedure uses the [create-function-definition-version](#) CLI command to configure the spooler to use persistent sessions. In this procedure, we assume that you're updating the configuration of the latest group version of an existing group.

1. Get the IDs of the target Greengrass group and group version. This procedure assumes that this is the latest group and group version. The following query returns the most recently created group.

```
aws greengrass list-groups --query "reverse(sort_by(Groups, &CreationTimestamp))
[0]"
```

Or, you can query by name. Group names are not required to be unique, so multiple groups might be returned.

```
aws greengrass list-groups --query "Groups[?Name=='MyGroup']"
```

Note

You can also find these values in the AWS IoT console. The group ID is displayed on the group's **Settings** page. Group version IDs are displayed on the group's **Deployments** tab.

2. Copy the Id and LatestVersion values from the target group in the output.
3. Get the latest group version.
 - Replace *group-id* with the Id that you copied.
 - Replace *latest-group-version-id* with the LatestVersion that you copied.

```
aws greengrass get-group-version \  
--group-id group-id \  
--group-version-id latest-group-version-id
```

4. From the Definition object in the output, copy the CoreDefinitionVersionArn and the ARNs of all other group components except FunctionDefinitionVersionArn. You use these values when you create a new group version.
5. From the FunctionDefinitionVersionArn in the output, copy the ID of the function definition. The ID is the GUID that follows the functions segment in the ARN, as shown in the following example.

```
arn:aws:greengrass:us-west-2:123456789012:/greengrass/  
definition/functions/bfc6b49-beb0-4396-b703-6dEXAMPLEcu5/  
versions/0f7337b4-922b-45c5-856f-1aEXAMPLEsf6
```

Note

Or, you can create a function definition by running the [create-function-definition](#) command, and then copying the ID from the output.

6. Add a function definition version to the function definition.
 - Replace *function-definition-id* with the Id that you copied for the function definition.

- Replace *arbitrary-function-id* with a name for the function, such as **spooler-function**.
- Add any Lambda functions that you want to include in this version to the functions array. You can use the [get-function-definition-version](#) command to get the Greengrass Lambda functions from an existing function definition version.

```
aws greengrass create-function-definition-version \  
--function-definition-id function-definition-id \  
--functions '[{"FunctionArn":  
  "arn:aws:lambda::function:GGCloudSpooler:1", "FunctionConfiguration":  
  {"Environment": {"Variables":  
    {"GG_CONFIG_SUBSCRIPTION_QUALITY": "AtLeastOncePersistent"}}}, "Executable":  
  "spooler", "MemorySize": 32768, "Pinned": true, "Timeout": 3}, {"Id": "arbitrary-  
function-id"}]'
```

 Note

If you previously set the GG_CONFIG_STORAGE_TYPE or GG_CONFIG_MAX_SIZE_BYTES environment variables to [define storage settings](#), include them in this function instance.

7. Copy the Arn of the function definition version from the output.
8. Create a group version that contains the system Lambda function.
 - Replace *group-id* with the Id for the group.
 - Replace *core-definition-version-arn* with the CoreDefinitionVersionArn that you copied from the latest group version.
 - Replace *function-definition-version-arn* with the Arn that you copied for the new function definition version.
 - Replace the ARNs for other group components (for example, SubscriptionDefinitionVersionArn or DeviceDefinitionVersionArn) that you copied from the latest group version.
 - Remove any unused parameters. For example, remove the `--resource-definition-version-arn` if your group version doesn't contain any resources.

```
aws greengrass create-group-version \  
--group-id group-id \  
--core-definition-version-arn core-definition-version-arn \  
--function-definition-version-arn function-definition-version-arn \  
--device-definition-version-arn device-definition-version-arn \  
--logger-definition-version-arn logger-definition-version-arn \  
--resource-definition-version-arn resource-definition-version-arn \  
--subscription-definition-version-arn subscription-definition-version-arn
```

9. Copy the Version from the output. This is the ID of the new group version.
10. Deploy the group with the new group version.
 - Replace *group-id* with the Id that you copied for the group.
 - Replace *group-version-id* with the Version that you copied for the new group version.

```
aws greengrass create-deployment \  
--group-id group-id \  
--group-version-id group-version-id \  
--deployment-type NewDeployment
```

11. (Optional) Increase the [maxWorkItemCount](#) property in the core configuration file. This can help the core handle increased MQTT traffic and communication with local targets.

To update the core with these configuration changes, you use the AWS IoT Greengrass API to create a new function definition version that contains the `GGCloudSpooler` function with the updated configuration. Then add the function definition version to a new group version (along with your other group components) and deploy the group version. If you want to restore the default configuration, you can create a function definition version that doesn't include the `GGCloudSpooler` function.

This system Lambda function isn't visible in the console. However, after the function is added to the latest group version, it's included in deployments that you make from the console, unless you use the API to replace or remove it.

Client IDs for MQTT connections with AWS IoT

This feature is available for AWS IoT Greengrass Core v1.8 and later.

The Greengrass core opens MQTT connections with AWS IoT Core for operations such as shadow sync and certificate management. For these connections, the core generates predictable client IDs based on the core thing name. Predictable client IDs can be used with monitoring, auditing, and pricing features, including AWS IoT Device Defender and [AWS IoT lifecycle events](#). You can also create logic around predictable client IDs (for example, [subscribe policy](#) templates based on certificate attributes).

GGC v1.9 and later

Two Greengrass system components open MQTT connections with AWS IoT Core. These components use the following patterns to generate the client IDs for the connections.

Operation	Client ID pattern
Deployments	<i>core-thing-name</i> Example: MyCoreThing Use this client ID for connect, disconnect, subscribe, and unsubscribe lifecycle event notifications.
Subscriptions	<i>core-thing-name</i> -cn Example: MyCoreThing-c01 <i>n</i> is an integer that starts at 00 and increments with each new connection to a maximum number of 250. The number of connections is determined by the number of devices that sync their shadow state with AWS IoT Core (maximum 2,500 per group) and the number of subscriptions with ccloud as their source in the group (maximum 10,000 per group). The spooler system component connects with AWS IoT Core to exchange messages for subscriptions with a cloud source or target.

Operation	Client ID pattern
	The spooler also acts as proxy for message exchange between AWS IoT Core and the local shadow service and device certificate manager.

To calculate the number of MQTT connections per group, use the following formula:

number of MQTT connections per group = number of connections for Deployment Agent + number of connections for Subscriptions

Where,

- number of connections for Deployment Agent = 1.
- number of connections for Subscriptions = (2 subscriptions for supporting certificate generation + number of MQTT topics in AWS IoT Core + number of device shadows synced) / 50.
- Where, 50 = the maximum number of subscriptions per connection that AWS IoT Core can support.

Note

If you enable [persistent sessions](#) for subscription with AWS IoT Core, the core opens at least one additional connection to use in a persistent session. The system components don't support persistent sessions, so they can't share that connection.

To reduce the number of MQTT connections and help reduce costs, you can use local Lambda functions to aggregate data at the edge. Then you send the aggregated data to the AWS Cloud. As a result, you use fewer MQTT topics in AWS IoT Core. For more information, see [AWS IoT Greengrass Pricing](#).

GGC v1.8

Several Greengrass system components open MQTT connections with AWS IoT Core. These components use the following patterns to generate the client IDs for the connections.

Operation	Client ID pattern
Deployments	<i>core-thing-name</i> Example: MyCoreThing Use this client ID for connect, disconnect, subscribe, and unsubscribe lifecycle event notifications.
MQTT message exchange with AWS IoT Core	<i>core-thing-name</i> -spr Example: MyCoreThing-spr
Shadow sync	<i>core-thing-name</i> -snn Example: MyCoreThing-s01 <i>nn</i> is an integer that starts at 00 and increments with each new connection to a maximum of 03. The number of connections is determined by the number of devices (maximum 200 devices per group) that sync their shadow state with AWS IoT Core (maximum 50 subscriptions per connection).
Device certificate management	<i>core-thing-name</i> -dcm Example: MyCoreThing-dcm

Note

Duplicate client IDs used in simultaneous connections can cause an infinite connect-disconnect loop. This can happen if another device is hardcoded to use the core device name as the client ID in connections. For more information, see this [troubleshooting step](#).

Greengrass devices are also fully integrated with the Fleet Indexing service of AWS IoT Device Management. This allows you to index and search for devices based on device attributes, shadow state, and connection state in the cloud. For example, Greengrass devices establish at least one connection that uses the thing name as the client ID, so you can use device connectivity indexing to discover which Greengrass devices are currently connected or disconnected to AWS IoT Core. For more information, see [Fleet indexing service](#) in the *AWS IoT Developer Guide*.

Configure the MQTT port for local messaging

This feature requires AWS IoT Greengrass Core v1.10 or later.

The Greengrass core acts as the local message broker for MQTT messaging between local Lambda functions, connectors, and [client devices](#). By default, the core uses port 8883 for MQTT traffic on the local network. You might want to change the port to avoid a conflict with other software that runs on port 8883.

To configure the port number that the core uses for local MQTT traffic

1. Run the following command to stop the Greengrass daemon:

```
cd /greengrass-root/ggc/core/  
sudo ./greengrassd stop
```

2. Open *greengrass-root*/config/config.json for editing as the su user.
3. In the coreThing object, add the ggMqttPort property and set the value to the port number you want to use. Valid values are 1024 to 65535. The following example sets the port number to 9000.

```
{  
  "coreThing" : {  
    "caPath" : "root.ca.pem",  
    "certPath" : "12345abcde.cert.pem",  
    "keyPath" : "12345abcde.private.key",  
    "thingArn" : "arn:aws:iot:us-west-2:123456789012:thing/core-thing-name",  
    "iotHost" : "abcd123456wxyz-ats.iot.us-west-2.amazonaws.com",  
    "ggHost" : "greengrass-ats.iot.us-west-2.amazonaws.com",  
    "ggMqttPort" : 9000,  
    "keepAlive" : 600  
  },  
  ...  
}
```

```
}
```

4. Start the daemon.

```
cd /greengrass-root/ggc/core/
sudo ./greengrassd start
```

5. If [automatic IP detection](#) is enabled for the core, the configuration is complete.

If automatic IP detection is not enabled, you must update the connectivity information for the core. This allows client devices to receive the correct port number during discovery operations to acquire core connectivity information. You can use the AWS IoT console or AWS IoT Greengrass API to update the core connectivity information. For this procedure, you update the port number only. The local IP address for the core remains the same.

To update the connectivity information for the core (console)

1. On the group configuration page, choose the Greengrass core.
2. On the core details page, choose the **MQTT broker endpoints** tab.
3. Choose **Manage endpoints** and then choose **Add endpoint**
4. Enter your current local IP address and the new port number. The following example sets the port number 9000 for the IP address 192.168.1.8.
5. Remove the obsolete endpoint, and then choose **Update**

To update the connectivity information for the core (API)

- Use the [UpdateConnectivityInfo](#) action. The following example uses update-connectivity-info in the AWS CLI to set the port number 9000 for the IP address 192.168.1.8.

```
aws greengrass update-connectivity-info \
  --thing-name "MyGroup_Core" \
  --connectivity-info "[{\"Metadata\":\"\",\"PortNumber\":9000,
  \"HostAddress\":\"192.168.1.8\",\"Id\":\"localIP_192.168.1.8\"},{\"Metadata
  \":\", \"PortNumber\":8883, \"HostAddress\":\"127.0.0.1\", \"Id\":
  \"localhost_127.0.0.1_0\"}]\""
```

Note

You can also configure the port that the core uses for MQTT messaging with AWS IoT Core. For more information, see [the section called “Connect on port 443 or through a network proxy”](#).

Timeout for publish, subscribe, unsubscribe operations in MQTT connections with the AWS Cloud

This feature is available in AWS IoT Greengrass v1.10.2 or later.

You can configure the amount of time (in seconds) to allow the Greengrass core to complete a publish, subscribe, or unsubscribe operation in MQTT connections to AWS IoT Core. You might want to adjust this setting if the operations time out because of bandwidth constraints or high latency. To configure this setting in the [config.json](#) file, add or change the `mqttOperationTimeout` property in the `coreThing` object. For example:

```
{
  "coreThing": {
    "mqttOperationTimeout": 10,
    "caPath": "root-ca.pem",
    "certPath": "hash.cert.pem",
    "keyPath": "hash.private.key",
    ...
  },
  ...
}
```

The default timeout is 5 seconds. The minimum timeout is 5 seconds.

Activate automatic IP detection

You can configure AWS IoT Greengrass to enable client devices in a Greengrass group to automatically discover the Greengrass core. When enabled, the core watches for changes to its IP addresses. If an address changes, the core publishes an updated list of addresses. These addresses are made available to client devices that are in the same Greengrass group as the core.

Note

The AWS IoT policy for client devices must grant the `greengrass:Discover` permission to allow devices to retrieve connectivity information for the core. For more information about the policy statement, see [the section called “Discovery authorization”](#).

To enable this feature from the AWS IoT Greengrass console, choose **Automatic detection** when you deploy your Greengrass group for the first time. You can also enable or disable this feature on the group configuration page by choosing the **Lambda functions** tab and selecting the **IP detector**. Automatic IP detection is enabled if **Automatically detect and override MQTT broker endpoints** is selected.

To manage automatic discovery with the AWS IoT Greengrass API, you must configure the `IPDetector` system Lambda function. The following procedure shows how to use the [create-function-definition-version](#) CLI command to configure automatic discovery of the Greengrass core.

1. Get the IDs of the target Greengrass group and group version. This procedure assumes that this is the latest group and group version. The following query returns the most recently created group.

```
aws greengrass list-groups --query "reverse(sort_by(Groups, &CreationTimestamp))
[0]"
```

Or, you can query by name. Group names are not required to be unique, so multiple groups might be returned.

```
aws greengrass list-groups --query "Groups[?Name=='MyGroup']"
```

Note

You can also find these values in the AWS IoT console. The group ID is displayed on the group's **Settings** page. Group version IDs are displayed on the group's **Deployments** tab.

2. Copy the `Id` and `LatestVersion` values from the target group in the output.
3. Get the latest group version.

- Replace *group-id* with the Id that you copied.
- Replace *latest-group-version-id* with the LatestVersion that you copied.

```
aws greengrass get-group-version \  
--group-id group-id \  
--group-version-id latest-group-version-id
```

4. From the Definition object in the output, copy the CoreDefinitionVersionArn and the ARNs of all other group components except FunctionDefinitionVersionArn. You use these values when you create a new group version.
5. From the FunctionDefinitionVersionArn in the output, copy the ID of the function definition and the function definition version:

```
arn:aws:greengrass:region:account-id:/greengrass/groups/function-definition-id/  
versions/function-definition-version-id
```

 Note

You can optionally create a function definition by running the [create-function-definition](#) command, and then copy the ID from the output.

6. Use the [get-function-definition-version](#) command to get the current definition state. Use the *function-definition-id* you copied for the function definition. For example, *4d941bc7-92a1-4f45-8d64-EXAMPLEf76c3*.

```
aws greengrass get-function-definition-version  
--function-definition-id function-definition-id  
--function-definition-version-id function-definition-version-id
```

Make a note of the listed function configurations. You will need to include these when creating a new function definition version in order to prevent loss of your current definition settings.

7. Add a function definition version to the function definition.
 - Replace *function-definition-id* with the Id that you copied for the function definition. For example, *4d941bc7-92a1-4f45-8d64-EXAMPLEf76c3*.

- Replace *arbitrary-function-id* with a name for the function, such as **auto-detection-function**.
- Add all Lambda functions that you want to include in this version to the functions array, such as any listed in the previous step.

```
aws greengrass create-function-definition-version \  
--function-definition-id function-definition-id \  
--functions  
'[{"FunctionArn":"arn:aws:lambda::function:GGIPDetector:1","Id": "arbitrary-  
function-id", "FunctionConfiguration":  
{"Pinned":true, "MemorySize":32768, "Timeout":3}}]' \  
--region us-west-2
```

8. Copy the Arn of the function definition version from the output.
9. Create a group version that contains the system Lambda function.
 - Replace *group-id* with the Id for the group.
 - Replace *core-definition-version-arn* with the CoreDefinitionVersionArn that you copied from the latest group version.
 - Replace *function-definition-version-arn* with the Arn that you copied for the new function definition version.
 - Replace the ARNs for other group components (for example, SubscriptionDefinitionVersionArn or DeviceDefinitionVersionArn) that you copied from the latest group version.
 - Remove any unused parameters. For example, remove the `--resource-definition-version-arn` if your group version doesn't contain any resources.

```
aws greengrass create-group-version \  
--group-id group-id \  
--core-definition-version-arn core-definition-version-arn \  
--function-definition-version-arn function-definition-version-arn \  
--device-definition-version-arn device-definition-version-arn \  
--logger-definition-version-arn logger-definition-version-arn \  
--resource-definition-version-arn resource-definition-version-arn \  
--subscription-definition-version-arn subscription-definition-version-arn
```

10. Copy the Version from the output. This is the ID of the new group version.

11. Deploy the group with the new group version.

- Replace *group-id* with the Id that you copied for the group.
- Replace *group-version-id* with the Version that you copied for the new group version.

```
aws greengrass create-deployment \  
--group-id group-id \  
--group-version-id group-version-id \  
--deployment-type NewDeployment
```

If you want to manually input the IP address of your Greengrass core, you can complete this tutorial with a different function definition that does not include the `IPDetector` function. This will prevent the detection function from locating and automatically inputting your Greengrass core IP address.

This system Lambda function isn't visible in the Lambda console. After the function is added to the latest group version, it's included in deployments that you make from the console, unless you use the API to replace or remove it.

Configure the init system to start the Greengrass daemon

It's a good practice to set up your init system to start the Greengrass daemon during boot, especially when managing large fleets of devices.

Note

If you used `apt` to install the AWS IoT Greengrass Core software, you can use the `systemd` scripts to enable start on boot. For more information, see [the section called "Use systemd scripts to manage the Greengrass daemon lifecycle"](#).

There are different types of init system, such as `initd`, `systemd`, and `SystemV`, and they use similar configuration parameters. The following example is a service file for `systemd`. The `Type` parameter is set to `forking` because `greengrassd` (which is used to start Greengrass) forks the Greengrass daemon process, and the `Restart` parameter is set to `on-failure` to direct `systemd` to restart Greengrass if Greengrass enters a failed state.

Note

To see if your device uses systemd, run the `check_ggc_dependencies` script as described in [Module 1](#). Then to use systemd, make sure that the `useSystemd` parameter in [config.json](#) is set to `yes`.

```
[Unit]
Description=Greengrass Daemon

[Service]
Type=forking
PIDFile=/var/run/greengrassd.pid
Restart=on-failure
ExecStart=/greengrass/ggc/core/greengrassd start
ExecReload=/greengrass/ggc/core/greengrassd restart
ExecStop=/greengrass/ggc/core/greengrassd stop

[Install]
WantedBy=multi-user.target
```

See also

- [What is AWS IoT Greengrass?](#)
- [the section called “Supported platforms and requirements”](#)
- [Getting started with AWS IoT Greengrass](#)
- [the section called “Overview of the group object model”](#)
- [the section called “Hardware security integration”](#)

AWS IoT Greengrass Version 1 maintenance policy

Use this AWS IoT Greengrass V1 maintenance policy to understand the different levels of maintenance and updates for the AWS IoT Greengrass V1 service and the AWS IoT Greengrass Core software v1.x.

Topics

- [AWS IoT Greengrass versioning scheme](#)
- [Lifecycle phases for major versions of the AWS IoT Greengrass Core software](#)
- [Maintenance policy for AWS IoT Greengrass Core software](#)
- [Deprecation schedule](#)
- [Support policy for AWS Lambda functions on Greengrass core devices](#)
- [Support policy for AWS IoT Device Tester for AWS IoT Greengrass V1](#)
- [End of maintenance schedule](#)

AWS IoT Greengrass versioning scheme

AWS IoT Greengrass uses [semantic versioning](#) for the AWS IoT Greengrass Core software. Semantic versions follow a *major.minor.patch* number system. The major version increments for functional and API changes that aren't backward-compatible with previous major versions. The minor version increments for releases that add new backward-compatible functionality. The patch version increments for security patches or bug fixes. Since its first major release, v1.0.0, AWS IoT Greengrass has released 11 minor versions of the AWS IoT Greengrass Core software v1.x, where v1.11.6 is the latest release. We recommend that you update your AWS IoT Greengrass Core software to the latest available version to take advantage of new features, enhancements, and bug fixes.

In December 2020, AWS IoT Greengrass released its first major version update. This update included the AWS IoT Greengrass V2 service and version 2.0.3 of the AWS IoT Greengrass Core software. For new applications, we strongly recommend that you use AWS IoT Greengrass Version 2 and the AWS IoT Greengrass Core software v2.x. Version 2 receives new features, includes all key V1 features, and supports additional platforms and continuous deployments to large fleets of devices. For more information, see [What is AWS IoT Greengrass V2?](#).

Lifecycle phases for major versions of the AWS IoT Greengrass Core software

Each major version of the AWS IoT Greengrass Core software has the following three sequential lifecycle phases. Each lifecycle phase provides different levels of maintenance over a period of time after the initial release date.

- **Release phase** – AWS IoT Greengrass may release the following updates:
 - Minor version updates that provide new features or enhancements to existing features
 - Patch version updates that provide security patches and bug fixes
- **Maintenance phase** – AWS IoT Greengrass may release patch version updates that provide security patches and bug fixes. AWS IoT Greengrass won't release new features or enhancements to existing features during the maintenance phase.
- **Extended life phase** – AWS IoT Greengrass won't release updates that provide features, enhancements to existing features, security patches, or bug fixes. However, the AWS Cloud endpoints and API operations will remain available and operate according to the [AWS IoT Greengrass Service Level Agreement](#). Devices that run the AWS IoT Greengrass Core software v1.x can continue to connect to the AWS Cloud and operate.

After the extended life phase ends for a major version of AWS IoT Greengrass, the AWS Cloud endpoints and API operations will be deprecated and no longer available. Devices that run the AWS IoT Greengrass Core software v1.x won't be able to connect to AWS Cloud services to operate.

Maintenance policy for AWS IoT Greengrass Core software

The AWS IoT Greengrass Core software v1.x entered the extended life phase on June 30, 2023. After this date, the AWS IoT Greengrass Core software v1.x will remain in the extended life phase until further notice.

The AWS IoT Greengrass Core software v2.x is currently in the release phase, and it will remain in the release phase until further notice. AWS IoT Greengrass continues to add new features and enhancements to the AWS IoT Greengrass Core software v2.x. For example, AWS IoT Greengrass released Windows support in v2.5.0 of the AWS IoT Greengrass Core software. AWS IoT Greengrass releases security patches and bug fixes for all minor versions of AWS IoT Greengrass Core v2.x for at least 1 year after the release date. For more information, see [What's new in AWS IoT Greengrass V2](#).

Maintenance phase schedule

On June 30, 2023, the maintenance phase ended for the AWS IoT Greengrass Core software v1.11.x. On March 31, 2022, the maintenance phase ended for the AWS IoT Greengrass Core software v1.10.x. The maintenance phase ends for certain AWS IoT Greengrass Core software v1.x artifacts and features earlier than these dates. For more information, see [End of maintenance schedule](#).

If you have an AWS Support plan, the maintenance phase for AWS IoT Greengrass Core software v1.x doesn't affect your AWS Support plan. You can continue to open AWS Support tickets even after the maintenance phase ends. If you have questions or concerns, contact your AWS Support contact, or ask a question on [AWS re:Post](#) using the **AWS IoT Greengrass** tag.

Deprecation schedule

Currently, there is no plan to stop supporting the AWS IoT Greengrass Core software v1.x. The AWS IoT Greengrass V1 endpoints and API operations will remain available until further notice. The AWS IoT Greengrass Core software v1.11.6 entered the extended life phase on June 30, 2023. During this phase, devices that run the AWS IoT Greengrass Core software v1.x can continue to connect to the AWS IoT Greengrass V1 service to operate until further notice.

If AWS IoT Greengrass V1 stops being supported in the future, AWS IoT Greengrass will provide 12 months advance notice before this happens. This will help you plan the update of your applications to use AWS IoT Greengrass V2 and the AWS IoT Greengrass Core software v2.x. For more information about how to update your applications to V2, see [Move from AWS IoT Greengrass V1 to V2](#).

Support policy for AWS Lambda functions on Greengrass core devices

AWS IoT Greengrass enables you to run AWS Lambda functions on IoT devices. AWS Lambda provides a support policy and timelines that determine support for Lambda runtimes in AWS IoT Greengrass. After a Lambda runtime reaches the end of support phase, AWS IoT Greengrass also ends support for that runtime. For more information, see [Runtime support policy](#) in the *AWS Lambda Developer Guide*.

When a Lambda runtime reaches end of support, you can't create or update Lambda functions that use that runtime. However, you can continue to deploy these Lambda functions to Greengrass core devices and invoke deployed Lambda functions. This policy also applies to AWS IoT Greengrass V2.

Support policy for AWS IoT Device Tester for AWS IoT Greengrass V1

AWS IoT Device Tester (IDT) for AWS IoT Greengrass V1 enables you to validate and [qualify](#) your AWS IoT Greengrass devices for inclusion in the [AWS Partner Device Catalog](#). As of April 4, 2022, AWS IoT Device Tester (IDT) for AWS IoT Greengrass V1 no longer generates signed qualification reports. You can no longer qualify new AWS IoT Greengrass V1 devices to list in the [AWS Partner Device Catalog](#) through the [AWS Device Qualification Program](#). While you can't qualify Greengrass V1 devices, you can continue to use IDT for AWS IoT Greengrass V1 to test your Greengrass V1 devices. We recommend that you use [IDT for AWS IoT Greengrass V2](#) to qualify and list Greengrass devices in the [AWS Partner Device Catalog](#). For more information, see [Support policy for AWS IoT Device Tester for AWS IoT Greengrass V1](#).

End of maintenance schedule

The following table lists end of maintenance dates for AWS IoT Greengrass Core v1.x artifacts and features. If you have questions about the maintenance schedule or policy, contact [AWS Support](#).

Artifact or feature	End of maintenance date
Greengrass APT repository installation	February 11, 2022
ML Image Classification connector	March 31, 2022
ML Object Detection connector	March 31, 2022
ML Feedback connector	March 31, 2022
AWS IoT Analytics connector	March 31, 2022
Twilio Notifications connector	March 31, 2022
Splunk Integration connector	March 31, 2022

Artifact or feature	End of maintenance date
Serial Stream connector	March 31, 2022
ServiceNow MetricBase Integration connector	March 31, 2022
Raspberry Pi GPIO connector	March 31, 2022
AWS IoT Greengrass Core software v1.10.x	March 31, 2022
AWS IoT Greengrass Core software v1.x Docker images	June 30, 2022
AWS IoT Greengrass Core software v1.11.x	June 30, 2023
AWS IoT Greengrass Core software v1.11.x Snap	December 31, 2023

End of maintenance for AWS IoT Greengrass Core software v1.x Docker images

On June 30, 2022, AWS IoT Greengrass ended maintenance for AWS IoT Greengrass Core software v1.x Docker images that are published to Amazon Elastic Container Registry (Amazon ECR) and Docker Hub. You can continue to download these Docker images from Amazon ECR and Docker Hub until June 30, 2023, which is 1 year after maintenance ended. However, the AWS IoT Greengrass Core software v1.x Docker images no longer receive security patches or bug fixes after maintenance ended on June 30, 2022. If you run a production workload that depends on these Docker images, we recommend that you build your own Docker images using the Dockerfiles that AWS IoT Greengrass provides. For more information, see [AWS IoT Greengrass Docker software](#).

End of maintenance for AWS IoT Greengrass Core software v1.x APT repository

On February 11, 2022, AWS IoT Greengrass ended maintenance for the option to [install the AWS IoT Greengrass Core software v1.x from an APT repository](#). The APT repository was removed on this date, so you can no longer to use the APT repository to update the AWS IoT Greengrass Core software or install the AWS IoT Greengrass Core software on new devices. On devices where you

added the AWS IoT Greengrass repository, you must [remove the repository from the sources list](#). We recommend that you update the AWS IoT Greengrass Core software v1.x using [tar files](#).

End of maintenance for AWS IoT Greengrass Core software v1.11.x Snap

On December 31, 2023, AWS IoT Greengrass will end maintenance for the AWS IoT Greengrass core software version 1.11.x Snap that is published on [snapcraft.io](#). Devices currently running the Snap will continue to work until further notice. However, the AWS IoT Greengrass core Snap will no longer receive security patches or bug fixes after maintenance ends.

Getting started with AWS IoT Greengrass

This Getting Started tutorial includes several modules designed to show you AWS IoT Greengrass basics and help you get started using AWS IoT Greengrass. This tutorial covers fundamental concepts, such as:

- Configuring AWS IoT Greengrass cores and groups.
- The deployment process for running AWS Lambda functions at the edge.
- Connecting AWS IoT devices, called client devices, to the AWS IoT Greengrass core.
- Creating subscriptions to allow MQTT communication between local Lambda functions, client devices, and AWS IoT.

Choose how to get started with AWS IoT Greengrass


You can choose how to use this tutorial to set up your core device:

- Run [Greengrass device setup](#) on your core device, which takes you from installing AWS IoT Greengrass dependencies to testing a Hello World Lambda function in minutes. This script reproduces the steps in Module 1 through Module 3-1.

- or -

- Walk through the steps in Module 1 through Module 3-1 to examine Greengrass requirements and processes more closely. These steps set up your core device, create and configure a Greengrass group that contains a Hello World Lambda function, and deploy your Greengrass group. Typically, this takes an hour or two to complete.

Quick Start: Greengrass Device Setup



Quick Start

[Greengrass device setup](#) configures your core device and Greengrass resources. The script:

- Installs AWS IoT Greengrass dependencies.
- Downloads the root CA certificate and core device certificate and keys.
- Downloads, installs, and configures the AWS IoT Greengrass Core software on your device.
- Starts the Greengrass daemon process on the core device.
- Creates or updates the [Greengrass service role](#), if needed.
- Creates a Greengrass group and Greengrass core.
- (Optional) Creates a Hello World Lambda function, subscription, and local logging configuration.
- (Optional) Deploys the Greengrass group.

Modules 1 and 2

[Module 1](#) and [Module 2](#) describe how to set up your environment. (Or, use [Greengrass device setup](#) to run these modules for you.)

- Configure your core device for Greengrass.
- Run the dependency checker script.
- Create a Greengrass group and Greengrass core.
- Download and install the latest AWS IoT Greengrass Core software from a tar.gz file.
- Start the Greengrass daemon process on the core.

Note

AWS IoT Greengrass also provides other options for installing the AWS IoT Greengrass Core software, including apt installations on supported Debian platforms. For more information, see [the section called “Install the AWS IoT Greengrass Core software”](#).

Modules 3-1 and 3-2

[Module 3-1](#) and [Module 3-2](#) describe how to use local Lambda functions. (Or, use [Greengrass device setup](#) to run Module 3-1 for you.)

- Create Hello World Lambda functions in AWS Lambda.
- Add Lambda functions to your Greengrass group.
- Create subscriptions that allow MQTT communication between the Lambda functions and AWS IoT.
- Configure local logging for Greengrass system components and Lambda functions.
- Deploy a Greengrass group that contains your Lambda functions and subscriptions.
- Send messages from local Lambda functions to AWS IoT.
- Invoke local Lambda functions from AWS IoT.
- Test on-demand and long-lived functions.

Modules 4 and 5

[Module 4](#) shows how client devices connect to the core and communicate with each other.

[Module 5](#) shows how client devices can use shadows to control state.

- Register and provision AWS IoT devices (represented by command-line terminals).
- Install the AWS IoT Device SDK for Python. This is used by client devices to discover the Greengrass core.
- Add the client devices to your Greengrass group.
- Create subscriptions that allow MQTT communication.
- Deploy a Greengrass group that contains your client devices.
- Test device-to-device communication.
- Test shadow state updates.

Module 6

[Module 6](#) shows you how Lambda functions can access the AWS Cloud.

- Create a Greengrass group role that allows access to Amazon DynamoDB resources.
- Add a Lambda function to your Greengrass group. This function uses the AWS SDK for Python to interact with DynamoDB.
- Create subscriptions that allow MQTT communication.
- Test the interaction with DynamoDB.

Module 7

[Module 7](#) shows you how to configure a simulated hardware security module (HSM) for use with a Greengrass core.

Important

This advanced module is provided only for experimentation and initial testing. It is not for production use of any kind.

- Install and configure a software-based HSM and private key.
- Configure the Greengrass core to use hardware security.
- Test the hardware security configuration.

Requirements

To complete this tutorial, you need the following:

- A Mac, Windows PC, or UNIX-like system.
- An AWS account. If you don't have one, see [the section called "Create an AWS account"](#).
- The use of an AWS [Region](#) that supports AWS IoT Greengrass. For the list of supported regions for AWS IoT Greengrass, see [AWS endpoints and quotas](#) in the *AWS General Reference*.

Note

Make a note of your AWS Region and make sure that it is consistently used throughout this tutorial. If you switch your AWS Region during the tutorial, you might experience problems completing the steps.

- A Raspberry Pi 4 Model B, or Raspberry Pi 3 Model B/B+, with a 8 GB microSD card, or an Amazon EC2 instance. Because AWS IoT Greengrass should ideally be used with physical hardware, we recommend that you use a Raspberry Pi.

Note

Run the following command to get the model of your Raspberry Pi:

```
cat /proc/cpuinfo
```

Near the bottom of the listing, make a note of the value of the Revision attribute and then consult the [Which Pi have I got?](#) table. For example, if the value of Revision is a02082, the table shows the Pi is a 3 Model B.

Run the following command to determine the architecture of your Raspberry Pi:

```
uname -m
```

For this tutorial, the result should be greater than or equal to armv71.

- Basic familiarity with Python.

Although this tutorial is intended to run AWS IoT Greengrass on a Raspberry Pi, AWS IoT Greengrass also supports other platforms. For more information, see [the section called "Supported platforms and requirements"](#).

Create an AWS account

If you don't have an AWS account, follow these steps to create and activate an AWS account:

Sign up for an AWS account

If you do not have an AWS account, complete the following steps to create one.

To sign up for an AWS account

1. Open <https://portal.aws.amazon.com/billing/signup>.
2. Follow the online instructions.

Part of the sign-up procedure involves receiving a phone call or text message and entering a verification code on the phone keypad.

When you sign up for an AWS account, an *AWS account root user* is created. The root user has access to all AWS services and resources in the account. As a security best practice, assign administrative access to a user, and use only the root user to perform [tasks that require root user access](#).

AWS sends you a confirmation email after the sign-up process is complete. At any time, you can view your current account activity and manage your account by going to <https://aws.amazon.com/> and choosing **My Account**.

Create a user with administrative access

After you sign up for an AWS account, secure your AWS account root user, enable AWS IAM Identity Center, and create an administrative user so that you don't use the root user for everyday tasks.

Secure your AWS account root user

1. Sign in to the [AWS Management Console](#) as the account owner by choosing **Root user** and entering your AWS account email address. On the next page, enter your password.

For help signing in by using root user, see [Signing in as the root user](#) in the *AWS Sign-In User Guide*.

2. Turn on multi-factor authentication (MFA) for your root user.

For instructions, see [Enable a virtual MFA device for your AWS account root user \(console\)](#) in the *IAM User Guide*.

Create a user with administrative access

1. Enable IAM Identity Center.

For instructions, see [Enabling AWS IAM Identity Center](#) in the *AWS IAM Identity Center User Guide*.

2. In IAM Identity Center, grant administrative access to a user.

For a tutorial about using the IAM Identity Center directory as your identity source, see [Configure user access with the default IAM Identity Center directory](#) in the *AWS IAM Identity Center User Guide*.

Sign in as the user with administrative access

- To sign in with your IAM Identity Center user, use the sign-in URL that was sent to your email address when you created the IAM Identity Center user.

For help signing in using an IAM Identity Center user, see [Signing in to the AWS access portal](#) in the *AWS Sign-In User Guide*.

Assign access to additional users

1. In IAM Identity Center, create a permission set that follows the best practice of applying least-privilege permissions.

For instructions, see [Create a permission set](#) in the *AWS IAM Identity Center User Guide*.

2. Assign users to a group, and then assign single sign-on access to the group.

For instructions, see [Add groups](#) in the *AWS IAM Identity Center User Guide*.


Important

For this tutorial, we assume that your IAM user account has administrator access permissions.

Quick start: Greengrass device setup

Greengrass device setup is a script that sets up your core device in minutes, so that you can start using AWS IoT Greengrass. Use this script to:

1. Configure your device and install the AWS IoT Greengrass Core software.
2. Configure your cloud-based resources.
3. Optionally deploy a Greengrass group with a Hello World Lambda function that sends MQTT messages to AWS IoT from the AWS IoT Greengrass core. This sets up the Greengrass environment shown in the following diagram.



Requirements

Greengrass device setup has the following requirements:

- Your core device must use a [supported platform](#). The device must have an appropriate package manager installed: apt, yum, or opkg.
- The Linux user who runs the script must have permissions to run as sudo.

- You must provide your AWS account credentials. For more information, see [the section called “Provide AWS account credentials”](#).

Note

Greengrass device setup installs the [latest version](#) of the AWS IoT Greengrass Core software on the device. By installing the AWS IoT Greengrass Core software, you agree to the [Greengrass Core Software License Agreement](#).

Run Greengrass device setup

You can run Greengrass device setup in just a few steps. After you provide your AWS account credentials, the script provisions your Greengrass core device and deploys a Greengrass group in minutes. Run the following commands in a terminal window on the target device.

Note

These steps show you how to run the script in interactive mode, which prompts you to enter or accept each input value. For information about how to run the script silently, see [the section called “Run Greengrass device setup in silent mode”](#).

1. [Provide your credentials](#). In this procedure, we assume you provide temporary security credentials as environment variables.

```
export AWS_ACCESS_KEY_ID=AKIAIOSFODNN7EXAMPLE  
export AWS_SECRET_ACCESS_KEY=wJalrXUtnFEMI/K7MDENG/bPxrFiCYEXAMPLEKEY  
export AWS_SESSION_TOKEN=AQoDYXdzEJr1K...o50ytwEXAMPLE=
```

Note

If you're running Greengrass device setup on a Raspbian or OpenWrt platform, make a copy of these commands. You must provide them again after you reboot the device.

2. Download and start the script. You can use `wget` or `curl` to download the script.

`wget`:

```
wget -q -O ./gg-device-setup-latest.sh https://d1onfpft10uf5o.cloudfront.net/greengrass-device-setup/downloads/gg-device-setup-latest.sh && chmod +x ./gg-device-setup-latest.sh && sudo -E ./gg-device-setup-latest.sh bootstrap-greengrass-interactive
```

`curl`:

```
curl https://d1onfpft10uf5o.cloudfront.net/greengrass-device-setup/downloads/gg-device-setup-latest.sh > gg-device-setup-latest.sh && chmod +x ./gg-device-setup-latest.sh && sudo -E ./gg-device-setup-latest.sh bootstrap-greengrass-interactive
```

3. Proceed through the command prompts for [input values](#). You can press the **Enter** key to use the default value or type a custom value and then press **Enter**.

The script writes status messages to the terminal that are similar to the following.

```
##### Greengrass Device Setup v1.0.0 #####
[GreengrassDeviceSetup] The Greengrass Device Setup bootstrap log is available at: /tmp/greengrass-device-setup-bootstrap-1575933831.log
[GreengrassDeviceSetup] Using package management tool: yum...
[GreengrassDeviceSetup] Using runtime: python3.7...
[GreengrassDeviceSetup] Installing a dedicated pip for Greengrass Device Setup...
[GreengrassDeviceSetup] Validating and installing required dependencies...
[GreengrassDeviceSetup] The Greengrass Device Setup configuration is complete. Starting the Greengrass environment setup...
[GreengrassDeviceSetup] Forwarding command-line parameters: bootstrap-greengrass-interactive

[GreengrassDeviceSetup] Validating the device environment...
[GreengrassDeviceSetup] Validation of the device environment is complete.

[GreengrassDeviceSetup] Running the Greengrass environment setup...
[GreengrassDeviceSetup] The Greengrass environment setup is complete.

[GreengrassDeviceSetup] Configuring cloud-based Greengrass group management...
[GreengrassDeviceSetup] The Greengrass group configuration is complete.

[GreengrassDeviceSetup] Preparing the Greengrass core software...
[GreengrassDeviceSetup] The Greengrass core software is running.

[GreengrassDeviceSetup] Configuring the group deployment...
[GreengrassDeviceSetup] The group deployment is complete.
```

4. If your core device is running Raspbian or OpenWrt, reboot the device when prompted, provide your credentials, and then restart the script.
 - a. When prompted to reboot the device, run one of the following commands.

For Raspbian platforms:

```
sudo reboot
```

For OpenWrt platforms:

```
reboot
```

- b. After the device reboots, open the terminal and provide your credentials as environment variables.

```
export AWS_ACCESS_KEY_ID=AKIAIOSFODNN7EXAMPLE  
export AWS_SECRET_ACCESS_KEY=wJalrXUtnFEMI/K7MDENG/bPxrFiCYEXAMPLEKEY  
export AWS_SESSION_TOKEN=AQoDYXdzEJr1K...o50ytwEXAMPLE=
```

- c. Restart the script.

```
sudo -E ./gg-device-setup-latest.sh bootstrap-greengrass-interactive
```

- d. When prompted whether to use your input values from the previous session or start a new installation, enter yes to reuse your input values.

 Note

On platforms that require a reboot, your input values from the previous session, excluding credentials, are temporarily stored in the `GreengrassDeviceSetup.config.info` file.

When the setup is complete, the terminal displays a success status message that's similar to the following.

```

=====
Your device is running the Greengrass core software.
Your Greengrass group and Hello World Lambda function were deployed to the core device.

Setup information:

Device info: Linux-4.14.152-127.182.amzn2.x86_64-x86_64-with-glibc2.2.5
Greengrass core software location: /
Installed Greengrass core software version: 1.10.0
Greengrass core: arn:aws:iot:us-west-2:012345678910:thing/GreengrassDeviceSetup_Core_d46a0ea4-18ae-4376-8f44-4a504cdea608
Greengrass core IoT certificate: arn:aws:iot:us-west-2:012345678910:cert/23fbf0f4b6a5ea369f2b97f1a1b558180a240faa8e059ce19dc58f4a4c0d3b77
Greengrass core IoT certificate location: /greengrass/certs/23fbf0f4b6.cert.pem
Greengrass core IoT key location: /greengrass/certs/23fbf0f4b6.private.key
Deployed Greengrass group name: GreengrassDeviceSetup_Group_ee70f777-9af0-43b6-8612-a18b418e8b4a
Deployed Greengrass group ID: 6f5c8410-f3a6-43a2-acf3-33158e10fb8e
Deployed Greengrass group version: arn:aws:greengrass:us-west-2:012345678910:/greengrass/groups/6f5c8410-f3a6-43a2-acf3-33158e10fb8e/vers
Greengrass service role: arn:aws:iam::012345678910:role/GreengrassServiceRole_mui1v
GreengrassDeviceSetup log location: GreengrassDeviceSetup-20191209-232356.log
Deployed hello-world Lambda function: arn:aws:lambda:us-west-2:012345678910:function:Greengrass_HelloWorld_uNTf2:1
Hello-world subscriber topic: hello/world

You can now use the AWS IoT Console to subscribe
to the 'hello/world' topic to receive messages published from your
Greengrass core.
=====

```

5. Review the new Greengrass group that the script configures using the input values you provide.
 - a. Sign in to the [AWS Management Console](#) on your computer and open the AWS IoT console.
- Note**

Make sure that the AWS Region selected in the console is the same one that you used to configure your Greengrass environment. By default, the Region is US West (Oregon).
- b. In the navigation pane, expand **Greengrass devices**, then choose **Groups (V1)** to locate the newly created group.
 6. If you included the Hello World Lambda function, Greengrass device setup deploys the Greengrass group to your core device. To test the Lambda function, or for information about how to remove the Lambda function from the group, continue to [the section called “Verify the Lambda function is running on the core device”](#) in Module 3-1 of the Getting Started tutorial.

Note

Make sure that the AWS Region selected in the console is the same one that you used to configure your Greengrass environment. By default, the Region is US West (Oregon).

If you didn't include the Hello World Lambda function, you can [create your own Lambda function](#) or try other Greengrass features. For example, you can add the [Docker application deployment](#) connector to your group and use it to deploy Docker containers to your core device.

Troubleshooting issues

You can use the following information to troubleshoot issues with the AWS IoT Greengrass device setup.

Error: Python (python3.7) not found. Attempting to install it...

Solution: You might see this error when working with an Amazon EC2 instance. This error occurs when Python is not installed in the `/usr/bin/python3.7` folder. To resolve this error, move Python in the correct directory after installing it:

```
sudo ln -s /usr/local/bin/python3.7 /usr/bin/python3.7
```

Additional troubleshooting

To troubleshoot additional issues with the AWS IoT Greengrass device setup, you can look for debug information in the log files:

- For issues with the Greengrass device setup configuration, check the `/tmp/greengrass-device-setup-bootstrap-epoch-timestamp.log` file.
- For issues with the Greengrass group or core environment setup, check the `GreengrassDeviceSetup-date-time.log` file in the same directory as `gg-device-setup-latest.sh` or in the location you specified.

For more troubleshooting help, see [Troubleshooting](#) or check the [AWS IoT Greengrass tag on AWS re:Post](#).

Greengrass device setup configuration options

You configure Greengrass device setup to access your AWS resources and set up your Greengrass environment.

Provide AWS account credentials

Greengrass device setup uses your AWS account credentials to access your AWS resources. It supports long-term credentials for an IAM user or temporary security credentials from an IAM role.

First, get your credentials.

- To use long-term credentials, provide the access key ID and secret access key for your IAM user. For information about creating access keys for long-term credentials, see [Managing access keys for IAM users](#) in the *IAM User Guide*.
- To use temporary security credentials (recommended), provide the access key ID, secret access key, and session token from an assumed IAM role. For information about extracting temporary security credentials from the AWS STS `assume-role` command, see [Using temporary security credentials with the AWS CLI](#) in the *IAM User Guide*.

Note

For the purposes of this tutorial, we assume that the IAM user or IAM role has administrator access permissions.

Then, provide your credentials to Greengrass device setup in one of two ways:

- **As environment variables.** Set the `AWS_ACCESS_KEY_ID`, `AWS_SECRET_ACCESS_KEY`, and `AWS_SESSION_TOKEN` (if required) environment variables before you start the script, as shown in step 1 of [the section called “Run Greengrass device setup”](#).
- **As input values.** Enter your access key ID, secret access key, and session token (if required) values directly in the terminal after you start the script.

Greengrass device setup doesn't save or store your credentials.

Provide input values

In interactive mode, Greengrass device setup prompts you for input values. You can press the **Enter** key to use the default value or type a custom value and then press **Enter**. In silent mode, you provide input values after you start the script.

Input values

AWS access key ID

The access key ID from the long-term or temporary security credentials. Specify this option as an input value only if you don't provide your credentials as environment variables. For more information, see [the section called "Provide AWS account credentials"](#).

Option name for silent mode: `--aws-access-key-id`

AWS secret access key

The secret access key from the long-term or temporary security credentials. Specify this option as an input value only if you don't provide your credentials as environment variables. For more information, see [the section called "Provide AWS account credentials"](#).

Option name for silent mode: `--aws-secret-access-key`

AWS session token

The session token from the temporary security credentials. Specify this option as an input value only if you don't provide your credentials as environment variables. For more information, see [the section called "Provide AWS account credentials"](#).

Option name for silent mode: `--aws-session-token`

AWS Region

The AWS Region where you want to create the Greengrass group. For the list of supported AWS Regions, see [AWS IoT Greengrass](#) in the *Amazon Web Services General Reference*.

Default value: `us-west-2`

Option name for silent mode: `--region`

Group name

The name for the Greengrass group.

Default value: `GreengrassDeviceSetup_Group_`*guid*

Option name for silent mode: `--group-name`

Core name

The name for the Greengrass core. The core is an AWS IoT device (thing) that runs the AWS IoT Greengrass Core software. The core is added to the AWS IoT registry and the Greengrass group. If you provide a name, it must be unique in the AWS account and AWS Region.

Default value: `GreengrassDeviceSetup_Core_`*guid*

Option name for silent mode: `--core-name`

AWS IoT Greengrass Core software installation path

The location in the device file system where you want to install the AWS IoT Greengrass Core software.

Default value: `/`

Option name for silent mode: `--ggc-root-path`

Hello World Lambda function

Indicates whether to include a Hello World Lambda function in the Greengrass group. The function publishes an MQTT message to the `hello/world` topic every five seconds.

The script creates and publishes this user-defined Lambda function in AWS Lambda and adds it to your Greengrass group. The script also creates a subscription in the group that allows the function to send MQTT messages to AWS IoT.

Note

This is a Python 3.7 Lambda function. If Python 3.7 isn't installed on the device and the script is unable to install it, the script prints an error message in the terminal. To include the Lambda function in the group, you must install Python 3.7 manually and restart the

script. To create the Greengrass group without the Lambda function, restart the script and enter no when prompted to include the function.

Default value: no

Option name for silent mode: `--hello-world-lambda` - This option doesn't take a value. Include it in your command if you want to create the function.

Deployment timeout

The number of seconds before Greengrass device setup stops checking the status of the [Greengrass group deployment](#). This is used only when the group includes the Hello World Lambda function. Otherwise, the group is not deployed.

The deployment time depends on your network speed. For slow network speeds, you can increase this value.

Default value: 180

Option name for silent mode: `--deployment-timeout`

Log path

The location of the log file that contains information about Greengrass group and core setup operations. Use this log to troubleshoot deployment and other issues with the Greengrass group and core setup.

Default value: `./`

Option name for silent mode: `--log-path`

Verbosity

Indicates whether to print detailed log information in the terminal while the script runs. You can use this information to troubleshoot device setup.

Default value: no

Option name for silent mode: `--verbose` - This option doesn't take a value. Include it in your command if you want to print detailed log information.

Run Greengrass device setup in silent mode

You can run Greengrass device setup in silent mode so that the script doesn't prompt you for any values. To run in silent mode, specify `bootstrap-greengrass` mode and your [input values](#) after you start the script. You can omit input values if you want to use their defaults.

The procedure depends on whether you provide your AWS account credentials as environment variables before you start the script, or as input values after you start the script.

Provide credentials as environment variables

1. [Provide your credentials](#) as environment variables. The following example exports temporary credentials, which include the session token.

```
export AWS_ACCESS_KEY_ID=AKIAIOSFODNN7EXAMPLE
export AWS_SECRET_ACCESS_KEY=wJalrXUtnFEMI/K7MDENG/bPxRfiCYEXAMPLEKEY
export AWS_SESSION_TOKEN=AQoDYXdzEJr1K...o50ytwEXAMPLE=
```

Note

If you're running Greengrass device setup on a Raspbian or OpenWrt platform, make a copy of these commands. You must provide them again after you reboot the device.

2. Download and start the script. Provide input values as needed. For example:

- To use all default values:

```
wget -q -O ./gg-device-setup-latest.sh https://d1onfpft10uf5o.cloudfront.net/
greengrass-device-setup/downloads/gg-device-setup-latest.sh && chmod +x ./
gg-device-setup-latest.sh && sudo -E ./gg-device-setup-latest.sh bootstrap-
greengrass
```

- To specify custom values:

```
wget -q -O ./gg-device-setup-latest.sh https://d1onfpft10uf5o.cloudfront.net/
greengrass-device-setup/downloads/gg-device-setup-latest.sh && chmod +x ./
gg-device-setup-latest.sh && sudo -E ./gg-device-setup-latest.sh bootstrap-
greengrass
--region us-east-1
--group-name Custom_Group_Name
--core-name Custom_Core_Name
```

```
--ggc-root-path /custom/ggc/root/path
--deployment-timeout 300
--log-path /customized/log/path
--hello-world-lambda
--verbose
```

Note

To use `curl` to download the script, replace `wget -q -O` with `curl` in the command.

3. If your core device is running Raspbian or OpenWrt, reboot the device when prompted, provide your credentials, and then restart the script.
 - a. When prompted to reboot the device, run one of the following commands.

For Raspbian platforms:

```
sudo reboot
```

For OpenWrt platforms:

```
reboot
```

- b. After the device reboots, open the terminal and provide your credentials as environment variables.

```
export AWS_ACCESS_KEY_ID=AKIAIOSFODNN7EXAMPLE
export AWS_SECRET_ACCESS_KEY=wJalrXUtnFEMI/K7MDENG/bPxRfiCYEXAMPLEKEY
export AWS_SESSION_TOKEN=AQoDYXdzEJr1K...o50ytwEXAMPLE=
```

- c. Restart the script.

```
sudo -E ./gg-device-setup-latest.sh bootstrap-greengrass
```

- d. When prompted whether to use your input values from the previous session or start a new installation, enter `yes` to reuse your input values.

Note

On platforms that require a reboot, your input values from the previous session, excluding credentials, are temporarily stored in the `GreengrassDeviceSetup.config.info` file.

When the setup is complete, the terminal displays a success status message that's similar to the following.

```
=====
Your device is running the Greengrass core software.
Your Greengrass group and Hello World Lambda function were deployed to the core device.

Setup information:

Device info: Linux-4.14.152-127.182.amzn2.x86_64-x86_64-with-glibc2.2.5
Greengrass core software location: /
Installed Greengrass core software version: 1.10.0
Greengrass core: arn:aws:iot:us-west-2:012345678910:thing/GreengrassDeviceSetup_Core_d46a0ea4-18ae-4376-8f44-4a504cdea608
Greengrass core IoT certificate: arn:aws:iot:us-west-2:012345678910:cert/23fbf0f4b6a5ea369f2b97f1a1b558180a240faa8e059ce19dc58f4a4c0d3b77
Greengrass core IoT certificate location: /greengrass/certs/23fbf0f4b6.cert.pem
Greengrass core IoT key location: /greengrass/certs/23fbf0f4b6.private.key
Deployed Greengrass group name: GreengrassDeviceSetup_Group_ee70f777-9af0-43b6-8612-a18b418e8b4a
Deployed Greengrass group ID: 6f5c8410-f3a6-43a2-acf3-33158e10fb8e
Deployed Greengrass group version: arn:aws:greengrass:us-west-2:012345678910:/greengrass/groups/6f5c8410-f3a6-43a2-acf3-33158e10fb8e/vers1
Greengrass service role: arn:aws:iam::012345678910:role/GreengrassServiceRole_mu1lv
GreengrassDeviceSetup log location: GreengrassDeviceSetup-20191209-232356.log
Deployed hello-world Lambda function: arn:aws:lambda:us-west-2:012345678910:function:Greengrass_HelloWorld_uNTf2:1
Hello-world subscriber topic: hello/world

You can now use the AWS IoT Console to subscribe
to the 'hello/world' topic to receive messages published from your
Greengrass core.
=====
```

- If you included the Hello World Lambda function, Greengrass device setup deploys the Greengrass group to your core device. To test the Lambda function, or for information about how to remove the Lambda function from the group, continue to [the section called “Verify the Lambda function is running on the core device”](#) in Module 3-1 of the Getting Started tutorial.

Note

Make sure that the AWS Region selected in the console is the same one that you used to configure your Greengrass environment. By default, the Region is US West (Oregon).

If you didn't include the Hello World Lambda function, you can [create your own Lambda function](#) or try other Greengrass features. For example, you can add the [Docker application](#)

[deployment](#) connector to your group and use it to deploy Docker containers to your core device.

Provide credentials as input values

1. Download and start the script. [Provide your credentials](#) and any other input values that you want to specify. The following examples show how to provide temporary credentials, which include the session token.

- To use all default values:

```
wget -q -O ./gg-device-setup-latest.sh https://d1onfpft10uf5o.cloudfront.net/greengrass-device-setup/downloads/gg-device-setup-latest.sh && chmod +x ./gg-device-setup-latest.sh && sudo -E ./gg-device-setup-latest.sh bootstrap-greengrass
--aws-access-key-id AKIAIOSFODNN7EXAMPLE
--aws-secret-access-key wJalrXUtnFEMI/K7MDENG/bPxrFiCYEXAMPLEKEY
--aws-session-token AQoDYXdzEJr1K...o50ytwEXAMPLE=
```

- To specify custom values:

```
wget -q -O ./gg-device-setup-latest.sh https://d1onfpft10uf5o.cloudfront.net/greengrass-device-setup/downloads/gg-device-setup-latest.sh && chmod +x ./gg-device-setup-latest.sh && sudo -E ./gg-device-setup-latest.sh bootstrap-greengrass
--aws-access-key-id AKIAIOSFODNN7EXAMPLE
--aws-secret-access-key wJalrXUtnFEMI/K7MDENG/bPxrFiCYEXAMPLEKEY
--aws-session-token AQoDYXdzEJr1K...o50ytwEXAMPLE=
--region us-east-1
--group-name Custom_Group_Name
--core-name Custom_Core_Name
--ggc-root-path /custom/ggc/root/path
--deployment-timeout 300
--log-path /customized/log/path
--hello-world-lambda
--verbose
```

Note

If you're running Greengrass device setup on a Raspbian or OpenWrt platform, make a copy of your credentials. You must provide them again after you reboot the device. To use `curl` to download the script, replace `wget -q -O` with `curl` in the command.

2. If your core device is running Raspbian or OpenWrt, reboot the device when prompted, provide your credentials, and then restart the script.
 - a. When prompted to reboot the device, run one of the following commands.

For Raspbian platforms:

```
sudo reboot
```

For OpenWrt platforms:

```
reboot
```

- b. Restart the script. You must include your credentials in the command, but not the other input values. For example:

```
sudo -E ./gg-device-setup-latest.sh bootstrap-greengrass
--aws-access-key-id AKIAIOSFODNN7EXAMPLE
--aws-secret-access-key wJalrXUtnFEMI/K7MDENG/bPxRfiCYEXAMPLEKEY
--aws-session-token AQoDYXdzEJr1K...o50ytwEXAMPLE=
```

- c. When prompted whether to use your input values from the previous session or start a new installation, enter yes to reuse your input values.

Note

On platforms that require a reboot, your input values from the previous session, excluding credentials, are temporarily stored in the `GreengrassDeviceSetup.config.info` file.

When the setup is complete, the terminal displays a success status message that's similar to the following.

```
=====
Your device is running the Greengrass core software.
Your Greengrass group and Hello World Lambda function were deployed to the core device.

Setup information:

Device info: Linux-4.14.152-127.182.amzn2.x86_64-with-glibc2.2.5
Greengrass core software location: /
Installed Greengrass core software version: 1.10.0
Greengrass core: arn:aws:iot:us-west-2:012345678910:thing/GreengrassDeviceSetup_Core_d46a0ea4-18ae-4376-8f44-4a504cdea608
Greengrass core IoT certificate: arn:aws:iot:us-west-2:012345678910:cert/23fbf0f4b6a5ea369f2b97f1a1b558180a240faa8e059ce19dc58f4a4c0d3b77
Greengrass core IoT certificate location: /greengrass/certs/23fbf0f4b6.cert.pem
Greengrass core IoT key location: /greengrass/certs/23fbf0f4b6.private.key
Deployed Greengrass group name: GreengrassDeviceSetup_Group_ee70f777-9af0-43b6-8612-a18b418e8b4a
Deployed Greengrass group ID: 6f5c8410-f3a6-43a2-acf3-33158e10fb8e
Deployed Greengrass group version: arn:aws:greengrass:us-west-2:012345678910:/greengrass/groups/6f5c8410-f3a6-43a2-acf3-33158e10fb8e/vers
Greengrass service role: arn:aws:iam::012345678910:role/GreengrassServiceRole_mui1v
GreengrassDeviceSetup log location: GreengrassDeviceSetup-20191209-232356.log
Deployed hello-world Lambda function: arn:aws:lambda:us-west-2:012345678910:function:Greengrass_HelloWorld_uNTf2:1
Hello-world subscriber topic: hello/world

You can now use the AWS IoT Console to subscribe
to the 'hello/world' topic to receive messages published from your
Greengrass core.
=====
```

3. If you included the Hello World Lambda function, Greengrass device setup deploys the Greengrass group to your core device. To test the Lambda function, or for information about how to remove the Lambda function from the group, continue to [the section called “Verify the Lambda function is running on the core device”](#) in Module 3-1 of the Getting Started tutorial.

Note

Make sure that the AWS Region selected in the console is the same one that you used to configure your Greengrass environment. By default, the Region is US West (Oregon).

If you didn't include the Hello World Lambda function, you can [create your own Lambda function](#) or try other Greengrass features. For example, you can add the [Docker application deployment](#) connector to your group and use it to deploy Docker containers to your core device.

Module 1: Environment setup for Greengrass

This module shows you how to get an out-of-the-box Raspberry Pi, Amazon EC2 instance, or other device ready to be used by AWS IoT Greengrass as your AWS IoT Greengrass core device.

Tip

Or, to use a script that sets up your core device for you, see [the section called “Quick start: Greengrass device setup”](#).

This module should take less than 30 minutes to complete.

Before you begin, read the [requirements](#) for this tutorial. Then, follow the setup instructions in one of the following topics. Choose only the topic that applies to your core device type.

Topics

- [Setting up a Raspberry Pi](#)
- [Setting up an Amazon EC2 instance](#)
- [Setting up other devices](#)

Note

To learn how to use AWS IoT Greengrass running in a prebuilt Docker container, see [the section called “Run AWS IoT Greengrass in a Docker container”](#).

Setting up a Raspberry Pi

Follow the steps in this topic to set up a Raspberry Pi to use as an AWS IoT Greengrass core.

Tip

AWS IoT Greengrass also provides other options for installing the AWS IoT Greengrass Core software. For example, you can use [Greengrass device setup](#) to configure your environment and install the latest version of the AWS IoT Greengrass Core software. Or, on supported

Debian platforms, you can use the [APT package manager](#) to install or upgrade the AWS IoT Greengrass Core software. For more information, see [the section called “Install the AWS IoT Greengrass Core software”](#).

If you are setting up a Raspberry Pi for the first time, you must follow all of these steps. Otherwise, you can skip to [step 9](#). However, we recommend that you re-image your Raspberry Pi with the operating system as recommended in step 2.

1. Download and install an SD card formatter such as [SD Memory Card Formatter](#). Insert the SD card into your computer. Start the program and choose the drive where you have inserted your SD card. You can perform a quick format of the SD card.
2. Download the [Raspbian Buster](#) operating system as a zip file.
3. Using an SD card-writing tool (such as [Etcher](#)), follow the tool's instructions to flash the downloaded zip file onto the SD card. Because the operating system image is large, this step might take some time. Eject your SD card from your computer, and insert the microSD card into your Raspberry Pi.
4. For the first boot, we recommend that you connect the Raspberry Pi to a monitor (through HDMI), a keyboard, and a mouse. Next, connect your Pi to a micro USB power source and the Raspbian operating system should start up.
5. You might want to configure the Pi's keyboard layout before you continue. To do so, choose the Raspberry icon in the upper-right, choose **Preferences** and then choose **Mouse and Keyboard Settings**. Next, on the **Keyboard** tab, choose **Keyboard Layout**, and then choose an appropriate keyboard variant.
6. Next, [connect your Raspberry Pi to the internet through a Wi-Fi network](#) or an Ethernet cable.

Note


Connect your Raspberry Pi to the *same* network that your computer is connected to, and be sure that both your computer and Raspberry Pi have internet access before you continue. If you're in a work environment or behind a firewall, you might need to connect your Pi and your computer to the guest network to get both devices on the same network. However, this approach might disconnect your computer from local network resources, such as your intranet. One solution is to connect the Pi to the guest Wi-Fi network and to connect your computer to the guest Wi-Fi network *and* your local

network through an Ethernet cable. In this configuration, your computer should be able to connect to the Raspberry Pi through the guest Wi-Fi network and your local network resources through the Ethernet cable.

7. You must set up [SSH](#) on your Pi to remotely connect to it. On your Raspberry Pi, open a [terminal window](#) and run the following command:

```
sudo raspi-config
```

You should see the following:



Scroll down and choose **Interfacing Options** and then choose **P2 SSH**. When prompted, choose **Yes**. (Use the **Tab** key followed by **Enter**). SSH should now be enabled. Choose **OK**. Use the **Tab** key to choose **Finish** and then press **Enter**. If the Raspberry Pi doesn't reboot automatically, run the following command:

```
sudo reboot
```

8. On your Raspberry Pi, run the following command in the terminal:

```
hostname -I
```

This returns the IP address of your Raspberry Pi.

Note


For the following, if you receive an ECDSA key fingerprint message (Are you sure you want to continue connecting (yes/no)?), enter yes. The default password for the Raspberry Pi is **raspberrypi**.

If you are using macOS, open a terminal window and enter the following:


```
ssh pi@IP-address
```

IP-address is the IP address of your Raspberry Pi that you obtained by using the `hostname -I` command.

If you are using Windows, you need to install and configure [PuTTY](#). Expand **Connection**, choose **Data**, and make sure that **Prompt** is selected:




Next, choose **Session**, enter the IP address of the Raspberry Pi, and then choose **Open** using default settings.



If a PuTTY security alert is displayed, choose **Yes**.

The default Raspberry Pi login and password are **pi** and **raspberry**, respectively.



```
pi@raspberrypi: ~
login as: pi
pi@192.168.91.146's password:

The programs included with the Debian GNU/Linux system are free software;
the exact distribution terms for each program are described in the
individual files in /usr/share/doc/*/copyright.

Debian GNU/Linux comes with ABSOLUTELY NO WARRANTY, to the extent
permitted by applicable law.
Last login: Fri Dec 15 21:12:48 2017 from 192.168.93.69

SSH is enabled and the default password for the 'pi' user has not been changed.
This is a security risk - please login as the 'pi' user and type 'passwd' to set
a new password.

pi@raspberrypi:~ $
```

Note

If your computer is connected to a remote network using VPN, you might have difficulty connecting from the computer to the Raspberry Pi using SSH.

9. You are now ready to set up the Raspberry Pi for AWS IoT Greengrass. First, run the following commands from a local Raspberry Pi terminal window or an SSH terminal window:

Tip

AWS IoT Greengrass also provides other options for installing the AWS IoT Greengrass Core software. For example, you can use [Greengrass device setup](#) to configure your environment and install the latest version of the AWS IoT Greengrass Core software. Or, on supported Debian platforms, you can use the [APT package manager](#) to install or upgrade the AWS IoT Greengrass Core software. For more information, see [the section called "Install the AWS IoT Greengrass Core software"](#).

```
sudo adduser --system ggc_user
```

```
sudo addgroup --system ggc_group
```

10. To improve security on the Pi device, enable hardlink and softlink (symlink) protection on the operating system at startup.

- a. Navigate to the `98-rpi.conf` file.

```
cd /etc/sysctl.d
ls
```

 Note

If you don't see the `98-rpi.conf` file, follow the instructions in the `README.sysctl` file.

- b. Use a text editor (such as Leafpad, GNU nano, or vi) to add the following two lines to the end of the file. You might need to use the `sudo` command to edit as root (for example, `sudo nano 98-rpi.conf`).

```
fs.protected_hardlinks = 1
fs.protected_symlinks = 1
```

- c. Reboot the Pi.

```
sudo reboot
```

After about a minute, connect to the Pi using SSH and then run the following command to confirm the change:

```
sudo sysctl -a 2> /dev/null | grep fs.protected
```

You should see `fs.protected_hardlinks = 1` and `fs.protected_symlinks = 1`.

11. Edit your command line boot file to enable and mount memory cgroups. This allows AWS IoT Greengrass to set the memory limit for Lambda functions. Cgroups are also required to run AWS IoT Greengrass in the default [containerization](#) mode.

- a. Navigate to your boot directory.

```
cd /boot/
```

- b. Use a text editor to open `cmdline.txt`. Append the following to the end of the existing line, not as a new line. You might need to use the `sudo` command to edit as root (for example, `sudo nano cmdline.txt`).

```
cgroup_enable=memory cgroup_memory=1
```

- c. Now reboot the Pi.

```
sudo reboot
```

Your Raspberry Pi should now be ready for AWS IoT Greengrass.

12. Optional. Install the Java 8 runtime, which is required by [stream manager](#). This tutorial doesn't use stream manager, but it does use the **Default Group creation** workflow that enables stream manager by default. Use the following commands to install the Java 8 runtime on the core device, or disable stream manager before you deploy your group. Instructions for disabling stream manager are provided in Module 3.

```
sudo apt install openjdk-8-jdk
```

13. To make sure that you have all required dependencies, download and run the Greengrass dependency checker from the [AWS IoT Greengrass Samples](#) repository on GitHub. These commands unzip and run the dependency checker script in the Downloads directory.

Note

The dependency checker might fail if you are running version 5.4.51 of the Raspbian kernel. This version does not mount memory cgroups correctly. This might cause Lambda functions running in container mode to fail.

For more information on updating your kernel, see the [Cgroups not loaded after kernel upgrade](#) in the Raspberry Pi forums.

```
cd /home/pi/Downloads
mkdir greengrass-dependency-checker-GGCv1.11.x
```

```
cd greengrass-dependency-checker-GGCv1.11.x
wget https://github.com/aws-samples/aws-greengrass-samples/raw/master/greengrass-
dependency-checker-GGCv1.11.x.zip
unzip greengrass-dependency-checker-GGCv1.11.x.zip
cd greengrass-dependency-checker-GGCv1.11.x
sudo modprobe configs
sudo ./check_ggc_dependencies | more
```

Where more appears, press the **Spacebar** key to display another screen of text.

Important

This tutorial requires the Python 3.7 runtime to run local Lambda functions. When stream manager is enabled, it also requires the Java 8 runtime. If the `check_ggc_dependencies` script produces warnings about these missing runtime prerequisites, make sure to install them before you continue. You can ignore warnings about other missing optional runtime prerequisites.

For information about the **modprobe** command, run **man modprobe** in the terminal.

Your Raspberry Pi configuration is complete. Continue to [the section called “Module 2: Installing the AWS IoT Greengrass Core software”](#).

Setting up an Amazon EC2 instance


Follow the steps in this topic to set up an Amazon EC2 instance to use as your AWS IoT Greengrass core.

Tip

Or, to use a script that sets up your environment and installs the AWS IoT Greengrass Core software for you, see [the section called “Quick start: Greengrass device setup”](#).

Although you can complete this tutorial using an Amazon EC2 instance, AWS IoT Greengrass should ideally be used with physical hardware. We recommend that you [set up a Raspberry Pi](#) instead of using an Amazon EC2 instance when possible. If you're using a Raspberry Pi, you do not need to follow the steps in this topic.

1. Sign in to the [AWS Management Console](#) and launch an Amazon EC2 instance using an Amazon Linux AMI. For information about Amazon EC2 instances, see the [Amazon EC2 Getting Started Guide](#).
2. After your Amazon EC2 instance is running, enable port 8883 to allow incoming MQTT communications so that other devices can connect with the AWS IoT Greengrass core.
 - a. In the navigation pane of the Amazon EC2 console, choose **Security Groups**.



- b. Select the security group for the instance that you just launched, and then choose the **Inbound rules** tab.
 - c. Choose **Edit inbound rules**.

To enable port 8883, you add a custom TCP rule to the security group. For more information, see [Adding rules to a security group](#) in the *Amazon EC2 User Guide*.

- d. On the **Edit inbound rules** page, choose **Add rule**, enter the following settings, and then choose **Save**.

- For **Type**, choose **Custom TCP Rule**.
- For **Port range**, enter **8883**.
- For **Source**, choose **Anywhere**.
- For **Description**, enter **MQTT Communications**.

3. Connect to your Amazon EC2 instance.

- a. In the navigation pane, choose **Instances**, choose your instance, and then choose **Connect**.
- b. Follow the instructions on the **Connect To Your Instance** page to connect to your instance [by using SSH](#) and your private key file.

You can use [PuTTY](#) for Windows or Terminal for macOS. For more information, see [Connect to your Linux instance](#) in the *Amazon EC2 User Guide*.

You are now ready to set up your Amazon EC2 instance for AWS IoT Greengrass.

4. After you are connected to your Amazon EC2 instance, create the `ggc_user` and `ggc_group` accounts:

```
sudo adduser --system ggc_user
sudo groupadd --system ggc_group
```

 Note

If the `adduser` command isn't available on your system, use the following command.

```
sudo useradd --system ggc_user
```

5. To improve security, make sure that hardlink and softlink (symlink) protections are enabled on the operating system of the Amazon EC2 instance at startup.

Note

The steps for enabling hardlink and softlink protection vary by operating system. Consult the documentation for your distribution.

- a. Run the following command to check if hardlink and softlink protections are enabled:

```
sudo sysctl -a | grep fs.protected
```

If hardlinks and softlinks are set to 1, your protections are enabled correctly. Proceed to step 6.

Note

Softlinks are represented by `fs.protected_symlinks`.

- b. If hardlinks and softlinks are not set to 1, enable these protections. Navigate to your system configuration file.

```
cd /etc/sysctl.d
ls
```

- c. Using your favorite text editor (Leafpad, GNU nano, or vi), add the following two lines to the end of the system configuration file. On Amazon Linux 1, this is the `00-defaults.conf` file. On Amazon Linux 2, this is the `99-amazon.conf` file. You might need to change permissions (using the `chmod` command) to write to the file, or use the `sudo` command to edit as root (for example, `sudo nano 00-defaults.conf`).

```
fs.protected_hardlinks = 1
fs.protected_symlinks = 1
```

- d. Reboot the Amazon EC2 instance.

```
sudo reboot
```

After a few minutes, connect to your instance using SSH and then run the following command to confirm the change.

```
sudo sysctl -a | grep fs.protected
```

You should see that hardlinks and softlinks are set to 1.

6. Extract and run the following script to mount [Linux control groups](#) (cgroups). This allows AWS IoT Greengrass to set the memory limit for Lambda functions. Cgroups are also required to run AWS IoT Greengrass in the default [containerization](#) mode.

```
curl https://raw.githubusercontent.com/tianon/cgroupfs-mount/951c38ee8d802330454bdede20d85ec1c0f8d312/cgroupfs-mount > cgroupfs-mount.sh
chmod +x cgroupfs-mount.sh
sudo bash ./cgroupfs-mount.sh
```

Your Amazon EC2 instance should now be ready for AWS IoT Greengrass.

7. Optional. Install the Java 8 runtime, which is required by [stream manager](#). This tutorial doesn't use stream manager, but it does use the **Default Group creation** workflow that enables stream manager by default. Use the following commands to install the Java 8 runtime on the core device, or disable stream manager before you deploy your group. Instructions for disabling stream manager are provided in Module 3.

- For Debian-based distributions:

```
sudo apt install openjdk-8-jdk
```

- For Red Hat-based distributions:

```
sudo yum install java-1.8.0-openjdk
```

8. To make sure that you have all required dependencies, download and run the Greengrass dependency checker from the [AWS IoT Greengrass Samples](#) repository on GitHub. These commands download, unzip, and run the dependency checker script in your Amazon EC2 instance.

```
mkdir greengrass-dependency-checker-GGCv1.11.x
cd greengrass-dependency-checker-GGCv1.11.x
wget https://github.com/aws-samples/aws-greengrass-samples/raw/master/greengrass-dependency-checker-GGCv1.11.x.zip
unzip greengrass-dependency-checker-GGCv1.11.x.zip
cd greengrass-dependency-checker-GGCv1.11.x
```

```
sudo ./check_ggc_dependencies | more
```

Important

This tutorial requires the Python 3.7 runtime to run local Lambda functions. When stream manager is enabled, it also requires the Java 8 runtime. If the `check_ggc_dependencies` script produces warnings about these missing runtime prerequisites, make sure to install them before you continue. You can ignore warnings about other missing optional runtime prerequisites.

Your Amazon EC2 instance configuration is complete. Continue to [the section called “Module 2: Installing the AWS IoT Greengrass Core software”](#).

Setting up other devices

Follow the steps in this topic to set up a device (other than a Raspberry Pi) to use as your AWS IoT Greengrass core.

Tip

Or, to use a script that sets up your environment and installs the AWS IoT Greengrass Core software for you, see [the section called “Quick start: Greengrass device setup”](#).

If you're new to AWS IoT Greengrass, we recommend that you use a Raspberry Pi or an Amazon EC2 instance as your core device, and follow the [setup steps](#) appropriate for your device.

If you plan to build a custom Linux-based system using the Yocto Project, you can use the AWS IoT Greengrass Bitbake Recipe from the `meta-aws` project. This recipe also helps you develop a software platform that supports AWS edge software for embedded applications. The Bitbake build installs, configures, and automatically runs the AWS IoT Greengrass Core software on your device.

Yocto Project

An open source collaboration project that helps you build custom Linux-based systems for embedded applications regardless hardware architecture. For more information, see the [Yocto Project](#).

meta-aws

An AWS managed project that provides Yocto recipes. You can use the recipes to develop AWS edge software in Linux-based systems built with [OpenEmbedded](#) and Yocto Project. For more information about this community supported capability, see the [meta-aws](#) project on GitHub.

meta-aws-demos

An AWS managed project that contains demonstrations for the meta-aws project. For more examples about the integration process, see the [meta-aws-demos](#) project on GitHub.

To use a different device or [supported platform](#), follow the steps in this topic.

1. If your core device is an NVIDIA Jetson device, you must first flash the firmware with the JetPack 4.3 installer. If you're configuring a different device, skip to step 2.

Note

The JetPack installer version that you use is based on your target CUDA Toolkit version. The following instructions use JetPack 4.3 and CUDA Toolkit 10.0. For information about using the versions appropriate for your device, see [How to Install Jetpack](#) in the NVIDIA documentation.

- a. On a physical desktop that is running Ubuntu 16.04 or later, flash the firmware with the JetPack 4.3 installer, as described in [Download and Install JetPack](#) (4.3) in the NVIDIA documentation.

Follow the instructions in the installer to install all the packages and dependencies on the Jetson board, which must be connected to the desktop with a Micro-B cable.

- b. Reboot your board in normal mode, and connect a display to the board.

Note

When you use SSH to connect to the Jetson board, use the default user name (**nvidia**) and the default password (**nvidia**).

2. Run the following commands to create user `ggc_user` and group `ggc_group`. The commands you run differ, depending on the distribution installed on your core device.

- If your core device is running OpenWrt, run the following commands:

```
opkg install shadow-useradd
opkg install shadow-groupadd
useradd --system ggc_user
groupadd --system ggc_group
```

- Otherwise, run the following commands:

```
sudo adduser --system ggc_user
sudo addgroup --system ggc_group
```

 Note

If the `addgroup` command isn't available on your system, use the following command.

```
sudo groupadd --system ggc_group
```

3. Optional. Install the Java 8 runtime, which is required by [stream manager](#). This tutorial doesn't use stream manager, but it does use the **Default Group creation** workflow that enables stream manager by default. Use the following commands to install the Java 8 runtime on the core device, or disable stream manager before you deploy your group. Instructions for disabling stream manager are provided in Module 3.

- For Debian-based or Ubuntu-based distributions:

```
sudo apt install openjdk-8-jdk
```

- For Red Hat-based distributions:

```
sudo yum install java-1.8.0-openjdk
```

4. To make sure that you have all required dependencies, download and run the Greengrass dependency checker from the [AWS IoT Greengrass Samples](#) repository on GitHub. These commands unzip and run the dependency checker script.

```
mkdir greengrass-dependency-checker-GGCv1.11.x
```

```
cd greengrass-dependency-checker-GGCv1.11.x
wget https://github.com/aws-samples/aws-greengrass-samples/raw/master/greengrass-
dependency-checker-GGCv1.11.x.zip
unzip greengrass-dependency-checker-GGCv1.11.x.zip
cd greengrass-dependency-checker-GGCv1.11.x
sudo ./check_ggc_dependencies | more
```

Note

The `check_ggc_dependencies` script runs on AWS IoT Greengrass supported platforms and requires specific Linux system commands. For more information, see the dependency checker's [Readme](#).

5. Install all required dependencies on your device, as indicated by the dependency checker output. For missing kernel-level dependencies, you might have to recompile your kernel. For mounting Linux control groups (cgroups), you can run the [cgroups-mount](#) script. This allows AWS IoT Greengrass to set the memory limit for Lambda functions. Cgroups are also required to run AWS IoT Greengrass in the default [containerization](#) mode.

If no errors appear in the output, AWS IoT Greengrass should be able to run successfully on your device.

Important

This tutorial requires the Python 3.7 runtime to run local Lambda functions. When stream manager is enabled, it also requires the Java 8 runtime. If the `check_ggc_dependencies` script produces warnings about these missing runtime prerequisites, make sure to install them before you continue. You can ignore warnings about other missing optional runtime prerequisites.

For the list of AWS IoT Greengrass requirements and dependencies, see [the section called "Supported platforms and requirements"](#).

Module 2: Installing the AWS IoT Greengrass Core software

This module shows you how to install the AWS IoT Greengrass Core software on your chosen device. In this module, you first create a Greengrass group and core. Then, you download,

configure, and start the software on your core device. For more information about AWS IoT Greengrass Core software functionality, see [the section called “Configure the AWS IoT Greengrass core”](#).

Before you begin, make sure that you have completed the setup steps in [Module 1](#) for your chosen device.

Tip

AWS IoT Greengrass also provides other options for installing the AWS IoT Greengrass Core software. For example, you can use [Greengrass device setup](#) to configure your environment and install the latest version of the AWS IoT Greengrass Core software. Or, on supported Debian platforms, you can use the [APT package manager](#) to install or upgrade the AWS IoT Greengrass Core software. For more information, see [the section called “Install the AWS IoT Greengrass Core software”](#).

This module should take less than 30 minutes to complete.

Topics

- [Provision an AWS IoT thing to use as a Greengrass core](#)
- [Create an AWS IoT Greengrass group for the core](#)
- [Install and run AWS IoT Greengrass on the core device](#)

Provision an AWS IoT thing to use as a Greengrass core

Greengrass *cores* are devices that run the AWS IoT Greengrass Core software to manage local IoT processes. To set up a Greengrass core, you create an AWS IoT *thing*, which represents a device or logical entity that connects to AWS IoT. When you register a device as an AWS IoT thing, that device can use a digital certificate and keys that allow it to access AWS IoT. You use an [AWS IoT policy](#) to allow the device to communicate with the AWS IoT and AWS IoT Greengrass services.

In this section, you register your device as an AWS IoT thing to use it as a Greengrass core.

To create an AWS IoT thing

1. Navigate to the [AWS IoT console](#).

2. Under **Manage**, expand **All devices**, and then choose **Things**.
3. On the **Things** page, choose **Create things**.
4. On the **Create things** page, choose **Create single thing**, and then choose **Next**.
5. On the **Specify thing properties** page, do the following:
 - a. For **Thing name**, enter a name that represents your device, such as **MyGreengrassV1Core**.
 - b. Choose **Next**.
6. On the **Configure device certificate** page, choose **Next**.
7. On the **Attach policies to certificate** page, do one of the following:
 - Select an existing policy that grants permissions that cores require, and then choose **Create thing**.

A modal opens where you can download the certificates and keys that the device uses to connect to the AWS Cloud.

- Create an attach a new policy that grants core device permissions. Do the following:
 - a. Choose **Create policy**.

The **Create policy** page opens in a new tab.

- b. On the **Create policy** page, do the following:
 - i. For **Policy name**, enter a name that describes the policy, such as **GreengrassV1CorePolicy**.
 - ii. On the **Policy statements** tab, under **Policy document**, choose **JSON**.
 - iii. Enter the following policy document. This policy allows the core to communicate with the AWS IoT Core service, interact with device shadows, and communicate with the AWS IoT Greengrass service. For information about how to restrict this policy's access based on your use case, see [Minimal AWS IoT policy for the AWS IoT Greengrass core device](#).

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
```

```

        "iot:Publish",
        "iot:Subscribe",
        "iot:Connect",
        "iot:Receive"
    ],
    "Resource": [
        "*"
    ]
  },
  {
    "Effect": "Allow",
    "Action": [
      "iot:GetThingShadow",
      "iot:UpdateThingShadow",
      "iot:DeleteThingShadow"
    ],
    "Resource": [
        "*"
    ]
  },
  {
    "Effect": "Allow",
    "Action": [
      "greengrass:*"
    ],
    "Resource": [
        "*"
    ]
  }
]
}

```

- iv. Choose **Create** to create the policy.
- c. Return to the browser tab with the **Attach policies to certificate** page open. Do the following:
 - i. In the **Policies** list, select the policy that you created, such as **GreengrassV1CorePolicy**.

If you don't see the policy, choose the refresh button.

- ii. Choose **Create thing**.

A modal opens where you can download the certificates and keys that the core uses to connect to AWS IoT.

8. Return to the browser tab with the **Attach policies to certificate** page open. Do the following:

- a. In the **Policies** list, select the policy that you created, such as **GreengrassV1CorePolicy**.

If you don't see the policy, choose the refresh button.

- b. Choose **Create thing**.

A modal opens where you can download the certificates and keys that the core uses to connect to AWS IoT.

9. In the **Download certificates and keys** modal, download the device's certificates.

 Important

Before you choose **Done**, download the security resources.

Do the following:

- a. For **Device certificate**, choose **Download** to download the device certificate.
- b. For **Public key file**, choose **Download** to download the public key for the certificate.
- c. For **Private key file**, choose **Download** to download the private key file for the certificate.
- d. Review [Server Authentication](#) in the *AWS IoT Developer Guide* and choose the appropriate root CA certificate. We recommend that you use Amazon Trust Services (ATS) endpoints and ATS root CA certificates. Under **Root CA certificates**, choose **Download** for a root CA certificate.
- e. Choose **Done**.

Make a note of the certificate ID that's common in the file names for the device certificate and keys. You need it later.

Create an AWS IoT Greengrass group for the core

AWS IoT Greengrass *groups* contain settings and other information about its components, such as client devices, Lambda functions, and connectors. A group defines the configuration for a core, including how its components can interact with each other.

In this section, you create a group for your core.

Tip

For an example that uses the AWS IoT Greengrass API to create and deploy a group, see the [gg_group_setup](#) repository on GitHub.

To create a group for the core

1. Navigate to the [AWS IoT console](#).
2. Under **Manage**, expand **Greengrass devices**, and choose **Groups (V1)**.

Note

If you don't see the **Greengrass devices** menu, change to an AWS Region that supports AWS IoT Greengrass V1. For the list of supported Regions, see [AWS IoT Greengrass V1 endpoints and quotas](#) in the *AWS General Reference*. You must [create the AWS IoT thing for your core](#) in a Region where AWS IoT Greengrass V1 is available.

3. On the **Greengrass groups** page, choose **Create group**.
4. On the **Create Greengrass group** page, do the following:
 - a. For **Greengrass group name**, enter a name that describes the group, such as **MyGreengrassGroup**.
 - b. For **Greengrass core**, choose the AWS IoT thing that you created earlier, such as **MyGreengrassV1Core**.

The console automatically selects the thing's device certificate for you.

- c. Choose **Create group**.

Install and run AWS IoT Greengrass on the core device

Note

This tutorial provides instructions for you to run the AWS IoT Greengrass Core software on a Raspberry Pi, but you can use any supported device.

In this section, you configure, install, and run the AWS IoT Greengrass Core software on your core device.

To install and run AWS IoT Greengrass

1. From the [AWS IoT Greengrass Core software](#) section in this guide, download the AWS IoT Greengrass Core software installation package. Choose the package that best fits the CPU architecture, distribution, and OS of your core device.
 - For Raspberry Pi, download the package for the Armv7l architecture and Linux operating system.
 - For an Amazon EC2 instance, download the package for the x86_64 architecture and Linux operating system.
 - For NVIDIA Jetson TX2, download the package for the Armv8 (AArch64) architecture and Linux operating system.
 - For Intel Atom, download the package for the x86_64 architecture and Linux operating system.
2. In previous steps, you downloaded five files to your computer:
 - `greengrass-OS-architecture-1.11.6.tar.gz` – This compressed file contains the AWS IoT Greengrass Core software that runs on the core device.
 - `certificateId-certificate.pem.crt` – The device certificate file.
 - `certificateId-public.pem.key` – The device certificate's public key file.
 - `certificateId-private.pem.key` – The device certificate's private key file.
 - `AmazonRootCA1.pem` – The Amazon root certificate authority (CA) file.

In this step, you transfer these files from your computer to your core device. Do the following:

- a. If you don't know the IP address of your Greengrass core device, open a terminal on the core device and run the following command.

 Note

This command might not return the correct IP address for some devices. Consult the documentation for your device to retrieve your device IP address.

```
hostname -I
```

- b. Transfer these files from your computer to your core device. The file transfer steps vary depending on the operating system of your computer. Choose your operating system for steps that show how to transfer files to your Raspberry Pi device.

 Note

For a Raspberry Pi, the default user name is **pi** and the default password is **raspberrypi**.

For an NVIDIA Jetson TX2, the default user name is **nvidia** and the default password is **nvidia**.

Windows

To transfer the compressed files from your computer to a Raspberry Pi core device, use a tool such as [WinSCP](#) or the [PuTTY](#) **pscp** command. To use the **pscp** command, open a Command Prompt window on your computer and run the following:

```
cd path-to-downloaded-files
pscp -pw Pi-password greengrass-OS-architecture-1.11.6.tar.gz pi@IP-address:/home/pi
pscp -pw Pi-password certificateId-certificate.pem.crt pi@IP-address:/home/pi
pscp -pw Pi-password certificateId-public.pem.key pi@IP-address:/home/pi
pscp -pw Pi-password certificateId-private.pem.key pi@IP-address:/home/pi
pscp -pw Pi-password AmazonRootCA1.pem pi@IP-address:/home/pi
```

Note

The version number in this command must match the version of your AWS IoT Greengrass Core software package.

macOS

To transfer the compressed files from your Mac to a Raspberry Pi core device, open a Terminal window on your computer and run the following commands. The *path-to-downloaded-files* is typically ~/Downloads.

Note

You might be prompted for two passwords. If so, the first password is for the Mac's sudo command and the second is the password for the Raspberry Pi.

```
cd path-to-downloaded-files
scp greengrass-OS-architecture-1.11.6.tar.gz pi@IP-address:/home/pi
scp certificateId-certificate.pem.crt pi@IP-address:/home/pi
scp certificateId-public.pem.key pi@IP-address:/home/pi
scp certificateId-private.pem.key pi@IP-address:/home/pi
scp AmazonRootCA1.pem pi@IP-address:/home/pi
```

Note

The version number in this command must match the version of your AWS IoT Greengrass Core software package.

UNIX-like system

To transfer the compressed files from your computer to a Raspberry Pi core device, open a terminal window on your computer and run the following commands:

```
cd path-to-downloaded-files
scp greengrass-OS-architecture-1.11.6.tar.gz pi@IP-address:/home/pi
```

```
scp certificateId-certificate.pem.crt pi@IP-address:/home/pi  
scp certificateId-public.pem.key pi@IP-address:/home/pi  
scp certificateId-private.pem.key pi@IP-address:/home/pi  
scp AmazonRootCA1.pem pi@IP-address:/home/pi
```

Note

The version number in this command must match the version of your AWS IoT Greengrass Core software package.

Raspberry Pi web browser

If you used the Raspberry Pi's web browser to download the compressed files, the files should be in the Pi's ~/Downloads folder, such as /home/pi/Downloads. Otherwise, the compressed files should be in the Pi's ~ folder, such as /home/pi.

3. On the Greengrass core device, open a terminal, and navigate to the folder that contains the AWS IoT Greengrass Core software and certificates. Replace *path-to-transferred-files* with the path where you transferred the files on the core device. For example, on a Raspberry Pi, run `cd /home/pi`.

```
cd path-to-transferred-files
```

4. Unpack the AWS IoT Greengrass Core software on the core device. Run the following command to unpack the software archive that you transferred to the core device. This command uses the `-C /` argument to create the /greengrass folder in the root folder of the core device.

```
sudo tar -xzf greengrass-OS-architecture-1.11.6.tar.gz -C /
```

Note

The version number in this command must match the version of your AWS IoT Greengrass Core software package.

5. Move the certificates and keys to the AWS IoT Greengrass Core software folder. Run the following commands to create a folder for certificates and move the certificates and keys to it. Replace *path-to-transferred-files* with the path where you transferred the files

on the core device, and replace *certificateId* with the certificate ID in the file names. For example, on a Raspberry Pi, replace *path-to-transferred-files* with **/home/pi**

```
sudo mv path-to-transferred-files/certificateId-certificate.pem.crt /greengrass/certs
sudo mv path-to-transferred-files/certificateId-public.pem.key /greengrass/certs
sudo mv path-to-transferred-files/certificateId-private.pem.key /greengrass/certs
sudo mv path-to-transferred-files/AmazonRootCA1.pem /greengrass/certs
```

6. The AWS IoT Greengrass Core software uses a configuration file that specifies parameters for the software. This configuration file specifies the file paths for certificate files and the AWS Cloud endpoints to use. In this step, you create the AWS IoT Greengrass Core software configuration file for your core. Do the following:
 - a. Get the Amazon Resource Name (ARN) for your core's AWS IoT thing. Do the following:
 - i. In the [AWS IoT console](#), under **Manage**, under **Greengrass devices**, choose **Groups (V1)**.
 - ii. On the **Greengrass groups** page, choose the group that you created earlier.
 - iii. Under **Overview**, choose **Greengrass core**.
 - iv. On the core details page, copy the **AWS IoT thing ARN**, and save it to use in the AWS IoT Greengrass Core configuration file.
 - b. Get the AWS IoT device data endpoint for your AWS account in the current Region. Devices use this endpoint to connect to AWS as AWS IoT things. Do the following:
 - i. In the [AWS IoT console](#), choose **Settings**.
 - ii. Under **Device data endpoint**, copy the **Endpoint**, and save it to use in the AWS IoT Greengrass Core configuration file.
 - c. Create the AWS IoT Greengrass Core software configuration file. For example, you can run the following command to use GNU nano to create the file.

```
sudo nano /greengrass/config/config.json
```

Replace the contents of the file with the following JSON document.

```
{
  "coreThing" : {
    "caPath": "AmazonRootCA1.pem",
```

```

    "certPath": "certificateId-certificate.pem.crt",
    "keyPath": "certificateId-private.pem.key",
    "thingArn": "arn:aws:iot:region:account-id:thing/MyGreengrassV1Core",
    "iotHost": "device-data-prefix-ats.iot.region.amazonaws.com",
    "ggHost": "greengrass-ats.iot.region.amazonaws.com",
    "keepAlive": 600
  },
  "runtime": {
    "cgroup": {
      "useSystemd": "yes"
    }
  },
  "managedRespawn": false,
  "crypto": {
    "caPath": "file:///greengrass/certs/AmazonRootCA1.pem",
    "principals": {
      "SecretsManager": {
        "privateKeyPath": "file:///greengrass/certs/certificateId-private.pem.key"
      },
      "IoTCertificate": {
        "privateKeyPath": "file:///greengrass/certs/certificateId-private.pem.key",
        "certificatePath": "file:///greengrass/certs/certificateId-certificate.pem.crt"
      }
    }
  }
}

```

Then, do the following:

- If you downloaded a different Amazon root CA certificate than Amazon Root CA 1, replace each instance of *AmazonRootCA1.pem* with the name of the Amazon root CA file.
- Replace each instance of *certificateId* with the certificate ID in the name of the certificate and key files.
- Replace *arn:aws:iot:region:account-id:thing/MyGreengrassV1Core* with the ARN of your core's thing that you saved earlier.
- Replace *MyGreengrassV1core* with the name of your core's thing.

- Replace *device-data-prefix-ats.iot.region.amazonaws.com* with the AWS IoT device data endpoint that you saved earlier.
- Replace *region* with your AWS Region.

For more information about the configuration options that you can specify in this configuration file, see [AWS IoT Greengrass core configuration file](#).

7. Make sure that your core device is connected to the internet. Then, start AWS IoT Greengrass on your core device.

```
cd /greengrass/ggc/core/  
sudo ./greengrassd start
```

You should see a `Greengrass successfully started` message. Make a note of the PID.

 Note

To set up your core device to start AWS IoT Greengrass on system boot, see [the section called “Start Greengrass on system boot”](#).

You can run the following command to confirm that the AWS IoT Greengrass Core software (Greengrass daemon) is functioning. Replace *PID-number* with your PID:

```
ps aux | grep PID-number
```

You should see an entry for the PID with a path to the running Greengrass daemon (for example, `/greengrass/ggc/packages/1.11.6/bin/daemon`). If you run into issues starting AWS IoT Greengrass, see [Troubleshooting](#).

Module 3 (part 1): Lambda functions on AWS IoT Greengrass

This module shows you how to create and deploy a Lambda function that sends MQTT messages from your AWS IoT Greengrass core device. The module describes Lambda function configurations, subscriptions used to allow MQTT messaging, and deployments to a core device.

[Module 3 \(Part 2\)](#) covers the differences between on-demand and long-lived Lambda functions running on the AWS IoT Greengrass core.

Before you begin, make sure that you have completed [Module 1](#) and [Module 2](#) and have a running AWS IoT Greengrass core device.

 Tip

Or, to use a script that sets up your core device for you, see [the section called “Quick start: Greengrass device setup”](#). The script can also create and deploy the Lambda function used in this module.

This module should take about 30 minutes to complete.

Topics

- [Create and package a Lambda function](#)
- [Configure the Lambda function for AWS IoT Greengrass](#)
- [Deploy cloud configurations to a Greengrass core device](#)
- [Verify the Lambda function is running on the core device](#)

Create and package a Lambda function

The example Python Lambda function in this module uses the [AWS IoT Greengrass Core SDK](#) for Python to publish MQTT messages.

In this step, you:


- Download the AWS IoT Greengrass Core SDK for Python to your computer (not the AWS IoT Greengrass core device).
- Create a Lambda function deployment package that contains the function code and dependencies.
- Use the Lambda console to create a Lambda function and upload the deployment package.
- Publish a version of the Lambda function and create an alias that points to the version.

To complete this module, Python 3.7 must be installed on your core device.

1. From the [AWS IoT Greengrass Core SDK](#) downloads page, download the AWS IoT Greengrass Core SDK for Python to your computer.
2. Unzip the downloaded package to get the Lambda function code and the SDK.

The Lambda function in this module uses:

- The `greengrassHelloWorld.py` file in `examples\HelloWorld`. This is your Lambda function code. Every five seconds, the function publishes one of two possible messages to the `hello/world` topic.
 - The `greengrasssdk` folder. This is the SDK.
3. Copy the `greengrasssdk` folder into the `HelloWorld` folder that contains `greengrassHelloWorld.py`.
 4. To create the Lambda function deployment package, save `greengrassHelloWorld.py` and the `greengrasssdk` folder to a compressed zip file named `hello_world_python_lambda.zip`. The py file and `greengrasssdk` folder must be in the root of the directory.



On UNIX-like systems (including the Mac terminal), you can use the following command to package the file and folder:

```
zip -r hello_world_python_lambda.zip greengrasssdk greengrassHelloWorld.py
```

Note

Depending on your distribution, you might need to install `zip` first (for example, by running `sudo apt-get install zip`). The installation command for your distribution might be different.


Now you're ready to create your Lambda function and upload the deployment package.

5. Open the Lambda console and choose **Create function**.

6. Choose **Author from scratch**.
7. Name your function **Greengrass_HelloWorld**, and set the remaining fields as follows:
 - For **Runtime**, choose **Python 3.7**.
 - For **Permissions**, keep the default setting. This creates an execution role that grants basic Lambda permissions. This role isn't used by AWS IoT Greengrass.

Choose **Create function**.

8. Upload your Lambda function deployment package:
 - a. On the **Code** tab, under **Code source**, choose **Upload from**. From the dropdown, choose **.zip file**.



- b. Choose **Upload**, and then choose your `hello_world_python_lambda.zip` deployment package. Then, choose **Save**.
 - c. On the **Code** tab for the function, under **Runtime settings**, choose **Edit**, and then enter the following values.
 - For **Runtime**, choose **Python 3.7**.
 - For **Handler**, enter `greengrassHelloWorld.function_handler`




- d. Choose **Save**.

Note

The **Test** button on the AWS Lambda console doesn't work with this function. The AWS IoT Greengrass Core SDK doesn't contain modules that are required to run your Greengrass Lambda functions independently in the AWS Lambda console. These modules (for example, `greengrass_common`) are supplied to the functions after they are deployed to your Greengrass core.

9. Publish the Lambda function:

- a. From the **Actions** menu at the top of the page, choose **Publish new version**.



- b. For **Version description**, enter **First version**, and then choose **Publish**.

Publish new version from \$LATEST ✕

Publishing a new version will save a "snapshot" of the code and configuration of the \$LATEST version. You will be unable to edit the new version's code. Please click to confirm.

Version description


Cancel **Publish**

10. Create an [alias](#) for the Lambda function [version](#):

Note

Greengrass groups can reference a Lambda function by alias (recommended) or by version. Using an alias makes it easier to manage code updates because you don't have to change your subscription table or group definition when the function code is updated. Instead, you just point the alias to the new function version.

- a. From the **Actions** menu at the top of the page, choose **Create alias**.



- b. Name the alias **GG_HelloWorld**, set the version to **1** (which corresponds to the version that you just published), and then choose **Save**.

Note

AWS IoT Greengrass doesn't support Lambda aliases for **\$LATEST** versions.

Create alias

Alias configuration

An alias is a pointer to one or two versions. Choose each version that you want the alias to point to.

Name

Description - optional

Version

► Weighted alias

Cancel Save

Configure the Lambda function for AWS IoT Greengrass

You are now ready to configure your Lambda function for AWS IoT Greengrass.

In this step, you:

- Use the AWS IoT console to add the Lambda function to your Greengrass group.
- Configure group-specific settings for the Lambda function.
- Add a subscription to the group that allows the Lambda function to publish MQTT messages to AWS IoT.
- Configure local log settings for the group.

1. In the AWS IoT console navigation pane, under **Manage**, expand **Greengrass devices**, and then choose **Groups (V1)**.
2. Under **Greengrass groups**, choose the group that you created in [Module 2](#).

3. On the group configuration page, choose the **Lambda functions** tab, and then scroll down to the **My Lambda functions** section and choose **Add Lambda function**.
4. Select the name of the Lambda function you created in the previous step (**Greengrass_HelloWorld**, not the alias name).
5. For the version, choose **Alias: GG_HelloWorld**.
6. In the **Lambda function configuration** section, make the following changes:
 - Set the **System user and group** to **Use group default**.
 - Set the **Lambda function containerization** to **Use group default**.
 - Set **Timeout** to 25 seconds. This Lambda function sleeps for 5 seconds before each invocation.
 - For **Pinned**, choose **True**.

 Note

A *long-lived* (or *pinned*) Lambda function starts automatically after AWS IoT Greengrass starts and keeps running in its own container. This is in contrast to an *on-demand* Lambda function, which starts when invoked and stops when there are no tasks left to run. For more information, see [the section called “Lifecycle configuration”](#).

7. Choose **Add Lambda function** to save your changes. For information about Lambda function properties, see [the section called “Controlling Greengrass Lambda function execution”](#).

Next, create a subscription that allows the Lambda function to send [MQTT](#) messages to AWS IoT Core.

A Greengrass Lambda function can exchange MQTT messages with:

- [Devices](#) in the Greengrass group.
- [Connectors](#) in the group.
- Other Lambda functions in the group.
- AWS IoT Core.
- The local shadow service. For more information, see [the section called “Module 5: Interacting with device shadows”](#).

The group uses subscriptions to control how these entities can communicate with each other. Subscriptions provide predictable interactions and a layer of security.

A subscription consists of a source, target, and topic. The source is the originator of the message. The target is the destination of the message. The topic allows you to filter the data that is sent from the source to the target. The source or target can be a Greengrass device, Lambda function, connector, device shadow, or AWS IoT Core.

 Note

A subscription is directed in the sense that messages flow in a specific direction: from the source to the target. To allow two-way communication, you must set up two subscriptions.

 Note

Currently, the subscription topic filter does not allow more than a single + character in a topic. The topic filter only allows a single # character at the end of a topic.

The Greengrass_HelloWorld Lambda function sends messages only to the hello/world topic in AWS IoT Core, so you only need to create one subscription from the Lambda function to AWS IoT Core. You create this in the next step.

8. On the group configuration page, choose the **Subscriptions** tab, and then choose **Add subscription**.

For an example that shows you how to create a subscription using the AWS CLI, see [create-subscription-definition](#) in the *AWS CLI Command Reference*.

9. In the **Source type**, choose **Lambda function** and, for the **Source**, choose **Greengrass_HelloWorld**.
10. For the **Target type**, choose **Service** and, for the **Target** select **IoT Cloud**.
11. For **Topic filter**, enter **hello/world**, and then choose **Create subscription**.

12. Configure the group's logging settings. For this tutorial, you configure AWS IoT Greengrass system components and user-defined Lambda functions to write logs to the file system of the core device.
 - a. On the group configuration page, choose the **Logs** tab.
 - b. In the **Local logs configuration** section, choose **Edit**.
 - c. On the **Edit local logs configuration** dialog box, keep the default values for both log levels and storage sizes, and then choose **Save**.

You can use logs to troubleshoot any issues you might encounter when running this tutorial. When troubleshooting issues, you can temporarily change the logging level to **Debug**. For more information, see [the section called "Accessing file system logs"](#).

13. If the Java 8 runtime isn't installed on your core device, you must install it or disable stream manager.

 Note

This tutorial doesn't use stream manager, but it does use the **Default Group creation** workflow that enables stream manager by default. If stream manager is enabled but Java 8 isn't installed, the group deployment fails. For more information, see the [stream manager requirements](#).

To disable stream manager:

- a. On the group settings page, choose the **Lambda functions** tab.
- b. Under the **System Lambda functions** section, select **Stream manager** and choose **Edit**.
- c. Choose **Disable**, and then choose **Save**.

Deploy cloud configurations to a Greengrass core device

1. Make sure that your Greengrass core device is connected to the internet. For example, try successfully navigating to a webpage.
2. Make sure that the Greengrass daemon is running on your core device. In your core device terminal, run the following commands to check whether the daemon is running and start it, if needed.

- a. To check whether the daemon is running:

```
ps aux | grep -E 'greengrass.*daemon'
```

If the output contains a root entry for `/greengrass/ggc/packages/1.11.6/bin/daemon`, then the daemon is running.

- b. To start the daemon:

```
cd /greengrass/ggc/core/  
sudo ./greengrassd start
```

Now you're ready to deploy the Lambda function and subscription configurations to your Greengrass core device.

3. In the AWS IoT console navigation pane, under **Manage**, expand **Greengrass devices**, and then choose **Groups (V1)**.
4. Under **Greengrass groups**, choose the group that you created in [Module 2](#).
5. On the group configuration page, choose **Deploy**.
6. On the **Lambda functions** tab, in the **System Lambda functions** section, choose **IP detector**.
7. Choose **Edit** and select **Automatically detect and override MQTT broker endpoints**. This enables devices to automatically acquire connectivity information for the core, such as IP address, DNS, and port number. Automatic detection is recommended, but AWS IoT Greengrass also supports manually specified endpoints. You're only prompted for the discovery method the first time that the group is deployed.

The first deployment might take a few minutes. When the deployment is complete, you should see **Successfully completed** in the **Status** column on the **Deployments** page:

 Note

The deployment status is also displayed below the group's name on the page header.

For troubleshooting help, see [Troubleshooting](#).

Verify the Lambda function is running on the core device

1. From the navigation pane of the [AWS IoT console](#), under **Test**, choose **MQTT test client**.
2. Choose the **Subscribe to topic** tab.
3. Enter **hello/world** into the **Topic filter** and expand the **Additional configuration**.
4. Enter the information listed in each of the following fields:
 - For **Quality of Service**, choose **0**.
 - For **MQTT payload display**, choose **Display payloads as strings**.
5. Choose **Subscribe**.

Assuming the Lambda function is running on your device, it publishes messages similar to the following to the hello/world topic:

The screenshot shows the AWS IoT console interface for the MQTT test client. On the left, under the 'Subscriptions' tab, the 'hello/world' subscription is selected. On the right, the 'hello/world' topic is displayed with buttons for 'Pause', 'Clear', 'Export', and 'Edit'. Below the topic name, a message is shown with a timestamp of 'April 29, 2021, 17:35:40 (UTC-0400)'. The message content is a JSON object:

```
{  "message": "Hello world! Sent from Greengrass Core running on platform: Linux-4.9.30-v7+-armv7l-with-debian-8.0"}
```

Although the Lambda function continues to send MQTT messages to the hello/world topic, don't stop the AWS IoT Greengrass daemon. The remaining modules are written with the assumption that it's running.

You can delete the function and subscription from the group:

- On the groups configuration page, under the **Lambda functions** tab, select the Lambda function you want to remove and choose **Remove**.
- On the groups configuration page, under the **Subscriptions** tab, choose the subscription, and then choose **Delete**.

The function and subscription are removed from the core during the next group deployment.

Module 3 (part 2): Lambda functions on AWS IoT Greengrass

This module explores the differences between on-demand and long-lived Lambda functions running on the AWS IoT Greengrass core.

Before you begin, run the [Greengrass Device Setup](#) script or make sure you have completed [Module 1](#), [Module 2](#), and [Module 3 \(Part 1\)](#).

This module should take about 30 minutes to complete.

Topics

- [Create and package the Lambda function](#)
- [Configure long-lived Lambda functions for AWS IoT Greengrass](#)
- [Test long-lived Lambda functions](#)
- [Test on-demand Lambda functions](#)

Create and package the Lambda function

In this step, you:

- Create a Lambda function deployment package that contains the function code and dependencies.
- Use the Lambda console to create a Lambda function and upload the deployment package.
- Publish a version of the Lambda function and create an alias that points to the version.


1. On your computer, go to the AWS IoT Greengrass Core SDK for Python that you downloaded and extracted in [the section called “Create and package a Lambda function”](#) in Module 3-1.

The Lambda function in this module uses:

- The `greengrassHelloWorldCounter.py` file in `examples\HelloWorldCounter`. This is your Lambda function code.
- The `greengrasssdk` folder. This is the SDK.

2. Create a Lambda function deployment package:

- a. Copy the `greengrasssdk` folder into the `HelloWorldCounter` folder that contains `greengrassHelloWorldCounter.py`.
- b. Save `greengrassHelloWorldCounter.py` and the `greengrasssdk` folder to a zip file named `hello_world_counter_python_lambda.zip`. The py file and `greengrasssdk` folder must be in the root of the directory.



On UNIX-like systems (including the Mac terminal) that have `zip` installed, you can use the following command to package the file and folder:

```
zip -r hello_world_counter_python_lambda.zip greengrasssdk
greengrassHelloWorldCounter.py
```

Now you're ready to create your Lambda function and upload the deployment package.

3. Open the Lambda console and choose **Create function**.
4. Choose **Author from scratch**.
5. Name your function **Greengrass_HelloWorld_Counter**, and set the remaining fields as follows:
 - For **Runtime**, choose **Python 3.7**.
 - For **Permissions**, keep the default setting. This creates an execution role that grants basic Lambda permissions. This role isn't used by AWS IoT Greengrass. Or, you can reuse the role that you created in Module 3-1.

Choose **Create function**.

Basic information

Function name
Enter a name that describes the purpose of your function.

Use only letters, numbers, hyphens, or underscores with no spaces.

Runtime [Info](#)
Choose the language to use to write your function.

Permissions [Info](#)
By default, Lambda will create an execution role with permissions to upload logs to Amazon CloudWatch Logs. You can customize this default role later when adding triggers.


► **Change default execution role**

► **Advanced settings**

Cancel **Create function**

6. Upload your Lambda function deployment package.

- a. On the **Code** tab, under **Code source**, choose **Upload from**. From the dropdown, choose **.zip file**.




- b. Choose **Upload**, and then choose your `hello_world_counter_python_lambda.zip` deployment package. Then, choose **Save**.
- c. On the **Code** tab for the function, under **Runtime settings**, choose **Edit**, and then enter the following values.
- For **Runtime**, choose **Python 3.7**.

- For **Handler**, enter `greengrassHelloWorldCounter.function_handler`
- d. Choose **Save**.

Note

The **Test** button on the AWS Lambda console doesn't work with this function. The AWS IoT Greengrass Core SDK doesn't contain modules that are required to run your Greengrass Lambda functions independently in the AWS Lambda console. These modules (for example, `greengrass_common`) are supplied to the functions after they are deployed to your Greengrass core.

7. Publish the first version of the function.
- From the **Actions** menu at the top of the page, choose **Publish new version**. For **Version description**, enter **First version**.
 - Choose **Publish**.
8. Create an alias for the function version.
- From the **Actions** menu at the top of the page, choose **Create alias**.



- For **Name**, enter `GG_HW_Counter`.
- For **Version**, choose **1**.
- Choose **Save**.

Create alias

Alias configuration

An alias is a pointer to one or two versions. Choose each version that you want the alias to point to.

Name

Description - *optional*

Version

1

▼

► Weighted alias

Cancel

Save

Aliases create a single entity for your Lambda function that Greengrass devices can subscribe to. This way, you don't have to update subscriptions with new Lambda function version numbers every time the function is modified.

Configure long-lived Lambda functions for AWS IoT Greengrass

You are now ready to configure your Lambda function for AWS IoT Greengrass.

1. In the AWS IoT console navigation pane, under **Manage**, expand **Greengrass devices**, and then choose **Groups (V1)**.
2. Under **Greengrass groups**, choose the group that you created in [Module 2](#).
3. On the group configuration page, choose the **Lambda functions** tab, and then under **My Lambda functions**, choose **Add**.
4. For **Lambda function**, choose **Greengrass_HelloWorld_Counter**.
5. For **Lambda function version**, choose the alias to the version that you published.
6. For **Timeout (seconds)**, enter **25**. This Lambda function sleeps for 20 seconds before each invocation.

7. For **Pinned**, choose **True**.
8. Keep the default values for all other fields, and choose **Add Lambda function**.

Test long-lived Lambda functions

A [long-lived](#) Lambda function starts automatically when the AWS IoT Greengrass core starts and runs in a single container (or sandbox). Any variables and preprocessing logic defined outside of the function handler are retained for every invocation of the function handler. Multiple invocations of the function handler are queued until earlier invocations have been executed.

The `greengrassHelloWorldCounter.py` code used in this module defines a `my_counter` variable outside of the function handler.

Note

You can view the code in the AWS Lambda console or in the [AWS IoT Greengrass Core SDK for Python](#) on GitHub.

In this step, you create subscriptions that allow the Lambda function and AWS IoT to exchange MQTT messages. Then you deploy the group and test the function.

1. On the group configuration page, choose **Subscriptions**, and then choose **Add**.
2. Under **Source type**, choose **Lambda function**, and then choose **Greengrass_HelloWorld_Counter**.
3. Under **Target type**, choose **Service**, choose **IoT Cloud**.
4. For **Topic filter**, enter **hello/world/counter**.
5. Choose **Create subscription**.

This single subscription goes in one direction only: from the `Greengrass_HelloWorld_Counter` Lambda function to AWS IoT. To invoke (or trigger) this Lambda function from the cloud, you must create a subscription in the opposite direction.

6. Follow steps 1 - 5 to add another subscription that uses the following values. This subscription allows the Lambda function to receive messages from AWS IoT. You use this subscription when you send a message from the AWS IoT console that invokes the function.
 - For the source, choose **Service**, and then choose **IoT Cloud**.

- For the target, choose **Lambda function**, and then choose **Greengrass_HelloWorld_Counter**.
- For the topic filter, enter **hello/world/counter/trigger**.

The `/trigger` extension is used in this topic filter because you created two subscriptions and don't want them to interfere with each other.

7. Make sure that the Greengrass daemon is running, as described in [Deploy cloud configurations to a core device](#).
8. On the group configuration page, choose **Deploy**.
9. After your deployment is complete, return to the AWS IoT console home page and choose **Test**.
10. Configure the following fields:
 - For **Subscription topic**, enter **hello/world/counter**.
 - For **Quality of Service**, choose **0**.
 - For **MQTT payload display**, choose **Display payloads as strings**.
11. Choose **Subscribe**.

Unlike [Part 1](#) of this module, you shouldn't see any messages after you subscribe to `hello/world/counter`. This is because the `greengrassHelloWorldCounter.py` code that publishes to the `hello/world/counter` topic is inside the function handler, which runs only when the function is invoked.

In this module, you configured the `Greengrass_HelloWorld_Counter` Lambda function to be invoked when it receives an MQTT message on the `hello/world/counter/trigger` topic.

The **Greengrass_HelloWorld_Counter to IoT Cloud** subscription allows the function to send messages to AWS IoT on the `hello/world/counter` topic. The **IoT Cloud to Greengrass_HelloWorld_Counter** subscription allows AWS IoT to send messages to the function on the `hello/world/counter/trigger` topic.

12. To test the long-lived lifecycle, invoke the Lambda function by publishing a message to the `hello/world/counter/trigger` topic. You can use the default message.

Subscribe to a topic**Publish to a topic****Topic name**

The topic name identifies the message. The message payload will be published to this topic with a Quality of Service (QoS) of 0.

Message payload**Additional configuration****Publish****Note**

The `Greengrass_HelloWorld_Counter` function ignores the content of received messages. It just runs the code in `function_handler`, which sends a message to the `hello/world/counter` topic. You can review this code from the [AWS IoT Greengrass Core SDK for Python](#) on GitHub.

Every time a message is published to the `hello/world/counter/trigger` topic, the `my_counter` variable is incremented. This invocation count is shown in the messages sent from the Lambda function. Because the function handler includes a 20-second sleep cycle (`time.sleep(20)`), repeatedly triggering the handler queues up responses from the AWS IoT Greengrass core.

Subscriptions

hello/world/counter

hello/world/counter May 03, 2021, 10:05:00 (UTC-0400)

```
{
  "message": "Hello world! Sent from Greengrass Core running on platform: Linux-4.9.30-v7+-armv7l-with-Debian-8.0"
}
```

Invocation Count: 3

hello/world/counter May 03, 2021, 10:04:40 (UTC-0400)

```
{
  "message": "Hello world! Sent from Greengrass Core running on platform: Linux-4.9.30-v7+-armv7l-with-Debian-8.0"
}
```

Invocation Count: 2

hello/world/counter May 03, 2021, 10:04:20 (UTC-0400)

```
{
  "message": "Hello world! Sent from Greengrass Core running on platform: Linux-4.9.30-v7+-armv7l-with-Debian-8.0"
}
```

Invocation Count: 1

Test on-demand Lambda functions

An [on-demand](#) Lambda function is similar in functionality to a cloud-based AWS Lambda function. Multiple invocations of an on-demand Lambda function can run in parallel. An invocation of the Lambda function creates a separate container to process invocations or reuses an existing container, if resources permit. Any variables or preprocessing that are defined outside of the function handler are not retained when containers are created.

1. On the group configuration page, choose the **Lambda functions** tab.
2. Under **My Lambda functions**, choose the **Greengrass_HelloWorld_Counter** Lambda function.
3. On the **Greengrass_HelloWorld_Counter** details page, choose **Edit**.
4. For **Pinned**, choose **False**, and then choose **Save**.

5. On the group configuration page, choose **Deploy**.
6. After your deployment is complete, return to the AWS IoT console home page and choose **Test**.
7. Configure the following fields:
 - For **Subscription topic**, enter **hello/world/counter**.
 - For **Quality of Service**, choose **0**.
 - For **MQTT payload display**, choose **Display payloads as strings**.

Subscribe to a topic | **Publish to a topic**

Topic filter [Info](#)
The topic filter describes the topic(s) to which you want to subscribe. The topic filter can include MQTT wildcard characters.

hello/world/counter

▼ **Additional configuration**

Number of messages to keep
The MQTT test client keeps this many of the most recent messages published to a topic that matches this topic filter.

100

Quality of service
When subscribing to a topic, quality of service 0 will be chosen by default.

☒ Quality of Service 0 - Message will be delivered at most once

☐ Quality of Service 1 - Message will be delivered at least once

MQTT payload display

☐ Auto-format JSON payloads (improves readability)

☒ Display payloads as strings (more accurate)

☐ Display raw payloads (displays binary data as hexadecimal values)

Subscribe


8. Choose **Subscribe**.

Note

You should not see any messages after you subscribe.

9. To test the on-demand lifecycle, invoke the function by publishing a message to the hello/world/counter/trigger topic. You can use the default message.

- a. Choose **Publish** three times quickly, within five seconds of each press of the button.



The screenshot shows the AWS IoT Greengrass console interface for publishing a message. At the top, there are two tabs: 'Subscribe to a topic' and 'Publish to a topic'. The 'Publish to a topic' tab is selected. Below the tabs, there is a 'Topic name' field with the text 'hello/world/counter/trigger' entered. A red circle highlights this field. Below the topic name field is a 'Message payload' field, which is currently empty. At the bottom of the form, there is a section titled 'Additional configuration' with a 'Publish' button. The 'Publish' button is circled in red, and a red 'x3' is next to it, indicating that the button should be clicked three times.

Each publish invokes the function handler and creates a container for each invocation. The invocation count is not incremented for the three times you triggered the function because each on-demand Lambda function has its own container/sandbox.

The screenshot displays the AWS IoT Greengrass console interface. On the left, a sidebar titled "Subscriptions" lists the topic "hello/world/counter" with a heart icon and a close button. The main panel shows the details for this topic, including buttons for "Pause", "Clear", "Export", and "Edit". Below these, three subscription events are listed, each with a timestamp and a message. The message for each event is a JSON object containing a "message" field and an "Invocation Count" field, which is circled in red and shows the value "1".

Subscription	Timestamp	Message
hello/world/counter	May 03, 2021, 10:14:13 (UTC-0400)	<pre>{ "message": "Hello world! Sent from Greengrass Core running on platform: Linux-4.9.30-v7+-armv7l-with-Debian-8.0.", "Invocation Count: 1" }</pre>
hello/world/counter	May 03, 2021, 10:14:12 (UTC-0400)	<pre>{ "message": "Hello world! Sent from Greengrass Core running on platform: Linux-4.9.30-v7+-armv7l-with-Debian-8.0.", "Invocation Count: 1" }</pre>
hello/world/counter	May 03, 2021, 10:14:11 (UTC-0400)	<pre>{ "message": "Hello world! Sent from Greengrass Core running on platform: Linux-4.9.30-v7+-armv7l-with-Debian-8.0.", "Invocation Count: 1" }</pre>

- b. After approximately 30 seconds, choose **Publish to topic**. The invocation count should be incremented to 2. This shows that a container created from an earlier invocation is being reused, and that preprocessing variables outside of the function handler were stored.


The screenshot displays the AWS IoT Greengrass console interface. On the left, a sidebar titled 'Subscriptions' contains a single entry 'hello/world/counter' with a heart icon and a close button. The main panel is titled 'hello/world/counter' and includes 'Pause', 'Clear', 'Export', and 'Edit' buttons. It lists two subscription events:

- Event 1:** Timestamped 'May 03, 2021, 10:14:40 (UTC-0400)', the message body is `{ "message": "Hello world! Sent from Greengrass Core running on platform: Linux-4.9.30-v7+-armv7l-with-Debian-8.0. Invocation Count: 2" }`. The text 'Invocation Count: 2' is circled in red.
- Event 2:** Timestamped 'May 03, 2021, 10:14:11 (UTC-0400)', the message body is `{ "message": "Hello world! Sent from Greengrass Core running on platform: Linux-4.9.30-v7+-armv7l-with-Debian-8.0. Invocation Count: 1" }`. The text 'Invocation Count: 1' is circled in red.

You should now understand the two types of Lambda functions that can run on the AWS IoT Greengrass core. The next module, [Module 4](#), shows you how local IoT devices can interact in an AWS IoT Greengrass group.

Module 4: Interacting with client devices in an AWS IoT Greengrass group

This module shows you how local IoT devices, called *client devices* or *devices*, can connect to and communicate with an AWS IoT Greengrass core device. Client devices that connect to an AWS IoT Greengrass core are part of an AWS IoT Greengrass group and can participate in the AWS IoT Greengrass programming paradigm. In this module, one client device sends a Hello World message to another client device in the Greengrass group.



Before you begin, run the [Greengrass device setup](#) script or complete [Module 1](#) and [Module 2](#). This module creates two simulated client devices. You do not need other components or devices.

This module should take less than 30 minutes to complete.

Topics

- [Create client devices in an AWS IoT Greengrass group](#)
- [Configure subscriptions](#)
- [Install the AWS IoT Device SDK for Python](#)
- [Test communications](#)

Create client devices in an AWS IoT Greengrass group

In this step, you add two client devices to your Greengrass group. This process includes registering the devices as AWS IoT things and configuring certificates and keys to allow them to connect to AWS IoT Greengrass.

1. In the AWS IoT console navigation pane, under **Manage**, expand **Greengrass devices**, and then choose **Groups (V1)**.
2. Choose the target group.
3. On the group configuration page, choose **Client devices**, and then choose **Associate**.
4. In the **Associate a client device with this group** modal, choose **Create new AWS IoT thing**.

The **Create things** page opens in a new tab.

5. On the **Create things** page, choose **Create single thing**, and then choose **Next**.
6. On the **Specify thing properties** page, register this client device as **HelloWorld_Publisher**, and then choose **Next**.
7. On the **Configure device certificate** page, choose **Next**.
8. On the **Attach policies to certificate** page, do one of the following:
 - Select an existing policy that grants permissions that client devices require, and then choose **Create thing**.

A modal opens where you can download the certificates and keys that the device uses to connect to the AWS Cloud and the core.

- Create and attach a new policy that grants client device permissions. Do the following:
 - a. Choose **Create policy**.

The **Create policy** page opens in a new tab.

- b. On the **Create policy** page, do the following:
 - i. For **Policy name**, enter a name that describes the policy, such as **GreengrassV1ClientDevicePolicy**.
 - ii. On the **Policy statements** tab, under **Policy document**, choose **JSON**.
 - iii. Enter the following policy document. This policy allows the client device to discover Greengrass cores and communicate on all MQTT topics. For information

about how to restrict this policy's access, see [Device authentication and authorization for AWS IoT Greengrass](#).

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "iot:Publish",
        "iot:Subscribe",
        "iot:Connect",
        "iot:Receive"
      ],
      "Resource": [
        "*"
      ]
    },
    {
      "Effect": "Allow",
      "Action": [
        "greengrass:*"
      ],
      "Resource": [
        "*"
      ]
    }
  ]
}
```

- iv. Choose **Create** to create the policy.
- c. Return to the browser tab with the **Attach policies to certificate** page open. Do the following:
 - i. In the **Policies** list, select the policy that you created, such as **GreengrassV1ClientDevicePolicy**.

If you don't see the policy, choose the refresh button.

- ii. Choose **Create thing**.

A modal opens where you can download the certificates and keys that the device uses to connect to the AWS Cloud and the core.

9. In the **Download certificates and keys** modal, download the device's certificates.

 Important

Before you choose **Done**, download the security resources.

Do the following:

- a. For **Device certificate**, choose **Download** to download the device certificate.
- b. For **Public key file**, choose **Download** to download the public key for the certificate.
- c. For **Private key file**, choose **Download** to download the private key file for the certificate.
- d. Review [Server Authentication](#) in the *AWS IoT Developer Guide* and choose the appropriate root CA certificate. We recommend that you use Amazon Trust Services (ATS) endpoints and ATS root CA certificates. Under **Root CA certificates**, choose **Download** for a root CA certificate.
- e. Choose **Done**.

Make a note of the certificate ID that's common in the file names for the device certificate and keys. You need it later.

10. Return to the browser tab with the **Associate a client device with this group** modal open. Do the following:
 - a. For **AWS IoT thing name**, choose the **HelloWorld_Publisher** thing that you created.

If you don't see the thing, choose the refresh button.
 - b. Choose **Associate**.
11. Repeat steps 3 - 10 to add a second client device to the group.

Name this client device **HelloWorld_Subscriber**. Download the certificates and keys for this client device to your computer. Again, make a note of the certificate's ID that's common in the file names for the HelloWorld_Subscriber device.

You should now have two client devices in your Greengrass group:

- HelloWorld_Publisher
- HelloWorld_Subscriber

12. Create a folder on your computer for these client devices' security credentials. Copy the certificates and keys into this folder.

Configure subscriptions

In this step, you enable the HelloWorld_Publisher client device to send MQTT messages to the HelloWorld_Subscriber client device.

1. On the group configuration page, choose the **Subscriptions** tab, and then choose **Add**.
2. On the **Create a subscription** page, do the following to configure the subscription:
 - a. For **Source type**, choose **Client device**, and then choose **HelloWorld_Publisher**.
 - b. Under **Target type**, choose **Client device**, and then choose **HelloWorld_Subscriber**.
 - c. For **Topic filter**, enter **hello/world/pubsub**.

Note

You can delete subscriptions from the previous modules. On the group's **Subscriptions** page, select the subscriptions to delete, and then choose **Delete**.

- d. Choose **Create subscription**.
3. Make sure that automatic detection is enabled so the Greengrass core can publish a list of its IP addresses. Client devices use this information to discover the core. Do the following:
 - a. On the group configuration page, choose the **Lambda functions** tab.
 - b. Under **System Lambda functions**, choose **IP detector**, and then choose **Edit**.
 - c. In the **Edit IP detector settings**, choose **Automatically detect and override MQTT broker endpoints**, and then choose **Save**.
 4. Make sure that the Greengrass daemon is running, as described in [Deploy cloud configurations to a core device](#).
 5. On the group configuration page, choose **Deploy**.

The deployment status is displayed below the group name on the page header. To see deployment details, choose the **Deployments** tab.

Install the AWS IoT Device SDK for Python

Client devices can use the AWS IoT Device SDK for Python to communicate with AWS IoT and AWS IoT Greengrass core devices (using the Python programming language). For more information, including requirements, see the AWS IoT Device SDK for Python [Readme](#) on GitHub.

In this step, you install the SDK and get the `basicDiscovery.py` sample function used by the simulated client devices on your computer.

1. To install the SDK on your computer, with all required components, choose your operating system:

Windows

1. Open an [elevated command prompt](#) and run the following command:

```
python --version
```

If no version information is returned or if the version number is less than 2.7 for Python 2 or less than 3.3 for Python 3, follow the instructions in [Downloading Python](#) to install Python 2.7+ or Python 3.3+. For more information, see [Using Python on Windows](#).

2. Download the [AWS IoT Device SDK for Python](#) as a zip file and extract it to an appropriate location on your computer.

Make a note of the file path to the extracted `aws-iot-device-sdk-python-master` folder that contains the `setup.py` file. In the next step, this file path is indicated by *path-to-SDK-folder*.

3. From the elevated command prompt, run the following:

```
cd path-to-SDK-folder  
python setup.py install
```

macOS

1. Open a Terminal window and run the following command:

```
python --version
```

If no version information is returned or if the version number is less than 2.7 for Python 2 or less than 3.3 for Python 3, follow the instructions in [Downloading Python](#) to install Python 2.7+ or Python 3.3+. For more information, see [Using Python on a Macintosh](#).

2. In the Terminal window, run the following commands to determine the OpenSSL version:

```
python
>>>import ssl
>>>print ssl.OPENSSL_VERSION
```

Make a note of the OpenSSL version value.

 Note

If you're running Python 3, use **print(ssl.OPENSSL_VERSION)**.

To close the Python shell, run the following command:

```
>>>exit()
```

If the OpenSSL version is 1.0.1 or later, skip to [step c](#). Otherwise, follow these steps:

- From the Terminal window, run the following command to determine if the computer is using Simple Python Version Management:

```
which pyenv
```

If a file path is returned, then choose the **Using pyenv** tab. If nothing is returned, choose the **Not using pyenv** tab.

Using pyenv

1. See [Python Releases for Mac OS X](#) (or similar) to determine the latest stable Python version. In the following example, this value is indicated by *latest-Python-version*.

2. From the Terminal window, run the following commands:

```
pyenv install latest-Python-version  
pyenv global latest-Python-version
```

For example, if the latest version for Python 2 is 2.7.14, then these commands are:

```
pyenv install 2.7.14  
pyenv global 2.7.14
```

3. Close and then reopen the Terminal window and then run the following commands:

```
python  
>>>import ssl  
>>>print ssl.OPENSSL_VERSION
```

The OpenSSL version should be at least 1.0.1. If the version is less than 1.0.1, then the update failed. Check the Python version value used in the **pyenv install** and **pyenv global** commands and try again.

4. Run the following command to exit the Python shell:

```
exit()
```

Not using pyenv

1. From a Terminal window, run the following command to determine if [brew](#) is installed:

```
which brew
```

If a file path is not returned, install brew as follows:

```
/usr/bin/ruby -e "$(curl -fsSL https://raw.githubusercontent.com/  
Homebrew/install/master/install)"
```

 Note

Follow the installation prompts. The download for the Xcode command line tools can take some time.

2. Run the following commands:

```
brew update
brew install openssl
brew install python@2
```

The AWS IoT Device SDK for Python requires OpenSSL version 1.0.1 (or later) compiled with the Python executable. The **brew install python** command installs a python2 executable that meets this requirement. The python2 executable is installed in the `/usr/local/bin` directory, which should be part of the PATH environment variable. To confirm, run the following command:

```
python2 --version
```

If python2 version information is provided, skip to the next step. Otherwise, permanently add the `/usr/local/bin` path to your PATH environment variable by appending the following line to your shell profile:

```
export PATH="/usr/local/bin:$PATH"
```

For example, if you're using `.bash_profile` or do not yet have a shell profile, run the following command from a Terminal window:

```
echo 'export PATH="/usr/local/bin:$PATH"' >> ~/.bash_profile
```

Next, [source](#) your shell profile and confirm that `python2 --version` provides version information. For example, if you're using `.bash_profile`, run the following commands:

```
source ~/.bash_profile
python2 --version
```

python2 version information should be returned.

3. Append the following line to your shell profile:

```
alias python="python2"
```

For example, if you're using `.bash_profile` or do not yet have a shell profile, run the following command:

```
echo 'alias python="python2"' >> ~/.bash_profile
```

4. Next, [source](#) your shell profile. For example, if you're using `.bash_profile`, run the following command:

```
source ~/.bash_profile
```

Invoking the **python** command runs the Python executable that contains the required OpenSSL version (python2) .

5. Run the following commands:

```
python
import ssl
print ssl.OPENSSL_VERSION
```

The OpenSSL version should be 1.0.1 or later.

6. To exit the Python shell, run the following command:

```
exit()
```

3. Run the following commands to install the AWS IoT Device SDK for Python:

```
cd ~
git clone https://github.com/aws/aws-iot-device-sdk-python.git
cd aws-iot-device-sdk-python
sudo python setup.py install
```

UNIX-like system

1. From a terminal window, run the following command:

```
python --version
```

If no version information is returned or if the version number is less than 2.7 for Python 2 or less than 3.3 for Python 3, follow the instructions in [Downloading Python](#) to install Python 2.7+ or Python 3.3+. For more information, see [Using Python on Unix platforms](#).

2. In the terminal, run the following commands to determine the OpenSSL version:

```
python
>>>import ssl
>>>print ssl.OPENSSL_VERSION
```

Make a note of the OpenSSL version value.

Note

If you're running Python 3, use **print(ssl.OPENSSL_VERSION)**.

To close the Python shell, run the following command:

```
exit()
```

If the OpenSSL version is 1.0.1 or later, skip to the next step. Otherwise, run the command(s) to update OpenSSL for your distribution (for example, `sudo yum update openssl`, `sudo apt-get update`, and so on).

Confirm that the OpenSSL version is 1.0.1 or later by running the following commands:

```
python
>>>import ssl
>>>print ssl.OPENSSL_VERSION
>>>exit()
```

3. Run the following commands to install the AWS IoT Device SDK for Python:

```
cd ~
git clone https://github.com/aws/aws-iot-device-sdk-python.git
cd aws-iot-device-sdk-python
sudo python setup.py install
```

2. After the AWS IoT Device SDK for Python is installed, navigate to the `samples` folder and open the `greengrass` folder.

For this tutorial, you copy the `basicDiscovery.py` sample function, which uses the certificates and keys that you downloaded in [the section called “Create client devices in an AWS IoT Greengrass group”](#).

3. Copy `basicDiscovery.py` to the folder that contains the `HelloWorld_Publisher` and `HelloWorld_Subscriber` device certificates and keys.

Test communications

1. Make sure that your computer and the AWS IoT Greengrass core device are connected to the internet using the same network.
 - a. On the AWS IoT Greengrass core device, run the following command to find its IP address.

```
hostname -I
```

- b. On your computer, run the following command using the IP address of the core. You can use **Ctrl + C** to stop the **ping** command.

```
ping IP-address
```

Output similar to the following indicates successful communication between the computer and the AWS IoT Greengrass core device (0% packet loss):

```
$ping 176.32.103.205
PING 176.32.103.205 (176.32.103.205) 56(84) bytes of data.
64 bytes from 176.32.103.205: icmp_seq=1 ttl=230 time=77.2 ms
64 bytes from 176.32.103.205: icmp_seq=2 ttl=230 time=77.1 ms
64 bytes from 176.32.103.205: icmp_seq=3 ttl=230 time=77.1 ms
64 bytes from 176.32.103.205: icmp_seq=4 ttl=230 time=77.1 ms
64 bytes from 176.32.103.205: icmp_seq=5 ttl=230 time=77.1 ms
64 bytes from 176.32.103.205: icmp_seq=6 ttl=230 time=77.1 ms
^C
--- 176.32.103.205 ping statistics ---
6 packets transmitted, 6 received, 0% packet loss, time 5549ms
rtt min/avg/max/mdev = 77.107/77.172/77.256/0.361 ms
```

Note

If you're unable to ping an EC2 instance that's running AWS IoT Greengrass, make sure that the inbound security group rules for the instance allow ICMP traffic for [Echo request](#) messages. For more information, see [Adding rules to a security group](#) in the *Amazon EC2 User Guide*.

On Windows host computers, in the Windows Firewall with Advanced Security app, you might also need to enable an inbound rule that allows inbound echo requests (for example, **File and Printer Sharing (Echo Request - ICMPv4-In)**), or create one.

2. Get your AWS IoT endpoint.
 - a. From the [AWS IoT console](#) navigation pane, choose **Settings**.
 - b. Under **Device data endpoint**, make a note of the value of **Endpoint**. You use this value to replace the `AWS_IOT_ENDPOINT` placeholder in the commands in the following steps.

Note

Make sure that your [endpoints correspond to your certificate type](#).

3. On your computer (not the AWS IoT Greengrass core device), open two [command-line](#) (terminal or command prompt) windows. One window represents the HelloWorld_Publisher client device and the other represents the HelloWorld_Subscriber client device.

Upon execution, `basicDiscovery.py` attempts to collect information on the location of the AWS IoT Greengrass core at its endpoints. This information is stored after the client device

has discovered and successfully connected to the core. This allows future messaging and operations to be executed locally (without the need for an internet connection).

Note

Client IDs used for MQTT connections must match the thing name of the client device. The `basicDiscovery.py` script sets the client ID for MQTT connections to the thing name that you specify when you run the script.

Run the following command from the folder that contains the `basicDiscovery.py` file for detailed script usage information:

```
python basicDiscovery.py --help
```

4. From the `HelloWorld_Publisher` client device window, run the following commands.
 - Replace *path-to-certs-folder* with the path to the folder that contains the certificates, keys, and `basicDiscovery.py`.
 - Replace *AWS_IOT_ENDPOINT* with your endpoint.
 - Replace the two *publisherCertId* instances with the certificate ID in the file name for your `HelloWorld_Publisher` client device.

```
cd path-to-certs-folder
python basicDiscovery.py --endpoint AWS_IOT_ENDPOINT --rootCA AmazonRootCA1.pem
--cert publisherCertId-certificate.pem.crt --key publisherCertId-private.pem.key
--thingName HelloWorld_Publisher --topic 'hello/world/pubsub' --mode publish --
message 'Hello, World! Sent from HelloWorld_Publisher'
```

You should see output similar to the following, which includes entries such as Published topic 'hello/world/pubsub': {"message": "Hello, World! Sent from HelloWorld_Publisher", "sequence": 1}.

Note

If the script returns an error: `unrecognized arguments message`, change the single quotation marks to double quotation marks for the `--topic` and `--message` parameters and run the command again.

To troubleshoot a connection issue, you can try using [manual IP detection](#).

```
Published topic hello/world/pubsub: {"message": "Hello, World! Sent from HelloWorld_Publisher", "sequence": 0}
2017-11-13 21:12:26,296 - AWSIoTPythonSDK.core.protocol.internal.workers - DEBUG - Produced [puback] event
2017-11-13 21:12:26,297 - AWSIoTPythonSDK.core.protocol.internal.workers - DEBUG - Dispatching [puback] event
2017-11-13 21:12:27,301 - AWSIoTPythonSDK.core.protocol.mqtt_core - INFO - Performing sync publish...
Published topic hello/world/pubsub: {"message": "Hello, World! Sent from HelloWorld_Publisher", "sequence": 1}
2017-11-13 21:12:27,302 - AWSIoTPythonSDK.core.protocol.internal.workers - DEBUG - Produced [puback] event
2017-11-13 21:12:27,303 - AWSIoTPythonSDK.core.protocol.internal.workers - DEBUG - Dispatching [puback] event
2017-11-13 21:12:28,305 - AWSIoTPythonSDK.core.protocol.mqtt_core - INFO - Performing sync publish...
Published topic hello/world/pubsub: {"message": "Hello, World! Sent from HelloWorld_Publisher", "sequence": 2}
2017-11-13 21:12:28,306 - AWSIoTPythonSDK.core.protocol.internal.workers - DEBUG - Produced [puback] event
2017-11-13 21:12:28,307 - AWSIoTPythonSDK.core.protocol.internal.workers - DEBUG - Dispatching [puback] event
2017-11-13 21:12:29,310 - AWSIoTPythonSDK.core.protocol.mqtt_core - INFO - Performing sync publish...
Published topic hello/world/pubsub: {"message": "Hello, World! Sent from HelloWorld_Publisher", "sequence": 3}
```

5. From the HelloWorld_Subscriber client device window, run the following commands.

- Replace *path-to-certs-folder* with the path to the folder that contains the certificates, keys, and basicDiscovery.py.
- Replace *AWS_IOT_ENDPOINT* with your endpoint.
- Replace the two *subscriberCertId* instances with the certificate ID in the file name for your HelloWorld_Subscriber client device.

```
cd path-to-certs-folder
python basicDiscovery.py --endpoint AWS_IOT_ENDPOINT --rootCA AmazonRootCA1.pem --
cert subscriberCertId-certificate.pem.crt --key subscriberCertId-private.pem.key --
thingName HelloWorld_Subscriber --topic 'hello/world/pubsub' --mode subscribe
```

You should see the following output, which includes entries such as Received message on topic hello/world/pubsub: {"message": "Hello, World! Sent from HelloWorld_Publisher", "sequence": 1}.

```
Received message on topic hello/world/pubsub: {"message": "Hello, World! Sent from HelloWorld_Publisher", "sequence": 0}
2017-11-13 21:12:27,435 - AWSIoTPythonSDK.core.protocol.internal.workers - DEBUG - Produced [message] event
2017-11-13 21:12:27,435 - AWSIoTPythonSDK.core.protocol.internal.workers - DEBUG - Dispatching [message] event
2017-11-13 21:12:27,436 - AWSIoTPythonSDK.core.protocol.internal.clients - DEBUG - Invoking custom event callback...
Received message on topic hello/world/pubsub: {"message": "Hello, World! Sent from HelloWorld_Publisher", "sequence": 1}
2017-11-13 21:12:28,320 - AWSIoTPythonSDK.core.protocol.internal.workers - DEBUG - Produced [message] event
2017-11-13 21:12:28,324 - AWSIoTPythonSDK.core.protocol.internal.workers - DEBUG - Dispatching [message] event
2017-11-13 21:12:28,324 - AWSIoTPythonSDK.core.protocol.internal.clients - DEBUG - Invoking custom event callback...
Received message on topic hello/world/pubsub: {"message": "Hello, World! Sent from HelloWorld_Publisher", "sequence": 2}
2017-11-13 21:12:29,547 - AWSIoTPythonSDK.core.protocol.internal.workers - DEBUG - Produced [message] event
2017-11-13 21:12:29,552 - AWSIoTPythonSDK.core.protocol.internal.workers - DEBUG - Dispatching [message] event
2017-11-13 21:12:29,552 - AWSIoTPythonSDK.core.protocol.internal.clients - DEBUG - Invoking custom event callback...
```

Close the HelloWorld_Publisher window to stop messages from accruing in the HelloWorld_Subscriber window.


Testing on a corporate network might interfere with connecting to the core. As a workaround, you can manually enter the endpoint. This ensures that the `basicDiscovery.py` script connects to the correct IP address of the AWS IoT Greengrass core device.

To manually enter the endpoint

1. In the AWS IoT console navigation pane, under **Manage**, expand **Greengrass devices**, and then choose **Groups (V1)**.
2. Under **Greengrass groups**, choose your group.
3. Configure the core to manually manage MQTT broker endpoints. Do the following:
 - a. On the group configuration page, choose the **Lambda functions** tab.
 - b. Under **System Lambda functions**, choose **IP detector**, and then choose **Edit**.
 - c. In the **Edit IP detector settings**, choose **Manually manage MQTT broker endpoints**, and then choose **Save**.
4. Enter the MQTT broker endpoint for the core. Do the following:
 - a. Under **Overview**, choose the **Greengrass core**.
 - b. Under **MQTT broker endpoints**, choose **Manage endpoints**.
 - c. Choose **Add endpoint** and make sure that you have only one endpoint value. This value must be the IP address endpoint for port 8883 of your AWS IoT Greengrass core device (for example, `192.168.1.4`).
 - d. Choose **Update**.

Module 5: Interacting with device shadows

This advanced module shows you how client devices can interact with [AWS IoT device shadows](#) in an AWS IoT Greengrass group. A *shadow* is a JSON document that is used to store current or desired state information for a thing. In this module, you discover how one client device (GG_Switch) can modify the state of another client device (GG_TrafficLight) and how these states can be synced to the AWS IoT Greengrass cloud:



Before you begin, run the [Greengrass device setup](#) script, or make sure that you have completed [Module 1](#) and [Module 2](#). You should also understand how to connect client devices to an AWS IoT Greengrass core ([Module 4](#)). You do not need other components or devices.

This module should take about 30 minutes to complete.

Topics

- [Configure devices and subscriptions](#)
- [Download required files](#)
- [Test communications \(device syncs disabled\)](#)
- [Test communications \(device syncs enabled\)](#)

Configure devices and subscriptions

Shadows can be synced to AWS IoT when the AWS IoT Greengrass core is connected to the internet. In this module, you first use local shadows without syncing to the cloud. Then, you enable cloud syncing.

Each client device has its own shadow. For more information, see [Device shadow service for AWS IoT](#) in the *AWS IoT Developer Guide*.

1. On the group configuration page, choose the **Client devices** tab.
2. From the **Client devices** tab, add two new client devices in your AWS IoT Greengrass group. For detailed steps of this process, see [the section called "Create client devices in an AWS IoT Greengrass group"](#).
 - Name the client devices **GG_Switch** and **GG_TrafficLight**.
 - Generate and download the security resources for both client devices.
 - Make a note of the certificate ID in the file names of the security resources for the client devices. You use these values later.
3. Create a folder on your computer for these client devices' security credentials. Copy the certificates and keys into this folder.
4. Make sure that the client devices are set to use local shadows and not sync with the AWS Cloud. If not, select the client device, choose **Sync shadow**, and then choose **Disable shadow sync with cloud**.
5. Add the subscriptions in the following table to your group. For example, to create the first subscription:
 - a. On the group configuration page, choose the **Subscriptions** tab, and then choose **Add**.
 - b. For **Source type**, choose **Client device**, and then choose **GG_Switch**.
 - c. For **Target type**, choose **Service**, and then choose **Local Shadow Service**.
 - d. For **Topic filter**, enter `$aws/things/GG_TrafficLight/shadow/update`
 - e. Choose **Create subscription**.

The topics must be entered exactly as shown in the table. Although it's possible to use wildcards to consolidate some of the subscriptions, we don't recommend this practice. For more information, see [Shadow MQTT topics](#) in the *AWS IoT Developer Guide*.

Source	Target	Topic	Notes
GG_Switch	Local Shadow Service	\$aws/things/GG_TrafficLight/shadow/update	The GG_Switch sends an update request to update topic.
Local Shadow Service	GG_Switch	\$aws/things/GG_TrafficLight/shadow/update/accepted	The GG_Switch needs to know whether the update request was accepted.
Local Shadow Service	GG_Switch	\$aws/things/GG_TrafficLight/shadow/update/rejected	The GG_Switch needs to know whether the update request was rejected.
GG_TrafficLight	Local Shadow Service	\$aws/things/GG_TrafficLight/shadow/update	The GG_TrafficLight sends an update of its state to the update topic.
Local Shadow Service	GG_TrafficLight	\$aws/things/GG_TrafficLight/shadow/update/delta	The Local Shadow Service sends a received update to GG_TrafficLight through the delta topic.
Local Shadow Service	GG_TrafficLight	\$aws/things/GG_TrafficLight/shadow/update/accepted	The GG_TrafficLight needs to know whether its state update was accepted.

Source	Target	Topic	Notes
Local Shadow Service	GG_TrafficLight	\$aws/things/GG_TrafficLight/shadow/update/rejected	The GG_TrafficLight needs to know whether its state update was rejected.

The new subscriptions are displayed on the **Subscriptions** tab.

 Note

For information about the \$ character, see [Reserved topics](#).

6. Make sure that automatic detection is enabled so the Greengrass core can publish a list of its IP addresses. Client devices use this information to discover the core. Do the following:
 - a. On the group configuration page, choose the **Lambda functions** tab.
 - b. Under **System Lambda functions**, choose **IP detector**, and then choose **Edit**.
 - c. In the **Edit IP detector settings**, choose **Automatically detect and override MQTT broker endpoints**, and then choose **Save**.
7. Make sure that the Greengrass daemon is running, as described in [Deploy cloud configurations to a core device](#).
8. On the group configuration page, choose **Deploy**.

Download required files

1. If you haven't already done so, install the AWS IoT Device SDK for Python. For instructions, see step 1 in [the section called "Install the AWS IoT Device SDK for Python"](#).

This SDK is used by client devices to communicate with AWS IoT and with AWS IoT Greengrass core devices.

2. From the [TrafficLight](#) examples folder on GitHub, download the `lightController.py` and `trafficLight.py` files to your computer. Save them in the folder that contains the GG_Switch and GG_TrafficLight client device certificates and keys.

The `lightController.py` script corresponds to the `GG_Switch` client device, and the `trafficLight.py` script corresponds to the `GG_TrafficLight` client device.

-  `7aa87aa1cf.cert.pem`
-  `7aa87aa1cf.private.key`
-  `7aa87aa1cf.public.key`
-  `a27b261ea9.cert.pem`
-  `a27b261ea9.private.key`
-  `a27b261ea9.public.key`
-  `lightController.py`
-  `root-ca-cert.pem`
-  `trafficLight.py`

Note

The example Python files are stored in the AWS IoT Greengrass Core SDK for Python repository for convenience, but they don't use the AWS IoT Greengrass Core SDK.

Test communications (device syncs disabled)

1. Make sure that your computer and the AWS IoT Greengrass core device are connected to the internet using the same network.
 - a. On the AWS IoT Greengrass core device, run the following command to find its IP address.

```
hostname -I
```

- b. On your computer, run the following command using the IP address of the core. You can use **Ctrl + C** to stop the **ping** command.

```
ping IP-address
```

Output similar to the following indicates successful communication between the computer and the AWS IoT Greengrass core device (0% packet loss):


```
$ping 176.32.103.205
PING 176.32.103.205 (176.32.103.205) 56(84) bytes of data.
64 bytes from 176.32.103.205: icmp_seq=1 ttl=230 time=77.2 ms
64 bytes from 176.32.103.205: icmp_seq=2 ttl=230 time=77.1 ms
64 bytes from 176.32.103.205: icmp_seq=3 ttl=230 time=77.1 ms
64 bytes from 176.32.103.205: icmp_seq=4 ttl=230 time=77.1 ms
64 bytes from 176.32.103.205: icmp_seq=5 ttl=230 time=77.1 ms
64 bytes from 176.32.103.205: icmp_seq=6 ttl=230 time=77.1 ms
^C
--- 176.32.103.205 ping statistics ---
6 packets transmitted, 6 received, 0% packet loss, time 5549ms
rtt min/avg/max/mdev = 77.107/77.172/77.256/0.361 ms
```

Note

If you're unable to ping an EC2 instance that's running AWS IoT Greengrass, make sure that the inbound security group rules for the instance allow ICMP traffic for [Echo request](#) messages. For more information, see [Adding rules to a security group](#) in the *Amazon EC2 User Guide*.

On Windows host computers, in the Windows Firewall with Advanced Security app, you might also need to enable an inbound rule that allows inbound echo requests (for example, **File and Printer Sharing (Echo Request - ICMPv4-In)**), or create one.

2. Get your AWS IoT endpoint.
 - a. From the [AWS IoT console](#) navigation pane, choose **Settings**.
 - b. Under **Device data endpoint**, make a note of the value of **Endpoint**. You use this value to replace the `AWS_IOT_ENDPOINT` placeholder in the commands in the following steps.

Note

Make sure that your [endpoints correspond to your certificate type](#).

3. On your computer (not the AWS IoT Greengrass core device), open two [command-line](#) (terminal or command prompt) windows. One window represents the GG_Switch client device and the other represents the GG_TrafficLight client device.
 - a. From the GG_Switch client device window, run the following commands.

- Replace *path-to-certs-folder* with the path to the folder that contains the certificates, keys, and Python files.
- Replace *AWS_IOT_ENDPOINT* with your endpoint.
- Replace the two *switchCertId* instances with the certificate ID in the file name for your GG_Switch client device.

```
cd path-to-certs-folder
python lightController.py --endpoint AWS_IOT_ENDPOINT --rootCA
  AmazonRootCA1.pem --cert switchCertId-certificate.pem.crt --key switchCertId-
  private.pem.key --thingName GG_TrafficLight --clientId GG_Switch
```

b. From the GG_TrafficLight client device window, run the following commands.

- Replace *path-to-certs-folder* with the path to the folder that contains the certificates, keys, and Python files.
- Replace *AWS_IOT_ENDPOINT* with your endpoint.
- Replace the two *lightCertId* instances with the certificate ID in the file name for your GG_TrafficLight client device.

```
cd path-to-certs-folder
python trafficLight.py --endpoint AWS_IOT_ENDPOINT --rootCA AmazonRootCA1.pem
  --cert lightCertId-certificate.pem.crt --key lightCertId-private.pem.key --
  thingName GG_TrafficLight --clientId GG_TrafficLight
```

Every 20 seconds, the switch updates the shadow state to G, Y, and R, and the light displays its new state, as shown next.

GG_Switch output:

```
{"state":{"desired":{"property":"R"}}}
2018-12-20 12:23:01,446 - AWSIoTPythonSDK.core.protocol.mqtt_core - INFO - Performing sync publish...
~~~~~Shadow Update Accepted~~~~~
Update request with token: 3b22e27c-930d-4c6a-8562-9f86088249f4 accepted!
property: R
~~~~~
```

GG_TrafficLight output:

```

+++++ Received Shadow Delta +++++
{u'state': {u'property': u'R'}, u'metadata': {u'property': {u'timestamp': 1545337381}}, u'version': 33, u'clientToken':
u'3b22e27c-930d-4c6a-8562-9f86088249f4'}
property: R
version: 33
+++++

Light changed to: R
{"state":{"reported":{"property":"R"}}}
2018-12-20 12:23:01,539 - AWSIoTPythonSDK.core.protocol.mqtt_core - INFO - Performing sync publish...
~~~~~ Shadow Update Accepted ~~~~~
Update request with token: f552109f-c1c2-4ae6-a841-8443506eefcb accepted!
property: R
~~~~~

```

When executed for the first time, each client device script runs the AWS IoT Greengrass discovery service to connect to the AWS IoT Greengrass core (through the internet). After a client device has discovered and successfully connected to the AWS IoT Greengrass core, future operations can be executed locally.

Note

The `lightController.py` and `trafficLight.py` scripts store connection information in the `groupCA` folder, which is created in the same folder as the scripts. If you receive connection errors, make sure that the IP address in the `ggc-host` file matches the IP address endpoint for your core.

4. In the AWS IoT console, choose your AWS IoT Greengrass group, choose the **Client devices** tab, and then choose **GG_TrafficLight** to open the client device's AWS IoT thing details page.
5. Choose the **Device Shadows** tab. After the GG_Switch changes states, there should not be any updates to this shadow. That's because the GG_TrafficLight is set to **Disable shadow sync with cloud**.
6. Press **Ctrl + C** in the GG_Switch (`lightController.py`) client device window. You should see that the GG_TrafficLight (`trafficLight.py`) window stops receiving state change messages.

Keep these windows open so you can run the commands in the next section.

Test communications (device syncs enabled)

For this test, you configure the GG_TrafficLight device shadow to sync to AWS IoT. You run the same commands as in the previous test, but this time the shadow state in the cloud is updated when GG_Switch sends an update request.

1. In the AWS IoT console, choose your AWS IoT Greengrass group, and then choose the **Client devices** tab.
2. Select the GG_TrafficLight device, choose **Sync shadow**, and then choose **Enable shadow sync with cloud**.

You should receive a notification that the device shadow sync status was updated.

3. On the group configuration page, choose **Deploy**.
4. In your two command-line windows, run the commands from the previous test for the [GG_Switch](#) and [GG_TrafficLight](#) client devices.
5. Now, check the shadow state in the AWS IoT console. Choose your AWS IoT Greengrass group, choose the **Client devices** tab, choose **GG_TrafficLight**, choose the **Device Shadows** tab, and then choose **Classic Shadow**.

Because you enabled sync of the GG_TrafficLight shadow to AWS IoT, the shadow state in the cloud should be updated whenever GG_Switch sends an update. This functionality can be used to expose the state of a client device to AWS IoT.

Note

If necessary, you can troubleshoot issues by viewing the AWS IoT Greengrass core logs, particularly `runtime.log`:


```
cd /greengrass/ggc/var/log
sudo cat system/runtime.log | more
```

You can also view `GGShadowSyncManager.log` and `GGShadowService.log`. For more information, see [Troubleshooting](#).

Keep the client devices and subscriptions set up. You use them in the next module. You also run the same commands.

Module 6: Accessing other AWS services

This advanced module shows you how AWS IoT Greengrass cores can interact with other AWS services in the cloud. It builds on the traffic light example from [Module 5](#) and adds a Lambda function that processes shadow states and uploads a summary to an Amazon DynamoDB table.



Before you begin, run the [Greengrass device setup](#) script, or make sure that you have completed [Module 1](#) and [Module 2](#). You should also complete [Module 5](#). You do not need other components or devices.

This module should take about 30 minutes to complete.

Note

This module creates and updates a table in DynamoDB. Although most of the operations are small and fall within the Amazon Web Services Free Tier, performing some of the steps

in this module might result in charges to your account. For information about pricing, see [DynamoDB pricing documentation](#).

Topics

- [Configure the group role](#)
- [Create and configure the Lambda function](#)
- [Configure subscriptions](#)
- [Test communications](#)

Configure the group role

The group role is an [IAM role](#) that you create and attach to your Greengrass group. This role contains the permissions that deployed Lambda functions (and other AWS IoT Greengrass features) use to access AWS services. For more information, see [the section called “Greengrass group role”](#).

You use the following high-level steps to create a group role in the IAM console.

1. Create a policy that allows or denies actions on one or more resources.
2. Create a role that uses the Greengrass service as a trusted entity.
3. Attach your policy to the role.

Then, in the AWS IoT console, you add the role to the Greengrass group.

Note

A Greengrass group has one group role. If you want to add permissions, you can edit attached policies or attach more policies.

For this tutorial, you create a permissions policy that allows describe, create, and update actions on an Amazon DynamoDB table. Then, you attach the policy to a new role and associate the role with your Greengrass group.

First, create a customer-managed policy that grants permissions required by the Lambda function in this module.

1. In the IAM console, in the navigation pane, choose **Policies**, and then choose **Create policy**.
2. On the **JSON** tab, replace the placeholder content with the following policy. The Lambda function in this module uses these permissions to create and update a DynamoDB table named `CarStats`.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "PermissionsForModule6",
      "Effect": "Allow",
      "Action": [
        "dynamodb:DescribeTable",
        "dynamodb:CreateTable",
        "dynamodb:PutItem"
      ],
      "Resource": "arn:aws:dynamodb:*:*:table/CarStats"
    }
  ]
}
```

3. Choose **Next: Tags**, and then choose **Next: Review**. Tags aren't used in this tutorial.
4. For **Name**, enter **greengrass_CarStats_Table**, and then choose **Create policy**.

Next, create a role that uses the new policy.

5. In the navigation pane, choose **Roles**, and then choose **Create role**.
6. Under **Trusted entity type**, choose **AWS service**.
7. Under **Use case**, **Use cases for other AWS services** choose **Greengrass**, select **Greengrass**, and then choose **Next**.
8. Under **Permissions policies**, select the new **greengrass_CarStats_Table** policy, and then choose **Next**.
9. For **Role name**, enter **Greengrass_Group_Role**.
10. For **Description**, enter **Greengrass group role for connectors and user-defined Lambda functions**.

11. Choose **Create role**.

Now, add the role to your Greengrass group.

12. In the AWS IoT console navigation pane, under **Manage**, expand **Greengrass devices**, and then choose **Groups (V1)**.

13. Under **Greengrass groups**, choose your group.

14. Choose **Settings**, and then choose **Associate role**.

15. Choose **Greengrass_Group_Role** from your list of roles, and then choose **Associate role**.

Create and configure the Lambda function

In this step, you create a Lambda function that tracks the number of cars that pass the traffic light. Every time that the GG_TrafficLight shadow state changes to G, the Lambda function simulates the passing of a random number of cars (from 1 to 20). On every third G light change, the Lambda function sends basic statistics, such as min and max, to a DynamoDB table.

1. On your computer, create a folder named `car_aggregator`.
2. From the [TrafficLight](#) examples folder on GitHub, download the `carAggregator.py` file to the `car_aggregator` folder. This is your Lambda function code.

Note

This example Python file is stored in the AWS IoT Greengrass Core SDK repository for convenience, but it doesn't use the AWS IoT Greengrass Core SDK.

3. If you aren't working in the US East (N. Virginia) Region, open `carAggregator.py` and change `region_name` in the following line to the AWS Region that's currently selected in the AWS IoT console. For the list of supported AWS Regions, see [AWS IoT Greengrass](#) in the *Amazon Web Services General Reference*.

```
dynamodb = boto3.resource('dynamodb', region_name='us-east-1')
```

4. Run the following command in a [command-line](#) window to install the [AWS SDK for Python \(Boto3\)](#) package and its dependencies in the `car_aggregator` folder. Greengrass Lambda functions use the AWS SDK to access other AWS services. (For Windows, use an [elevated command prompt](#).)


```
pip install boto3 -t path-to-car_aggregator-folder
```

This results in a directory listing similar to the following:

Name	Date modified	Type
 bin	12/31/2018 2:27 PM	File folder
 boto3	12/31/2018 2:27 PM	File folder
 boto3-1.9.71.dist-info	12/31/2018 2:27 PM	File folder
 botocore	12/31/2018 2:27 PM	File folder
 botocore-1.12.71.dist-info	12/31/2018 2:27 PM	File folder
 concurrent	12/31/2018 2:27 PM	File folder
 dateutil	12/31/2018 2:27 PM	File folder
 docutils	12/31/2018 2:27 PM	File folder
 docutils-0.14.dist-info	12/31/2018 2:27 PM	File folder
 futures-3.2.0.dist-info	12/31/2018 2:27 PM	File folder
 jmespath	12/31/2018 2:27 PM	File folder
 jmespath-0.9.3.dist-info	12/31/2018 2:27 PM	File folder
 python_dateutil-2.7.5.dist-info	12/31/2018 2:27 PM	File folder
 s3transfer	12/31/2018 2:27 PM	File folder
 s3transfer-0.1.13.dist-info	12/31/2018 2:27 PM	File folder
 six-1.12.0.dist-info	12/31/2018 2:27 PM	File folder
 urllib3	12/31/2018 2:27 PM	File folder
 urllib3-1.24.1.dist-info	12/31/2018 2:27 PM	File folder
 carAggregator.py	12/31/2018 2:25 PM	PY File
 six.py	12/31/2018 2:27 PM	PY File
 six.pyc	12/31/2018 2:27 PM	Compiled Python ...

- Compress the contents of the `car_aggregator` folder into a `.zip` file named `car_aggregator.zip`. (Compress the folder's contents, not the folder.) This is your Lambda function deployment package.
- In the Lambda console, create a function named **GG_Car_Aggregator**, and set the remaining fields as follows:
 - For **Runtime**, choose **Python 3.7**.
 - For **Permissions**, keep the default setting. This creates an execution role that grants basic Lambda permissions. This role isn't used by AWS IoT Greengrass.

Choose **Create function**.

Basic information

Function name
Enter a name that describes the purpose of your function.
GG_Car_Aggregator
Use only letters, numbers, hyphens, or underscores with no spaces.


Runtime [Info](#)
Choose the language to use to write your function.
Python 3.7

Permissions [Info](#)
Lambda will create an execution role with permission to upload logs to Amazon CloudWatch Logs. You can configure and modify permissions further when you add triggers.
▶ Choose or create an execution role

Cancel **Create function**

7. Upload your Lambda function deployment package:

- a. On the **Code** tab, under **Code source**, choose **Upload from**. From the dropdown, choose **.zip file**.



- b. Choose upload, and then choose your `car_aggregator.zip` deployment package. Then, choose **Save**.
 - c. On the **Code** tab for the function, under **Runtime settings**, choose **Edit**, and then enter the following values.
 - For **Runtime**, choose **Python 3.7**.
 - For **Handler**, enter `carAggregator.function_handler`
 - d. Choose **Save**.
8. Publish the Lambda function, and then create an alias named **GG_CarAggregator**. For step-by-step instructions, see the steps to [publish the Lambda function](#) and [create an alias](#) in Module 3 (Part 1).
9. In the AWS IoT console, add the Lambda function that you just created to your AWS IoT Greengrass group:

- a. On the group configuration page, choose **Lambda functions**, and then under **My Lambda functions**, choose **Add**.
- b. For **Lambda function**, choose **GG_Car_Aggregator**.
- c. For **Lambda function version**, choose the alias to the version that you published.
- d. For **Memory limit**, enter **64 MB**.
- e. For **Pinned**, choose **True**.
- f. Choose **Add Lambda function**.

 Note

You can remove other Lambda functions from earlier modules.

Configure subscriptions

In this step, you create a subscription that enables the GG_TrafficLight shadow to send updated state information to the GG_Car_Aggregator Lambda function. This subscription is added to the subscriptions that you created in [Module 5](#), which are all required for this module.

1. On the group configuration page, choose the **Subscriptions** tab, and then choose **Add**.
2. On the **Create a subscription** page, do the following:
 - a. For **Source type**, choose **Service**, and then choose **Local Shadow Service**.
 - b. For **Target type**, choose **Lambda function**, and then choose **GG_Car_Aggregator**.
 - c. For **Topic filter**, enter **\$aws/things/GG_TrafficLight/shadow/update/documents**
 - d. Choose **Create subscription**.

This module requires the new subscription and the [subscriptions](#) that you created in Module 5.

3. Make sure that the Greengrass daemon is running, as described in [Deploy cloud configurations to a core device](#).
4. On the group configuration page, choose **Deploy**.

Test communications

1. On your computer, open two [command-line](#) windows. Just as in [Module 5](#), one window is for the GG_Switch client device and the other is for the GG_TrafficLight client device. You use them to run the same commands that you ran in Module 5.

Run the following commands for the GG_Switch client device:

```
cd path-to-certs-folder  
python lightController.py --endpoint AWS_IOT_ENDPOINT --rootCA AmazonRootCA1.pem  
--cert switchCertId-certificate.pem.crt --key switchCertId-private.pem.key --  
thingName GG_TrafficLight --clientId GG_Switch
```

Run the following commands for the GG_TrafficLight client device:

```
cd path-to-certs-folder  
python trafficLight.py --endpoint AWS_IOT_ENDPOINT --rootCA AmazonRootCA1.pem --  
cert lightCertId-certificate.pem.crt --key lightCertId-private.pem.key --thingName  
GG_TrafficLight --clientId GG_TrafficLight
```

Every 20 seconds, the switch updates the shadow state to G, Y, and R, and the light displays its new state.

2. The function handler of the Lambda function is triggered on every third green light (every three minutes), and a new DynamoDB record is created. After `lightController.py` and `trafficLight.py` have run for three minutes, go to the AWS Management Console, and open the DynamoDB console.
3. Choose **US East (N. Virginia)** in the AWS Region menu. This is the Region where the GG_Car_Aggregator function creates the table.
4. In the navigation pane, choose **Tables**, and then choose the **CarStats** table.
5. Choose **View items** to view the entries in the table.

You should see entries with basic statistics on cars passed (one entry for every three minutes). You might need to choose the refresh button to view updates to the table.

6. If the test is not successful, you can look for troubleshooting information in the Greengrass logs.

- a. Switch to the root user and navigate to the `log` directory. Access to AWS IoT Greengrass logs requires root permissions.

```
sudo su
cd /greengrass/ggc/var/log
```

- b. Check `runtime.log` for errors.

```
cat system/runtime.log | grep 'ERROR'
```

- c. Check the log generated by the Lambda function.

```
cat user/region/account-id/GG_Car_Aggregator.log
```

The `lightController.py` and `trafficLight.py` scripts store connection information in the `groupCA` folder, which is created in the same folder as the scripts. If you receive connection errors, make sure that the IP address in the `ggc-host` file matches the IP address endpoint for your core.

For more information, see [Troubleshooting](#).

This is the end of the basic tutorial. You should now understand the AWS IoT Greengrass programming model and its fundamental concepts, including AWS IoT Greengrass cores, groups, subscriptions, client devices, and the deployment process for Lambda functions running at the edge.

You can delete the DynamoDB table and the Greengrass Lambda functions and subscriptions. To stop communications between the AWS IoT Greengrass core device and the AWS IoT cloud, open a terminal on the core device and run one of the following commands:

- To shut down the AWS IoT Greengrass core device:

```
sudo halt
```

- To stop the AWS IoT Greengrass daemon:

```
cd /greengrass/ggc/core/
```

```
sudo ./greengrassd stop
```

Module 7: Simulating hardware security integration

This feature is available for AWS IoT Greengrass Core v1.7 and later.

This advanced module shows you how to configure a simulated hardware security module (HSM) for use with a Greengrass core. The configuration uses SoftHSM, which is a pure software implementation that uses the [PKCS#11](#) application programming interface (API). The purpose of this module is to allow you to set up an environment where you can learn and do initial testing against a software-only implementation of the PKCS#11 API. It is provided only for learning and initial testing, not for production use of any kind.

You can use this configuration to experiment with using a PKCS#11-compatible service to store your private keys. For more information about the software-only implementation, see [SoftHSM](#). For more information about integrating hardware security on an AWS IoT Greengrass core, including general requirements, see [the section called “Hardware security integration”](#).

Important

This module is intended for experimentation purposes only. We strongly discourage the use of SoftHSM in a production environment because it might provide a false sense of additional security. The resulting configuration doesn't provide any actual security benefits. The keys stored in SoftHSM are not stored more securely than any other means of secrets storage in the Greengrass environment.

The purpose of this module is to allow you to learn about the PKCS#11 specification and do initial testing of your software if you plan to use a real hardware-based HSM in the future. You must test your future hardware implementation separately and completely before any production usage because there might be differences between the PKCS#11 implementation provided in SoftHSM and a hardware-based implementation.

If you need assistance with the onboarding of a [supported hardware security module](#), contact your AWS Enterprise Support representative.

Before you begin, run the [Greengrass Device Setup](#) script, or make sure that you completed [Module 1](#) and [Module 2](#) of the Getting Started tutorial. In this module, we assume that your core is already provisioned and communicating with AWS. This module should take about 30 minutes to complete.

Install the SoftHSM software

In this step, you install SoftHSM and the pkcs11 tools, which are used to manage your SoftHSM instance.

- In a terminal on your AWS IoT Greengrass core device, run the following command:

```
sudo apt-get install softhsm2 libsofthsm2-dev pkcs11-dump
```

For more information about these packages, see [Install softhsm2](#), [Install libsofthsm2-dev](#), and [Install pkcs11-dump](#).

Note

If you encounter issues when using this command on your system, see [SoftHSM version 2](#) on GitHub. This site provides more installation information, including how to build from source.

Configure SoftHSM

In this step, you [configure SoftHSM](#).

1. Switch to the root user.

```
sudo su
```

2. Use the manual page to find the system-wide `softhsm2.conf` location. A common location is `/etc/softhsm/softhsm2.conf`, but the location might be different on some systems.

```
man softhsm2.conf
```

3. Create the directory for the `softhsm2` configuration file in the system-wide location. In this example, we assume the location is `/etc/softhsm/softhsm2.conf`.

```
mkdir -p /etc/softhsm
```

4. Create the token directory in the `/greengrass` directory.

Note

If this step is skipped, softhsm2-util reports ERROR: Could not initialize the library.

```
mkdir -p /greengrass/softhsm2/tokens
```

5. Configure the token directory.

```
echo "directories.tokenidir = /greengrass/softhsm2/tokens" > /etc/softhsm/softhsm2.conf
```

6. Configure a file-based backend.

```
echo "objectstore.backend = file" >> /etc/softhsm/softhsm2.conf
```

Note

These configuration settings are intended for experimentation purposes only. To see all configuration options, read the manual page for the configuration file.

```
man softhsm2.conf
```

Import the private key into SoftHSM

In this step, you initialize the SoftHSM token, convert the private key format, and then import the private key.

1. Initialize the SoftHSM token.

```
softhsm2-util --init-token --slot 0 --label greengrass --so-pin 12345 --pin 1234
```


Note

If prompted, enter an SO pin of 12345 and a user pin of 1234. AWS IoT Greengrass doesn't use the SO (supervisor) pin, so you can use any value.

If you receive the error `CKR_SLOT_ID_INVALID: Slot 0 does not exist`, try the following command instead:

```
softhsm2-util --init-token --free --label greengrass --so-pin 12345 --pin 1234
```

2. Convert the private key to a format that can be used by the SoftHSM import tool. For this tutorial, you convert the private key that you obtained from the **Default Group creation** option in [Module 2](#) of the Getting Started tutorial.

```
openssl pkcs8 -topk8 -inform PEM -outform PEM -nocrypt -in hash.private.key -out hash.private.pem
```

3. Import the private key into SoftHSM. Run only one of the following commands, depending on your version of softhsm2-util.

Raspbian softhsm2-util v2.2.0 syntax

```
softhsm2-util --import hash.private.pem --token greengrass --label iotkey --id 0000 --pin 12340
```

Ubuntu softhsm2-util v2.0.0 syntax

```
softhsm2-util --import hash.private.pem --slot 0 --label iotkey --id 0000 --pin 1234
```

This command identifies the slot as 0 and defines the key label as `iotkey`. You use these values in the next section.

After the private key is imported, you can optionally remove it from the `/greengrass/certs` directory. Make sure to keep the root CA and device certificates in the directory.

Configure the Greengrass core to use SoftHSM

In this step, you modify the Greengrass core configuration file to use SoftHSM.

1. Find the path to the SoftHSM provider library (`libsofthsm2.so`) on your system:

- a. Get the list of installed packages for the library.

```
sudo dpkg -L libsofthsm2
```

The `libsofthsm2.so` file is located in the `softhsm` directory.

- b. Copy the full path to the file (for example, `/usr/lib/x86_64-linux-gnu/softhsm/libsofthsm2.so`). You use this value later.

2. Stop the Greengrass daemon.

```
cd /greengrass/ggc/core/  
sudo ./greengrassd stop
```

3. Open the Greengrass configuration file. This is the [config.json](#) file in the `/greengrass/config` directory.

Note

The examples in this procedure are written with the assumption that the `config.json` file uses the format that's generated from the **Default Group creation** option in [Module 2](#) of the Getting Started tutorial.

4. In the `crypto.principals` object, insert the following MQTT server certificate object. Add a comma where needed to create a valid JSON file.

```
"MQTTServerCertificate": {  
  "privateKeyPath": "path-to-private-key"  
}
```

5. In the `crypto` object, insert the following PKCS11 object. Add a comma where needed to create a valid JSON file.

```
"PKCS11": {  
  "P11Provider": "/path-to-pkcs11-provider-so",
```

```

    "slotLabel": "crypto-token-name",
    "slotUserPin": "crypto-token-user-pin"
  }

```

Your file should look similar to the following:

```

{
  "coreThing" : {
    "caPath" : "root.ca.pem",
    "certPath" : "hash.cert.pem",
    "keyPath" : "hash.private.key",
    "thingArn" : "arn:partition:iot:region:account-id:thing/core-thing-name",
    "iotHost" : "host-prefix.iot.region.amazonaws.com",
    "ggHost" : "greengrass.iot.region.amazonaws.com",
    "keepAlive" : 600
  },
  "runtime" : {
    "cgroup" : {
      "useSystemd" : "yes"
    }
  },
  "managedRespawn" : false,
  "crypto": {
    "PKCS11": {
      "P11Provider": "/path-to-pkcs11-provider-so",
      "slotLabel": "crypto-token-name",
      "slotUserPin": "crypto-token-user-pin"
    },
    "principals" : {
      "MQTTServerCertificate": {
        "privateKeyPath": "path-to-private-key"
      },
      "IoTCertificate" : {
        "privateKeyPath" : "file:///greengrass/certs/hash.private.key",
        "certificatePath" : "file:///greengrass/certs/hash.cert.pem"
      },
      "SecretsManager" : {
        "privateKeyPath" : "file:///greengrass/certs/hash.private.key"
      }
    },
    "caPath" : "file:///greengrass/certs/root.ca.pem"
  }
}

```

Note

To use over-the-air (OTA) updates with hardware security, the PKCS11 object must also contain the OpenSSLEngine property. For more information, see [the section called "Configure OTA updates"](#).

6. Edit the crypto object:**a. Configure the PKCS11 object.**

- For `P11Provider`, enter the full path to `libsofthsm2.so`.
- For `slotLabel`, enter `greengrass`.
- For `slotUserPin`, enter `1234`.

b. Configure the private key paths in the principals object. Do not edit the `certificatePath` property.

- For the `privateKeyPath` properties, enter the following RFC 7512 PKCS#11 path (which specifies the key's label). Do this for the `IoTCertificate`, `SecretsManager`, and `MQTTServerCertificate` principals.

```
pkcs11:object=iotkey;type=private
```

c. Check the crypto object. It should look similar to the following:

```
"crypto": {
  "PKCS11": {
    "P11Provider": "/usr/lib/x86_64-linux-gnu/softhsm/libsofthsm2.so",
    "slotLabel": "greengrass",
    "slotUserPin": "1234"
  },
  "principals": {
    "MQTTServerCertificate": {
      "privateKeyPath": "pkcs11:object=iotkey;type=private"
    },
    "SecretsManager": {
      "privateKeyPath": "pkcs11:object=iotkey;type=private"
    },
    "IoTCertificate": {
      "certificatePath": "file://certs/core.crt",
```

```
        "privateKeyPath": "pkcs11:object=iotkey;type=private"
      },
      "caPath": "file://certs/root.ca.pem"
    }
  }
```

7. Remove the `caPath`, `certPath`, and `keyPath` values from the `coreThing` object. It should look similar to the following:

```
"coreThing" : {
  "thingArn" : "arn:partition:iot:region:account-id:thing/core-thing-name",
  "iotHost" : "host-prefix-ats.iot.region.amazonaws.com",
  "ggHost" : "greengrass-ats.iot.region.amazonaws.com",
  "keepAlive" : 600
}
```

Note

For this tutorial, you specify the same private key for all principals. For more information about choosing the private key for the local MQTT server, see [Performance](#). For more information about the local secrets manager, see [Deploy secrets to the core](#).

Test the configuration

- Start the Greengrass daemon.

```
cd /greengrass/ggc/core/
sudo ./greengrassd start
```

If the daemon starts successfully, then your core is configured correctly.

You are now ready to learn about the PKCS#11 specification and do initial testing with the PKCS#11 API that's provided by the SoftHSM implementation.

⚠ Important

Again, it's extremely important to be aware that this module is intended for learning and testing only. It doesn't actually increase the security posture of your Greengrass environment.

Instead, the purpose of the module is to enable you to start learning and testing in preparation for using a true hardware-based HSM in the future. At that time, you must separately and completely test your software against the hardware-based HSM prior to any production usage, because there might be differences between the PKCS#11 implementation provided in SoftHSM and a hardware-based implementation.

See also

- *PKCS #11 Cryptographic Token Interface Usage Guide Version 2.40*. Edited by John Leiseboer and Robert Griffin. 16 November 2014. OASIS Committee Note 02. <http://docs.oasis-open.org/pkcs11/pkcs11-ug/v2.40/cn02/pkcs11-ug-v2.40-cn02.html>. Latest version: <http://docs.oasis-open.org/pkcs11/pkcs11-ug/v2.40/pkcs11-ug-v2.40.html>.
- [RFC 7512](#)

OTA updates of AWS IoT Greengrass Core software

The AWS IoT Greengrass Core software package includes an update agent that can perform over-the-air (OTA) updates of AWS IoT Greengrass software. You can use OTA updates to install the latest version of the AWS IoT Greengrass Core software or OTA update agent software on one or more cores. With OTA updates, your core devices don't have to be physically present.

We recommend that you use OTA updates when possible. They provide a mechanism you can use to track update status and update history. If a failed update occurs, the OTA update agent rolls back to the previous software version.

Note

OTA updates are not supported when you use apt to install the AWS IoT Greengrass Core software. For these installations, we recommend that you use apt to upgrade the software. For more information, see [the section called "Install from an APT repository"](#).

OTA updates make it more efficient to:

- Fix security vulnerabilities.
- Address software stability issues.
- Deploy new or improved features.

This feature integrates with [AWS IoT jobs](#).

Requirements

The following requirements apply for OTA updates of AWS IoT Greengrass software.

- The Greengrass core must have at least 400 MB of disk space available in local storage. The OTA update agent requires about three times the runtime usage requirement of the AWS IoT Greengrass Core software. For more information, see [Service quotas](#) for the Greengrass core in the *Amazon Web Services General Reference*.
- The Greengrass core must have a connection to the AWS Cloud.

- The Greengrass core must be correctly configured and provisioned with certificates and keys for authentication with AWS IoT Core and AWS IoT Greengrass. For more information, see [the section called “X.509 certificates”](#).
- The Greengrass core can't be configured to use a network proxy.

Note

Starting in AWS IoT Greengrass v1.9.3, OTA updates are supported on cores that configure MQTT traffic to use port 443 instead of the default port 8883. However, the OTA update agent does not support updates through a network proxy. For more information, see [the section called “Connect on port 443 or through a network proxy”](#).

- Trusted boot can't be enabled in the partition that contains the AWS IoT Greengrass Core software.

Note

You can install and run the AWS IoT Greengrass Core software on a partition that has trusted boot enabled, but OTA updates aren't supported.

- AWS IoT Greengrass must have read/write permissions on the partition that contains the AWS IoT Greengrass Core software.
- If you use an init system to manage your Greengrass core, you must configure OTA updates to integrate with the init system. For more information, see [the section called “Integration with init systems”](#).
- You must create a role that's used to presign the Amazon S3 URLs to AWS IoT Greengrass software update artifacts. This signer role allows AWS IoT Core to access software update artifacts stored in Amazon S3 on your behalf. For more information, see [the section called “IAM permissions for OTA updates”](#).

IAM permissions for OTA updates

When AWS IoT Greengrass releases a new version of the AWS IoT Greengrass Core software, AWS IoT Greengrass updates the software artifacts stored in Amazon S3 that are used for the OTA update.

Your AWS account must include an Amazon S3 URL signer role that can be used to access these artifacts. The role must have a permissions policy that allows the `s3:GetObject` action on the buckets in target AWS Regions. The role must also have a trust policy that allows `iot.amazonaws.com` to assume the role as a trusted entity.

Permissions policy

For role permissions, you can use the AWS managed policy or create a custom policy.

- **Use the AWS managed policy**

The [GreengrassOTAUpdateArtifactAccess](#) managed policy is provided by AWS IoT Greengrass. Use this policy if you want to allow access in all Amazon Web Services Regions supported by AWS IoT Greengrass, both current and future.

- **Create a custom policy**

You should create a custom policy if you want to explicitly specify the Amazon Web Services Regions where your cores are deployed. The following example policy allows access to AWS IoT Greengrass software updates in six Regions.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "AllowAccessToGreengrassOTAUpdateArtifacts",
      "Effect": "Allow",
      "Action": [
        "s3:GetObject"
      ],
      "Resource": [
        "arn:aws:s3:::us-east-1-greengrass-updates/*",
        "arn:aws:s3:::us-west-2-greengrass-updates/*",
        "arn:aws:s3:::ap-northeast-1-greengrass-updates/*",
        "arn:aws:s3:::ap-southeast-2-greengrass-updates/*",
        "arn:aws:s3:::eu-central-1-greengrass-updates/*",
        "arn:aws:s3:::eu-west-1-greengrass-updates/*"
      ]
    }
  ]
}
```

Trust policy

The trust policy attached to the role must allow the `sts:AssumeRole` action and define `iot.amazonaws.com` as a principal. This allows AWS IoT Core to assume the role as a trusted entity. Here's an example policy document:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "AllowIotToAssumeRole",
      "Action": "sts:AssumeRole",
      "Principal": {
        "Service": "iot.amazonaws.com"
      },
      "Effect": "Allow"
    }
  ]
}
```

In addition, the user who initiates an OTA update must have permissions to use `greengrass:CreateSoftwareUpdateJob` and `iot:CreateJob`, and to use `iam:PassRole` to pass the permissions of the signer role. Here's an example IAM policy:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Action": [
        "greengrass:CreateSoftwareUpdateJob"
      ],
      "Effect": "Allow",
      "Resource": "*"
    },
    {
      "Effect": "Allow",
      "Action": [
        "iot:CreateJob"
      ],
      "Resource": "*"
    }
  ],
}
```

```
{
  "Effect": "Allow",
  "Action": [
    "iam:PassRole"
  ],
  "Resource": "arn-of-s3-url-signer-role"
}
]
```

Considerations

Before you launch an OTA update of Greengrass Core software, be aware of the impact on the devices in your Greengrass group, both on the core device and on client devices connected locally to that core:

- The core shuts down during the update.
- Any Lambda functions running on the core are shut down. If those functions write to local resources, they might leave those resources in an incorrect state unless shut down properly.
- During the core's downtime, all its connections with the AWS Cloud are lost. Messages routed through the core by client devices are lost.
- Credential caches are lost.
- Queues that hold pending work for Lambda functions are lost.
- Long-lived Lambda functions lose their dynamic state information and all pending work is dropped.

The following state information is preserved during an OTA update:

- Core configuration
- Greengrass group configuration
- Local shadows
- Greengrass logs
- OTA update agent logs

Greengrass OTA update agent

The Greengrass OTA update agent is the software component on the device that handles update jobs created and deployed in the cloud. The OTA update agent is distributed in the same software package as the AWS IoT Greengrass Core software. The agent is located in `/greengrass-root/ota/ota_agent/ggc-ota`. It writes logs to `/var/log/greengrass/ota/ggc_ota.txt`.

Note

`greengrass-root` represents the path where the AWS IoT Greengrass Core software is installed on your device. Typically, this is the `/greengrass` directory.

You can start the OTA update agent by executing the binary manually or by integrating it as part of an init script, such as a systemd service file. If you execute the binary manually, it should be run as root. When it starts, the OTA update agent listens for AWS IoT Greengrass software update jobs from AWS IoT Core and executes them sequentially. The OTA update agent ignores all other AWS IoT job types.

The following excerpt shows an example of a systemd service file to start, stop, and restart the OTA update agent:

```
[Unit]
Description=Greengrass OTA Daemon

[Service]
Type=forking
Restart=on-failure
ExecStart=/greengrass/ota/ota_agent/ggc-ota

[Install]
WantedBy=multi-user.target
```

A core that is the target of an update must not run two instances of the OTA update agent. Doing so causes the two agents to process the same jobs, which creates conflicts.

Integration with init systems

During an OTA update, the OTA update agent restarts binaries on the core device. If the binaries are running, this might cause conflicts when an init system is monitoring the state of the AWS IoT Greengrass Core software or the agent during the update. To help integrate the OTA update mechanism with your init monitoring strategies, you can write shell scripts that run before and after an update. For example, you can use the `ggc_pre_update.sh` script to back up data or stop processes before the device shuts down.

To tell the OTA update agent to run these scripts, you must include the `"managedRespawn"` : `true` flag in the [config.json](#) file. This setting is shown in the following excerpt:

```
{
  "coreThing": {
    ...
  },
  "runtime": {
    ...
  },
  "managedRespawn": true
  ...
}
```

Managed respawn with OTA updates

The following requirements apply to OTA updates with `managedRespawn` set to `true`:

- The following shell scripts must be present in the `/greengrass-root/usr/scripts` directory:
 - `ggc_pre_update.sh`
 - `ggc_post_update.sh`
 - `ota_pre_update.sh`
 - `ota_post_update.sh`
- The scripts must return a successful return code.
- The scripts must be owned by root and executable by root only.
- The `ggc_pre_update.sh` script must stop the Greengrass daemon.
- The `ggc_post_update.sh` script must start the Greengrass daemon.

Note

Because the OTA update agent manages its own process, the `ota_pre_update.sh` and `ota_post_update.sh` scripts do not need to stop or start the OTA service.

The OTA update agent runs the scripts from the `/greengrass-root/usr/scripts`. The directory tree should look like the following:

```
<greengrass_root>
|-- certs
|-- config
|   |-- config.json
|-- ggc
|-- usr/scripts
|   |-- ggc_pre_update.sh
|   |-- ggc_post_update.sh
|   |-- ota_pre_update.sh
|   |-- ota_post_update.sh
|-- ota
```

When `managedRespawn` is set to `true`, the OTA update agent checks the `/greengrass-root/usr/scripts` directory for these scripts before and after the software update. If the scripts don't exist, the update fails. AWS IoT Greengrass does not validate the contents of these scripts. As a best practice, verify that your scripts function correctly and issue appropriate exit codes for errors.

For OTA updates of the AWS IoT Greengrass Core software:

- Before starting the update, the agent runs the `ggc_pre_update.sh` script. Use this script for commands that need to run before the OTA update agent starts the AWS IoT Greengrass Core software update, such as to back up data or stop any running processes. The following example shows a simple script to stop the Greengrass daemon.

```
#!/bin/bash
set -euo pipefail
systemctl stop greengrass
```

- After completing the update, the agent runs the `ggc_post_update.sh` script. Use this script for commands that need to run after the OTA update agent starts the AWS IoT Greengrass Core

software update, such as to restart processes. The following example shows a simple script to start the Greengrass daemon.

```
#!/bin/bash
set -euo pipefail
systemctl start greengrass
```

For OTA updates of the OTA update agent:

- Before starting the update, the agent runs the `ota_pre_update.sh` script. Use this script for commands that need to run before the OTA update agent updates itself, such as to back up data or stop any running processes.
- After completing the update, the agent runs the `ota_post_update.sh` script. Use this script for commands that need to run after the OTA update agent updates itself, such as to restart processes.

Note

If `managedRespawn` is set to `false`, the OTA update agent does not run the scripts.

Create an OTA update

Follow these steps to perform an OTA update of AWS IoT Greengrass software on one or more cores:

1. Make sure that your cores meet the [requirements](#) for OTA updates.

Note

If you configured an init system to manage the AWS IoT Greengrass Core software or the OTA update agent, verify the following on your cores:

- The [config.json](#) file specifies `"managedRespawn" : true`.
- The `/greengrass-root/usr/scripts` directory contains the following scripts:
 - `ggc_pre_update.sh`

- ggc_post_update.sh
- ota_pre_update.sh
- ota_post_update.sh

For more information, see [the section called "Integration with init systems"](#).

2. In a core device terminal, start the OTA update agent.

```
cd /greengrass-root/ota/ota_agent  
sudo ./ggc-ota
```

 Note

greengrass-root represents the path where the AWS IoT Greengrass Core software is installed on your device. Typically, this is the /greengrass directory.

Don't start multiple instances of the OTA update agent on a core because it might cause conflicts.

3. Use the AWS IoT Greengrass API to create a software update job.
 - a. Call the [CreateSoftwareUpdateJob](#) API. In this example procedure, we use AWS CLI commands.

The following command creates a job that updates the AWS IoT Greengrass Core software on one core. Replace the example values and then run the command.

Linux or macOS terminal

```
aws greengrass create-software-update-job \  
--update-targets-architecture x86_64 \  
--update-targets ["arn:aws:iot:region:123456789012:thing/myCoreDevice"] \  
--update-targets-operating-system ubuntu \  
--software-to-update core \  
--s3-url-signer-role arn:aws:iam::123456789012:role/myS3UrlSignerRole \  
--update-agent-log-level WARN \  
--amzn-client-token myClientToken1
```


Windows command prompt

```
aws greengrass create-software-update-job ^  
--update-targets-architecture x86_64 ^  
--update-targets ["arn:aws:iot:region:123456789012:thing/myCoreDevice\""] ^  
--update-targets-operating-system ubuntu ^  
--software-to-update core ^  
--s3-url-signer-role arn:aws:iam::123456789012:role/myS3UrlSignerRole ^  
--update-agent-log-level WARN ^  
--amzn-client-token myClientToken1
```

The command returns the following response.

```
{  
  "IotJobId": "GreengrassUpdateJob_c3bd7f36-ee80-4d42-8321-a1da0EXAMPLE",  
  "IotJobArn": "arn:aws:iot:region:123456789012:job/  
GreengrassUpdateJob_c3bd7f36-ee80-4d42-8321-a1da0EXAMPLE",  
  "PlatformSoftwareVersion": "1.10.1"  
}
```

- b. Copy the `IotJobId` from the response.
- c. Call [DescribeJob](#) in the AWS IoT Core API to see the job status. Replace the example value with your job ID and then run the command.

```
aws iot describe-job --job-id GreengrassUpdateJob_c3bd7f36-ee80-4d42-8321-  
a1da0EXAMPLE
```

The command returns a response object that contains information about the job, including status and `jobProcessDetails`.

```
{  
  "job": {  
    "jobArn": "arn:aws:iot:region:123456789012:job/  
GreengrassUpdateJob_c3bd7f36-ee80-4d42-8321-a1da0EXAMPLE",  
    "jobId": "GreengrassUpdateJob_c3bd7f36-ee80-4d42-8321-a1da0EXAMPLE",  
    "targetSelection": "SNAPSHOT",  
    "status": "IN_PROGRESS",  
    "targets": [  
      "arn:aws:iot:region:123456789012:thing/myCoreDevice"
```

```

    ],
    "description": "This job was created by Greengrass to update the
Greengrass Cores in the targets with version 1.10.1 of the core software
running on x86_64 architecture.",
    "presignedUrlConfig": {
        "roleArn": "arn:aws::iam::123456789012:role/myS3UrlSignerRole",
        "expiresInSec": 3600
    },
    "jobExecutionsRolloutConfig": {},
    "createdAt": 1588718249.079,
    "lastUpdatedAt": 1588718253.419,
    "jobProcessDetails": {
        "numberOfCanceledThings": 0,
        "numberOfSucceededThings": 0,
        "numberOfFailedThings": 0,
        "numberOfRejectedThings": 0,
        "numberOfQueuedThings": 1,
        "numberOfInProgressThings": 0,
        "numberOfRemovedThings": 0,
        "numberOfTimedOutThings": 0
    },
    "timeoutConfig": {}
}
}

```

For troubleshooting help, see [Troubleshooting](#).

CreateSoftwareUpdateJob API

You can use the CreateSoftwareUpdateJob API to update the AWS IoT Greengrass Core software or OTA update agent software on your core devices. This API creates an AWS IoT snapshot job that notifies devices when an update is available. After you call CreateSoftwareUpdateJob, you can use other AWS IoT job commands to track the software update. For more information, see [Jobs](#) in the *AWS IoT Developer Guide*.

The following example shows how to use the AWS CLI to create a job that updates the AWS IoT Greengrass Core software on a core device:

```

aws greengrass create-software-update-job \
--update-targets-architecture x86_64 \

```

```
--update-targets ["arn:aws:iot:region:123456789012:thing/myCoreDevice\" ] \  
--update-targets-operating-system ubuntu \  
--software-to-update core \  
--s3-url-signer-role arn:aws:iam::123456789012:role/myS3UrlSignerRole \  
--update-agent-log-level WARN \  
--amzn-client-token myClientToken1
```

The `create-software-update-job` command returns a JSON response that contains the job ID, job ARN, and software version that was installed by the update:

```
{  
  "IotJobId": "GreengrassUpdateJob_c3bd7f36-ee80-4d42-8321-a1da0EXAMPLE",  
  "IotJobArn": "arn:aws:iot:region:123456789012:job/GreengrassUpdateJob_c3bd7f36-  
ee80-4d42-8321-a1da0EXAMPLE",  
  "PlatformSoftwareVersion": "1.9.2"  
}
```

For steps that show you how to use `create-software-update-job` to update a core device, see [the section called “Create an OTA update”](#).

The `create-software-update-job` command has the following parameters:

`--update-targets-architecture`

The architecture of the core device.

Valid values: `armv7l`, `armv6l`, `x86_64`, or `aarch64`

`--update-targets`

The cores to update. The list can contain ARNs of individual cores and ARNs of thing groups whose members are cores. For more information about thing groups, see [Static thing groups](#) in the *AWS IoT Developer Guide*.

`--update-targets-operating-system`

The operating system of the core device.

Valid values: `ubuntu`, `amazon_linux`, `raspbian`, or `openwrt`

`--software-to-update`

Specifies whether the core's software or the OTA update agent software should be updated.

Valid values: core or ota_agent

`--s3-url-signer-role`

The ARN of the IAM role used to presign the Amazon S3 URL that links to the AWS IoT Greengrass software update artifacts. The role's attached permissions policy must allow the `s3:GetObject` action on the buckets in the target AWS Regions. The role must also allow `iot.amazonaws.com` to assume the role as a trusted entity. For more information, see [the section called "IAM permissions for OTA updates"](#).

`--amzn-client-token`

(Optional) A client token used to make idempotent requests. Provide a unique token to prevent duplicate updates from being created because of internal retries.

`--update-agent-log-level`

(Optional) The logging level for log statements generated by the OTA update agent. The default is ERROR.

Valid values: NONE, TRACE, DEBUG, VERBOSE, INFO, WARN, ERROR, or FATAL

Note

CreateSoftwareUpdateJob accepts requests only for the following supported architecture and operating system combinations:

- ubuntu/x86_64
- ubuntu/aarch64
- amazon_linux/x86_64
- raspbian/armv7l
- raspbian/armv6l
- openwrt/aarch64
- openwrt/armv7l


Deploy AWS IoT Greengrass groups to an AWS IoT Greengrass core

Use AWS IoT Greengrass groups to organize entities in your edge environment. You also use groups to control how the entities in the group interact with each other and with the AWS Cloud. For example, only the Lambda functions in the group are deployed for running locally, and only the devices in the group can communicate using the local MQTT server.

A group must include a [core](#), which is an AWS IoT device that runs the AWS IoT Greengrass Core software. The core acts as an edge gateway and provides AWS IoT Core capabilities in the edge environment. Depending on your business need, you can also add the following entities to a group:

- **Client devices.** Represented as things in the AWS IoT registry. These devices must run [FreeRTOS](#) or use the [AWS IoT Device SDK](#) or [AWS IoT Greengrass Discovery API](#) to get connection information for the core. Only client devices that are members of the group can connect to the core.
- **Lambda functions.** User-defined serverless applications that run code on the core. Lambda functions are authored in AWS Lambda and referenced from a Greengrass group. For more information, see [Run local Lambda functions](#).
- **Connectors.** Predefined serverless applications that run code on the core. Connectors can provide built-in integration with local infrastructure, device protocols, AWS, and other cloud services. For more information, see [Integrate with services and protocols using connectors](#).
- **Subscriptions.** Defines the publishers, subscribers, and MQTT topics (or subjects) that are authorized for MQTT communication.
- **Resources.** References to local [devices and volumes](#), [machine learning models](#), and [secrets](#), used for access control by Greengrass Lambda functions and connectors.
- **Logs.** Logging configurations for AWS IoT Greengrass system components and Lambda functions. For more information, see [the section called “Monitoring with AWS IoT Greengrass logs”](#).

You manage your Greengrass group in the AWS Cloud and then deploy it to a core. The deployment copies the group configuration to the `group.json` file on the core device. This file is located in `greengrass-root/ggc/deployments/group`.

**Note**

During a deployment, the Greengrass daemon process on the core device stops and then restarts.

Deploying groups from the AWS IoT console

You can deploy a group and manage its deployments from the group's configuration page in the AWS IoT console.

Note

To open this page in the console, choose **Greengrass devices**, then **Groups (V1)**, and then under **Greengrass groups**, choose your group.

To deploy the current version of the group

- From the group configuration page, choose **Deploy**.

To view the deployment history of the group

A group's deployment history includes the date and time, group version, and status of each deployment attempt.

1. From the group configuration page, choose the **Deployments** tab.
2. To see more information about a deployment, including error messages, choose **Deployments** from the AWS IoT console, under **Greengrass devices**.

To redeploy a group deployment

You might want to redeploy a deployment if the current deployment fails or revert to a different group version.

1. From the AWS IoT console, choose **Greengrass devices**, and then choose **Groups (V1)**.
2. Choose the **Deployments** tab.
3. Choose the deployment you want to redeploy and choose **Redeploy**.

To reset group deployments

You might want to reset group deployments to move or delete a group or to remove deployment information. For more information, see [the section called "Reset deployments"](#).

1. From the AWS IoT console, choose **Greengrass devices**, and then choose **Groups (V1)**.
2. Choose the **Deployments** tab.
3. Choose the deployment you want to reset and choose **Reset deployments**.

Deploying groups with the AWS IoT Greengrass API

The AWS IoT Greengrass API provides the following actions to deploy AWS IoT Greengrass groups and manage group deployments. You can call these actions from the AWS CLI, AWS IoT Greengrass API, or AWS SDK.

Action	Description
CreateDeployment	Creates a NewDeployment or Redeployment deployment. You might want to redeploy a deployment if the current deployment fails. Or you might want to redeploy to revert to a different group version.
GetDeploymentStatus	Returns the status of a deployment: Building, InProgress , Success, or Failure. You can configure Amazon EventBridge events to receive deployment notifications. For more information, see the section called "Get deployment notifications".
ListDeployments	Returns the deployment history for the group.
ResetDeployments	Resets the deployments for the group. You might want to reset group deployments to move or delete a group or to remove deployment information. For more information, see the section called "Reset deployments".

Note

For information about bulk deployment operations, see [the section called "Create bulk deployments"](#).

Getting the group ID

The group ID is commonly used in API actions. You can use the [ListGroups](#) action to find the ID of the target group from your list of groups. For example, in the AWS CLI, use the `list-groups` command.

```
aws greengrass list-groups
```

You can also include the `query` option to filter results. For example:

- To get the most recently created group:

```
aws greengrass list-groups --query "reverse(sort_by(Groups, &CreationTimestamp))[0]"
```

- To get a group by name:

```
aws greengrass list-groups --query "Groups[?Name=='MyGroup']"
```

Group names are not required to be unique, so multiple groups might be returned.

The following is an example `list-groups` response. The information for each group includes the ID of the group (in the `Id` property) and the ID of the most recent group version (in the `LatestVersion` property). To get other version IDs for a group, use the group ID with [ListGroupVersions](#).

Note

You can also find these values in the AWS IoT console. The group ID is displayed on the group's **Settings** page. Group version IDs are displayed on the group's **Deployments** tab.

```
{
  "Groups": [
    {
      "LatestVersionArn": "arn:aws:us-west-2:123456789012:/greengrass/
groups/00dedaaa-ac16-484d-ad77-c3eedEXAMPLE/versions/4cbc3f07-fc5e-48c4-
a50e-7d356EXAMPLE",
      "Name": "MyFirstGroup",
```

```

        "LastUpdatedTimestamp": "2019-11-11T05:47:31.435Z",
        "LatestVersion": "4cbc3f07-fc5e-48c4-a50e-7d356EXAMPLE",
        "CreationTimestamp": "2019-11-11T05:47:31.435Z",
        "Id": "00dedaaa-ac16-484d-ad77-c3eedEXAMPLE",
        "Arn": "arn:aws:us-west-2:123456789012:/greengrass/groups/00dedaaa-
ac16-484d-ad77-c3eedEXAMPLE"
    },
    {
        "LatestVersionArn": "arn:aws:us-west-2:123456789012:/greengrass/
groups/036ceaf9-9319-4716-ba2a-237f9EXAMPLE/versions/8fe9e8ec-64d1-4647-
b0b0-01dc8EXAMPLE",
        "Name": "GreenhouseSensors",
        "LastUpdatedTimestamp": "2020-01-07T19:58:36.774Z",
        "LatestVersion": "8fe9e8ec-64d1-4647-b0b0-01dc8EXAMPLE",
        "CreationTimestamp": "2020-01-07T19:58:36.774Z",
        "Id": "036ceaf9-9319-4716-ba2a-237f9EXAMPLE",
        "Arn": "arn:aws:us-west-2:123456789012:/greengrass/
groups/036ceaf9-9319-4716-ba2a-237f9EXAMPLE"
    },
    ...
]
}

```

If you don't specify an AWS Region, AWS CLI commands use the default Region from your profile. To return groups in a different Region, include the *region* option. For example:


```
aws greengrass list-groups --region us-east-1
```

Overview of the AWS IoT Greengrass group object model

When programming with the AWS IoT Greengrass API, it's helpful to understand the Greengrass group object model.


Groups

In the AWS IoT Greengrass API, the top-level Group object consists of metadata and a list of GroupVersion objects. GroupVersion objects are associated with a Group by ID.



Group versions

GroupVersion objects define group membership. Each **GroupVersion** references a **CoreDefinitionVersion** and other component versions by ARN. These references determine which entities to include in the group.



For example, to include three Lambda functions, one device, and two subscriptions in the group, the **GroupVersion** references:

- The **CoreDefinitionVersion** that contains the required core.
- The **FunctionDefinitionVersion** that contains the three functions.
- The **DeviceDefinitionVersion** that contains the client device.

- The `SubscriptionDefinitionVersion` that contains the two subscriptions.

The `GroupVersion` deployed to a core device determines the entities that are available in the local environment and how they can interact.

Group components


Components that you add to groups have a three-level hierarchy:

- A *Definition* that references a list of *DefinitionVersion* objects of a given type. For example, a `DeviceDefinition` references a list of `DeviceDefinitionVersion` objects.
- A *DefinitionVersion* that contains a set of entities of a given type. For example, a `DeviceDefinitionVersion` contains a list of `Device` objects.
- Individual entities that define their properties and behavior. For example, a `Device` defines the ARN of the corresponding client device in the AWS IoT registry, the ARN of its device certificate, and whether its local shadow syncs automatically with the cloud.

You can add the following types of entities to a group:

- [Connector](#)
- [Core](#)
- [Device](#)
- [Function](#)
- [Logger](#)
- [Resource](#)
- [Subscription](#)

The following example `DeviceDefinition` references three `DeviceDefinitionVersion` objects that each contain multiple `Device` objects. Only one `DeviceDefinitionVersion` at a time is used in a group.



Updating groups

In the AWS IoT Greengrass API, you use versions to update a group's configuration. Versions are immutable, so to add, remove, or change group components, you must create *DefinitionVersion* objects that contain new or updated entities.

You can associate new *DefinitionVersions* objects with new or existing *Definition* objects. For example, you can use the `CreateFunctionDefinition` action to create a `FunctionDefinition` that includes the `FunctionDefinitionVersion` as an initial version, or you can use the `CreateFunctionDefinitionVersion` action and reference an existing `FunctionDefinition`.

After you create your group components, you create a `GroupVersion` that contains all *DefinitionVersion* objects that you want to include in the group. Then, you deploy the `GroupVersion`.

To deploy a `GroupVersion`, it must reference a `CoreDefinitionVersion` that contains exactly one `Core`. All referenced entities must be members of the group. Also, a [Greengrass service role](#) must be associated with your AWS account in the AWS Region where you are deploying the `GroupVersion`.

Note

The Update actions in the API are used to change the name of a Group or component *Definition* object.

Updating entities that reference AWS resources

Greengrass Lambda functions and [secret resources](#) define Greengrass-specific properties and also reference corresponding AWS resources. To update these entities, you might make changes to the corresponding AWS resource instead of your Greengrass objects. For example, Lambda functions reference a function in AWS Lambda and also define lifecycle and other properties that are specific to the Greengrass group.

- To update Lambda function code or packaged dependencies, make your changes in AWS Lambda. During the next group deployment, these changes are retrieved from AWS Lambda and copied to your local environment.
- To update [Greengrass-specific properties](#), you create a `FunctionDefinitionVersion` that contains the updated Function properties.

Note

Greengrass Lambda functions can reference a Lambda function by alias ARN or version ARN. If you reference the alias ARN (recommended), you don't need to update your `FunctionDefinitionVersion` (or `SubscriptionDefinitionVersion`) when you publish a new function version in AWS Lambda. For more information, see [the section called "Reference functions by alias or version"](#).

See also

- [the section called "Get deployment notifications"](#)
- [the section called "Reset deployments"](#)
- [the section called "Create bulk deployments"](#)
- [Troubleshooting Deployment Issues](#)

- [AWS IoT Greengrass Version 1 API Reference](#)
- [AWS IoT Greengrass commands](#) in the *AWS CLI Command Reference*

Get deployment notifications

Amazon EventBridge event rules provide you with notifications about state changes for your Greengrass group deployments. EventBridge delivers a near real-time stream of system events that describes changes in AWS resources. AWS IoT Greengrass sends these events to EventBridge on an *at least once* basis. This means that AWS IoT Greengrass might send multiple copies of a given event to ensure delivery. Additionally, your event listeners might not receive the events in the order that the events occurred.

Note

Amazon EventBridge is an event bus service that you can use to connect your applications with data from a variety of sources, such as [Greengrass core devices](#) and deployment notifications. For more information, see [What is Amazon EventBridge?](#) in the *Amazon EventBridge User Guide*.

AWS IoT Greengrass emits an event when group deployments change state. You can create an EventBridge rule that runs for all state transitions or transitions to states you specify. When a deployment enters a state that initiates a rule, EventBridge invokes the target actions defined in the rule. This allows you to send notifications, capture event information, take corrective action, or initiate other events in response to a state change. For example, you can create rules for the following use cases:

- Initiate post-deployment operations, such as downloading assets and notifying personnel.
- Send notifications upon a successful or failed deployment.
- Publish custom metrics about deployment events.

AWS IoT Greengrass emits an event when a deployment enters the following states: `Building`, `InProgress`, `Success`, and `Failure`.

Note

Monitoring the status of a [bulk deployment](#) operation is not currently supported. However, AWS IoT Greengrass emits state-change events for individual group deployments that are part of a bulk deployment.

Group deployment status change event

The [event](#) for a deployment state change uses the following format:

```
{
  "version": "0",
  "id": " cd4d811e-ab12-322b-8255-EXAMPLEb1bc8",
  "detail-type": "Greengrass Deployment Status Change",
  "source": "aws.greengrass",
  "account": "123456789012",
  "time": "2018-03-22T00:38:11Z",
  "region": "us-west-2",
  "resources": [],
  "detail": {
    "group-id": "284dcd4e-24bc-4c8c-a770-EXAMPLEf03b8",
    "deployment-id": "4f38f1a7-3dd0-42a1-af48-EXAMPLE09681",
    "deployment-type": "NewDeployment|Redeployment|ResetDeployment|
ForceResetDeployment",
    "status": "Building|InProgress|Success|Failure"
  }
}
```

You can create rules that apply to one or more groups. You can filter rules by one or more of the following deployment types and deployment states:

Deployment types

- **NewDeployment.** The first deployment of a group version.
- **ReDeployment.** A redeployment of a group version.
- **ResetDeployment.** Deletes deployment information stored in the AWS Cloud and on the AWS IoT Greengrass core. For more information, see [the section called “Reset deployments”](#).

- **ForceResetDeployment.** Deletes deployment information stored in the AWS Cloud and reports success without waiting for the core to respond. Also deletes deployment information stored on the core if the core is connected or when it next connects.

Deployment states

- **Building.** AWS IoT Greengrass is validating the group configuration and building deployment artifacts.
- **InProgress.** The deployment is in progress on the AWS IoT Greengrass core.
- **Success.** The deployment was successful.
- **Failure.** The deployment failed.

It's possible that events might be duplicated or out of order. To determine the order of events, use the `time` property.

Note

AWS IoT Greengrass doesn't use the `resources` property, so it's always empty.

Prerequisites for creating EventBridge rules

Before you create an EventBridge rule for AWS IoT Greengrass, do the following:

- Familiarize yourself with events, rules, and targets in EventBridge.
- Create and configure the targets invoked by your EventBridge rules. Rules can invoke many types of targets, including:
 - Amazon Simple Notification Service (Amazon SNS)
 - AWS Lambda functions
 - Amazon Kinesis Video Streams
 - Amazon Simple Queue Service (Amazon SQS) queues

For more information, see [What is Amazon EventBridge?](#) and [Getting started with Amazon EventBridge](#) in the *Amazon EventBridge User Guide*.

Configure deployment notifications (console)

Use the following steps to create an EventBridge rule that publishes an Amazon SNS topic when the deployment state changes for a group. This allows web servers, email addresses, and other topic subscribers to respond to the event. For more information, see [Creating a EventBridge rule that triggers on an event from an AWS resource](#) in the *Amazon EventBridge User Guide*.

1. Open the [Amazon EventBridge console](#).
2. In the navigation pane, choose **Rules**.
3. Choose **Create rule**.
4. Enter a name and description for the rule.

A rule can't have the same name as another rule in the same Region and on the same event bus.

5. For **Event bus**, choose the event bus that you want to associate with this rule. If you want this rule to match events that come from your account, select **AWS default event bus**. When an AWS service in your account emits an event, it always goes to your account's default event bus.
6. For **Rule type**, choose **Rule with an event pattern**.
7. Choose **Next**.
8. For **Event source**, choose **AWS services**.
9. For **Event pattern**, choose **AWS services**.
10. For **AWS service**, choose Greengrass.
11. For **Event type**, choose **Greengrass Deployment Status Change**.

Note

The **AWS API Call via CloudTrail** event type is based on AWS IoT Greengrass integration with AWS CloudTrail. You can use this option to create rules initiated by read or write calls to the AWS IoT Greengrass API. For more information, see [the section called "Logging AWS IoT Greengrass API calls with AWS CloudTrail"](#).

12. Choose the deployment states that initiate a notification.
 - To receive notifications for all state change events, choose **Any state**.
 - To receive notifications for some state change events only, choose **Specific state(s)**, and then choose the target states.

13. Choose the deployment types that initiate a notification.
 - To receive notifications for all deployment types, choose **Any state**.
 - To receive notifications for some deployment types only, choose **Specific state(s)**, and then choose the target deployment types.
14. Choose **Next**.
15. For **Target types**, choose **AWS service**.
16. For **Select a target**, configure your target. This example uses an Amazon SNS topic, but you can configure other target types to send notifications.
 - a. For **Target**, choose **SNS topic**.
 - b. For **Topic**, choose your target topic.
 - c. Choose **Next**.
17. Under **Tags**, define tags for the rule or leave the fields empty.
18. Choose **Next**.
19. Review the details of the rule and choose **Create rule**.

Configure deployment notifications (CLI)

Use the following steps to create an EventBridge rule that publishes an Amazon SNS topic when the deployment state changes for a group. This allows web servers, email addresses, and other topic subscribers to respond to the event.

1. Create the rule.
 - Replace *group-id* with the ID of your AWS IoT Greengrass group.

```
aws events put-rule \  
  --name TestRule \  
  --event-pattern "{\"source\": [\"aws.greengrass\"], \"detail\": {\"group-id\":  
    [\"group-id\"]}}\"
```

Properties that are omitted from the pattern are ignored.

2. Add the topic as a rule target.
 - Replace *topic-arn* with the ARN of your Amazon SNS topic.

```
aws events put-targets \  
  --rule TestRule \  
  --targets "Id"="1", "Arn"="topic-arn"
```

Note

To allow Amazon EventBridge to call your target topic, you must add a resource-based policy to your topic. For more information, see [Amazon SNS permissions](#) in the *Amazon EventBridge User Guide*.

For more information, see [Events and event patterns in EventBridge](#) in the *Amazon EventBridge User Guide*.

Configure deployment notifications (AWS CloudFormation)

Use AWS CloudFormation templates to create EventBridge rules that send notifications about state changes for your Greengrass group deployments. For more information, see [Amazon EventBridge resource type reference](#) in the *AWS CloudFormation User Guide*.

See also

- [Deploy AWS IoT Greengrass groups](#)
- [What is Amazon EventBridge?](#) in the *Amazon EventBridge User Guide*

Reset deployments

This feature is available for AWS IoT Greengrass Core v1.1 and later.

You might want to reset a group's deployments to:

- Delete the group, such as when you want to move the group's core to another group, or the group's core has been reimaged. Before you delete a group, you must reset the group's deployments to use the core with another Greengrass group.
- Move the group's core to a different group.
- Revert the group to its state before any deployments.

- Remove the deployment configuration from the core device.
- Delete sensitive data from the core device or from the cloud.
- Deploy a new group configuration to a core without having to replace the core with another in the current group.

 Note

Reset deployments functionality is not available in AWS IoT Greengrass Core Software v1.0.0. You cannot delete a group that has been deployed using v1.0.0.

The reset deployments operation first cleans up all deployment information stored in the cloud for a given group. It then instructs the group's core device to clean up all of its deployment related information as well (Lambda functions, user logs, shadow database and server certificate, but not the user-defined `config.json` or the Greengrass core certificates). You cannot initiate a reset of deployments for a group if the group currently has a deployment with status of `In Progress` or `Building`.

Reset deployments from the AWS IoT console

You can reset group deployments from group configuration page in the AWS IoT console.

1. In the AWS IoT console navigation pane, under **Manage**, expand **Greengrass devices**, and then choose **Groups (V1)**.
2. Choose the target group.
3. From the **Deployments** tab, choose **Reset deployments**.
4. In the **Reset deployments for this Greengrass Group** dialog box, type **confirm** to agree, and choose **Reset deployment**.

Reset deployments with the AWS IoT Greengrass API

You can use the `ResetDeployments` action in the AWS CLI, AWS IoT Greengrass API, or AWS SDK to reset deployments. The examples in this topic use the CLI.

```
aws greengrass reset-deployments --group-id GroupId [--force]
```

Arguments for the `reset-deployments` CLI command:

`--group-id`

The group ID. Use the `list-groups` command to get this value.

`--force`

Optional. Use this parameter if the group's core device has been lost, stolen, or destroyed. This option causes the reset deployment process to report success after all deployment information in the cloud has been cleaned up, without waiting for a core device to respond. However, if the core device is or becomes active, it also performs cleanup operations.

The output of the `reset-deployments` CLI command looks like this:

```
{
  "DeploymentId": "4db95ef8-9309-4774-95a4-eea580b6ceef",
  "DeploymentArn": "arn:aws:greengrass:us-west-2:106511594199:/greengrass/groups/
b744ed45-a7df-4227-860a-8d4492caa412/deployments/4db95ef8-9309-4774-95a4-eea580b6ceef"
}
```

You can check the status of the reset deployment with the `get-deployment-status` CLI command:

```
aws greengrass get-deployment-status --deployment-id DeploymentId --group-id GroupId
```

Arguments for the `get-deployment-status` CLI command:

`--deployment-id`

The deployment ID.

`--group-id`

The group ID.

The output of the `get-deployment-status` CLI command looks like this:

```
{
  "DeploymentStatus": "Success",
  "UpdatedAt": "2017-04-04T00:00:00.000Z"
}
```

```
}
```

The `DeploymentStatus` is set to `Building` when the reset deployment is being prepared. When the reset deployment is ready but the AWS IoT Greengrass core has not picked up the reset deployment, the `DeploymentStatus` is `InProgress`.

If the reset operation fails, error information is returned in the response.

See also

- [Deploy AWS IoT Greengrass groups](#)
- [ResetDeployments](#) in the *AWS IoT Greengrass Version 1 API Reference*
- [GetDeploymentStatus](#) in the *AWS IoT Greengrass Version 1 API Reference*

Create bulk deployments for groups

You can use simple API calls to deploy large numbers of Greengrass groups at once. These deployments are triggered with an adaptive rate that has a fixed upper limit.

This tutorial describes how to use the AWS CLI to create and monitor a bulk group deployment in AWS IoT Greengrass. The bulk deployment example in this tutorial contains multiple groups. You can use the example in your implementation to add as many groups as you need.

The tutorial contains the following high-level steps:

1. [Create and upload the bulk deployment input file](#)
2. [Create and configure an IAM execution role for bulk deployments](#)
3. [Allow your execution role access to your S3 Bucket](#)
4. [Deploy the groups](#)
5. [Test the deployment](#)

Prerequisites

To complete this tutorial, you need:

- One or more deployable Greengrass groups. For more information about creating AWS IoT Greengrass groups and cores, see [Getting started with AWS IoT Greengrass](#).

- The AWS CLI installed and configured on your machine. For information, see the [AWS CLI User Guide](#).
- An S3 bucket created in the same AWS Region as AWS IoT Greengrass. For information, see [Creating and configuring an S3 bucket](#) in the *Amazon Simple Storage Service User Guide*.

Note

Currently, SSE KMS enabled buckets are not supported.

Step 1: Create and upload the bulk deployment input file

In this step, you create a deployment input file and upload it to your Amazon S3 bucket. This file is a serialized, line-delimited JSON file that contains information about each group in your bulk deployment. AWS IoT Greengrass uses this information to deploy each group on your behalf when you initialize your bulk group deployment.

1. Run the following command to get the groupId for each group you want to deploy. You enter the groupId into your bulk deployment input file so that AWS IoT Greengrass can identify each group to be deployed.

Note

You can also find these values in the AWS IoT console. The group ID is displayed on the group's **Settings** page. Group version IDs are displayed on the group's **Deployments** tab.

```
aws greengrass list-groups
```

The response contains information about each group in your AWS IoT Greengrass account:

```
{
  "Groups": [
    {
```



```
    "Name": "string",
    "Id": "string",
    "Arn": "string",
    "LastUpdatedTimestamp": "string",
    "CreationTimestamp": "string",
    "LatestVersion": "string",
    "LatestVersionArn": "string"
  }
],
"NextToken": "string"
}
```

Run the following command to get the `groupVersionId` of each group you want to deploy.

```
list-group-versions --group-id groupId
```

The response contains information about all of the versions in the group. Make a note of the `Version` value for the group version you want to use.

```
{
  "Versions": [
    {
      "Arn": "string",
      "Id": "string",
      "Version": "string",
      "CreationTimestamp": "string"
    }
  ],
  "NextToken": "string"
}
```

2. In your computer terminal or editor of choice, create a file, *MyBulkDeploymentInputFile*, from the following example. This file contains information about each AWS IoT Greengrass group to be included in a bulk deployment. Although this example defines multiple groups, for this tutorial, your file can contain just one.

Note

The size of this file must be less than 100 MB.

```
{"GroupId": "groupId1", "GroupVersionId": "groupVersionId1",  
  "DeploymentType": "NewDeployment"}  
{"GroupId": "groupId2", "GroupVersionId": "groupVersionId2",  
  "DeploymentType": "NewDeployment"}  
{"GroupId": "groupId3", "GroupVersionId": "groupVersionId3",  
  "DeploymentType": "NewDeployment"}  
...
```

Each record (or line) contains a group object. Each group object contains its corresponding GroupId and GroupVersionId and a DeploymentType. Currently, AWS IoT Greengrass supports NewDeployment bulk deployment types only.

Save and close your file. Make a note of the location of the file.

3. Use the following command in your terminal to upload your input file to your Amazon S3 bucket. Replace the file path with the location and name of your file. For information, see [Add an object to a bucket](#).

```
aws s3 cp path/MyBulkDeploymentInputFile s3://amzn-s3-demo-bucket/
```

Step 2: Create and configure an IAM execution role

In this step, you use the IAM console to create a standalone execution role. You then establish a trust relationship between the role and AWS IoT Greengrass and ensure that your IAM user has PassRole privileges for your execution role. This allows AWS IoT Greengrass to assume your execution role and create the deployments on your behalf.

1. Use the following policy to create an execution role. This policy document allows AWS IoT Greengrass to access your bulk deployment input file when it creates each deployment on your behalf.

For more information about creating an IAM role and delegating permissions, see [Creating IAM roles](#).

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "VisualEditor0",
      "Effect": "Allow",
      "Action": "greengrass:CreateDeployment",
      "Resource": [
        "arn:aws:greengrass:region:accountId:/greengrass/groups/groupId1",
        "arn:aws:greengrass:region:accountId:/greengrass/groups/groupId2",
        "arn:aws:greengrass:region:accountId:/greengrass/groups/groupId3",
        ...
      ]
    }
  ]
}
```

Note

This policy must have a resource for each group or group version in your bulk deployment input file to be deployed by AWS IoT Greengrass. To allow access to all groups, for Resource, specify an asterisk:

```
"Resource": ["*"]
```

2. Modify the trust relationship for your execution role to include AWS IoT Greengrass. This allows AWS IoT Greengrass to use your execution role and the permissions attached to it. For information, see [Editing the trust relationship for an existing role](#).

We recommend that you also include the `aws:SourceArn` and `aws:SourceAccount` global condition context keys in your trust policy to help prevent the *confused deputy* security problem. The condition context keys restrict access to allow only those requests that come from the specified account and Greengrass workspace. For more information about the confused deputy problem, see [Cross-service confused deputy prevention](#).

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "",
      "Effect": "Allow",
      "Principal": {
        "Service": "greengrass.amazonaws.com"
      },
      "Action": "sts:AssumeRole",
      "Condition": {
        "StringEquals": {
          "aws:SourceAccount": "account-id"
        },
        "ArnLike": {
          "aws:SourceArn": "arn:aws:greengrass:region:account-id:"
        }
      }
    }
  ]
}
```

3. Give IAM PassRole permissions for your execution role to your IAM user. This IAM user is the one used to initiate the bulk deployment. PassRole permissions allow your IAM user to pass your execution role to AWS IoT Greengrass for use. For more information, see [Granting a user permissions to pass a role to an AWS service](#).

Use the following example to update the IAM policy attached to your execution role. Modify this example, as necessary.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "Stmt1508193814000",
      "Effect": "Allow",
      "Action": [
        "iam:PassRole"
      ],
      "Resource": [
```

```

        "arn:aws:iam::account-id:user/executionRoleArn"
    ]
    "Condition": {
        "StringEquals": {
            "iam:PassedToService": "greengrass.amazonaws.com"
        }
    }
}

```

Step 3: Allow your execution role access to your S3 Bucket

To start your bulk deployment, your execution role must be able to read your bulk deployment input file from your Amazon S3 bucket. Attach the following example policy to your Amazon S3 bucket so its `GetObject` permissions are accessible to your execution role.

For more information, see [How do I add an S3 bucket policy?](#)

```

{
  "Version": "2008-10-17",
  "Id": "examplePolicy",
  "Statement": [
    {
      "Sid": "Stmt1535408982966",
      "Effect": "Allow",
      "Principal": {
        "AWS": [
          "executionRoleArn"
        ]
      },
      "Action": "s3:GetObject",
      "Resource":
        "arn:aws:s3:::amzn-s3-demo-bucket/objectKey"
    }
  ]
}

```

You can use the following command in your terminal to check your bucket's policy:

```
aws s3api get-bucket-policy --bucket amzn-s3-demo-bucket
```

Note

You can directly modify your execution role to grant it permission to your Amazon S3 bucket's `GetObject` permissions instead. To do this, attach the following example policy to your execution role.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "VisualEditor0",
      "Effect": "Allow",
      "Action": "s3:GetObject",
      "Resource": "arn:aws:s3:::amzn-s3-demo-bucket/objectKey"
    }
  ]
}
```

Step 4: Deploy the groups

In this step, you start a bulk deployment operation for all group versions configured in your bulk deployment input file. The deployment action for each of your group versions is of type `NewDeploymentType`.

Note

You cannot call **StartBulkDeployment** while another bulk deployment from the same account is still running. The request is rejected.

1. Use the following command to start the bulk deployment.

We recommend that you include an `X-Amzn-Client-Token` token in every **StartBulkDeployment** request. These requests are idempotent with respect to the token and

the request parameters. This token can be any unique, case-sensitive string of up to 64 ASCII characters.

```
aws greengrass start-bulk-deployment --cli-input-json "{
    \"InputFileUri\": \"URI of file in S3 bucket\",
    \"ExecutionRoleArn\": \"ARN of execution role\",
    \"AmznClientToken\": \"your Amazon client token\"
}"
```

The command should result in a successful status code of 200, along with the following response:

```
{
  \"bulkDeploymentId\": UUID
}
```

Make a note of the bulk deployment ID. It can be used to check the status of your bulk deployment.

 Note

Although bulk deployment operations are not currently supported, you can create Amazon EventBridge event rules to get notifications about deployment status changes for individual groups. For more information, see [the section called “Get deployment notifications”](#).

2. Use the following command to check the status of your bulk deployment.

```
aws greengrass get-bulk-deployment-status --bulk-deployment-id 1234567
```

The command should return a successful status code of 200 in addition to a JSON payload of information:

```
{
  \"BulkDeploymentStatus\": Running,
```

```
"Statistics": {
  "RecordsProcessed": integer,
  "InvalidInputRecords": integer,
  "RetryAttempts": integer
},
"CreatedAt": "string",
"ErrorMessage": "string",
"ErrorDetails": [
  {
    "DetailedErrorCode": "string",
    "DetailedErrorMessage": "string"
  }
]
```

BulkDeploymentStatus contains the current status of the bulk execution. The execution can have one of six different statuses:

- **Initializing.** The bulk deployment request has been received, and the execution is preparing to start.
- **Running.** The bulk deployment execution has started.
- **Completed.** The bulk deployment execution has finished processing all records.
- **Stopping.** The bulk deployment execution has received a command to stop and will terminate shortly. You can't start a new bulk deployment while a previous deployment is in the Stopping state.
- **Stopped.** The bulk deployment execution has been manually stopped.
- **Failed.** The bulk deployment execution has encountered an error and terminated. You can find error details in the ErrorDetails field.

The JSON payload also includes statistical information about the progress of the bulk deployment. You can use this information to determine how many groups have been processed and how many have failed. The statistical information includes:

- **RecordsProcessed:** The number of group records that were attempted.
- **InvalidInputRecords:** The total number of records that returned a non-retryable error. For example, this can occur if a group record from the input file uses an invalid format or

specifies a nonexistent group version, or if the execution doesn't grant permission to deploy a group or group version.

- **RetryAttempts:** The number of deployment attempts that returned a retryable error. For example, a retry is triggered if the attempt to deploy a group returns a throttling error. A group deployment can be retried up to five times.

In the case of a bulk deployment execution failure, this payload also includes an **ErrorDetails** section that can be used for troubleshooting. It contains information about the cause of the execution failure.

You can periodically check the status of the bulk deployment to confirm that it is progressing as expected. After the deployment is complete, **RecordsProcessed** should be equal to the number of deployment groups in your bulk deployment input file. This indicates that each record has been processed.

Step 5: Test the deployment

Use the **ListBulkDeployments** command to find the ID of your bulk deployment.

```
aws greengrass list-bulk-deployments
```

This command returns a list of all of your bulk deployments from most to least recent, including your **BulkDeploymentId**.

```
{
  "BulkDeployments": [
    {
      "BulkDeploymentId": 1234567,
      "BulkDeploymentArn": "string",
      "CreatedAt": "string"
    }
  ],
  "NextToken": "string"
}
```

Now call the **ListBulkDeploymentDetailedReports** command to gather detailed information about each deployment.

```
aws greengrass list-bulk-deployment-detailed-reports --bulk-deployment-id 1234567
```

The command should return a successful status code of 200 along with a JSON payload of information:

```
{
  "BulkDeploymentResults": [
    {
      "DeploymentId": "string",
      "GroupVersionedArn": "string",
      "CreatedAt": "string",
      "DeploymentStatus": "string",
      "ErrorMessage": "string",
      "ErrorDetails": [
        {
          "DetailedErrorCode": "string",
          "DetailedErrorMessage": "string"
        }
      ]
    }
  ],
  "NextToken": "string"
}
```

This payload usually contains a paginated list of each deployment and its deployment status from most to least recent. It also contains more information in the event of a bulk deployment execution failure. Again, the total number of deployments listed should be equal to the number of groups you identified in your bulk deployment input file.

The information returned can change until the deployments are in a terminal state (success or failure). You can call this command periodically until then.

Troubleshooting bulk deployments

If the bulk deployment is not successful, you can try the following troubleshooting steps. Run the commands in your terminal.

Troubleshoot input file errors

The bulk deployment can fail in the event of syntax errors in the bulk deployment input file. This returns a bulk deployment status of `Failed` with an error message indicating the line number of the first validation error. There are four possible errors:

- InvalidInputFile: Missing `GroupId` at line number: `line number`

This error indicates that the given input file line is unable to register the specified parameter. The possible missing parameters are the `GroupId` and the `GroupVersionId`.

- InvalidInputFile: Invalid deployment type at line number : `line number`. Only valid type is 'NewDeployment'.

This error indicates that the given input file line lists an invalid deployment type. At this time, the only supported deployment type is a `NewDeployment`.

- Line `%s` is too long in S3 File. Valid line is less than 256 chars.

This error indicates that the given input file line is too long and must be shortened.

- Failed to parse input file at line number: `line number`

This error indicates that the given input file line is not considered valid json.

Check for concurrent bulk deployments

You cannot start a new bulk deployment while another one is still running or in a non-terminal state. This can result in a `Concurrent Deployment Error`. You can use the

ListBulkDeployments command to verify that a bulk deployment is not currently running. This command lists your bulk deployments from most to least recent.

```
{
  "BulkDeployments": [
    {
      "BulkDeploymentId": BulkDeploymentId,
      "BulkDeploymentArn": "string",
      "CreatedAt": "string"
    }
  ],
  "NextToken": "string"
}
```

Use the `BulkDeploymentId` of the first listed bulk deployment to run the **GetBulkDeploymentStatus** command. If your most recent bulk deployment is in a running state (Initializing or Running), use the following command to stop the bulk deployment.

```
aws greengrass stop-bulk-deployment --bulk-deployment-id BulkDeploymentId
```

This action results in a status of Stopping until the deployment is Stopped. After the deployment has reached a Stopped status, you can start a new bulk deployment.

Check ErrorDetails

Run the `GetBulkDeploymentStatus` command to return a JSON payload that contains information about any bulk deployment execution failure.

```
"Message": "string",
"ErrorDetails": [
  {
    "DetailedErrorCode": "string",
    "DetailedErrorMessage": "string"
  }
]
```

When exiting with an error, the `ErrorDetails` JSON payload that is returned by this call contains more information about the bulk deployment execution failure. An error status code in the 400 series, for example, indicates an input error, either in the input parameters or the caller dependencies.

Check the AWS IoT Greengrass core log

You can troubleshoot issues by viewing the AWS IoT Greengrass core logs. Use the following commands to view `runtime.log`:

```
cd /greengrass/ggc/var/log
sudo cat system/runtime.log | more
```

For more information about AWS IoT Greengrass logging, see [Monitoring with AWS IoT Greengrass logs](#).

See also

For more information, see the following resources:

- [Deploy AWS IoT Greengrass groups](#)
- [Amazon S3 API commands](#) in the *AWS CLI Command Reference*
- [AWS IoT Greengrass commands](#) in the *AWS CLI Command Reference*

Run Lambda functions on the AWS IoT Greengrass core

AWS IoT Greengrass provides a containerized Lambda runtime environment for user-defined code that you author in AWS Lambda. Lambda functions that are deployed to an AWS IoT Greengrass core run in the core's local Lambda runtime. Local Lambda functions can be triggered by local events, messages from the cloud, and other sources, which brings local compute functionality to client devices. For example, you can use Greengrass Lambda functions to filter device data before transmitting the data to the cloud.

To deploy a Lambda function to a core, you add the function to a Greengrass group (by referencing the existing Lambda function), configure group-specific settings for the function, and then deploy the group. If the function accesses AWS services, you also must add any required permissions to the [Greengrass group role](#).

You can configure parameters that determine how the Lambda functions run, including permissions, isolation, memory limits, and more. For more information, see [the section called "Controlling Greengrass Lambda function execution"](#).

Note

These settings also make it possible to run AWS IoT Greengrass in a Docker container. For more information, see [the section called "Run AWS IoT Greengrass in a Docker container"](#).

The following table lists supported [AWS Lambda runtimes](#) and the versions of AWS IoT Greengrass Core software that they can run on.

Language or platform	GGC version
Python 3.8	1.11
Python 3.7	1.9 or later
Python 2.7 *	1.0 or later
Java 8	1.1 or later
Node.js 12.x *	1.10 or later

Language or platform	GGC version
Node.js 8.10 *	1.9 or later
Node.js 6.10 *	1.1 or later
C, C++	1.6 or later

* You can run Lambda functions that use these runtimes on supported versions of AWS IoT Greengrass, but you can't create them in AWS Lambda. If the runtime on your device is different from the AWS Lambda runtime specified for that function, you are able to choose your own runtime by using `FunctionRuntimeOverride` in `FunctionDefinitionVersion`. For more information, see [CreateFunctionDefinition](#). For more information about supported runtimes, see [Runtime support policy](#) in the *AWS Lambda Developer Guide*.

SDKs for Greengrass Lambda functions

AWS provides three SDKs that can be used by Greengrass Lambda functions running on an AWS IoT Greengrass core. These SDKs are contained in different packages, so functions can use them simultaneously. To use an SDK in a Greengrass Lambda function, include it in the Lambda function deployment package that you upload to AWS Lambda.

AWS IoT Greengrass Core SDK

Enables local Lambda functions to interact with the core to:

- Exchange MQTT messages with AWS IoT Core.
- Exchange MQTT messages with connectors, client devices, and other Lambda functions in the Greengrass group.
- Interact with the local shadow service.
- Invoke other local Lambda functions.
- Access [secret resources](#).
- Interact with [stream manager](#).

AWS IoT Greengrass provides the AWS IoT Greengrass Core SDK in the following languages and platforms on GitHub.

- [AWS IoT Greengrass Core SDK for Java](#)

- [AWS IoT Greengrass Core SDK for Node.js](#)
- [AWS IoT Greengrass Core SDK for Python](#)
- [AWS IoT Greengrass Core SDK for C](#)

To include the AWS IoT Greengrass Core SDK dependency in the Lambda function deployment package:

1. Download the language or platform of the AWS IoT Greengrass Core SDK package that matches the runtime of your Lambda function.
2. Unzip the downloaded package to get the SDK. The SDK is the `greengrasssdk` folder.
3. Include `greengrasssdk` in the Lambda function deployment package that contains your function code. This is the package you upload to AWS Lambda when you create the Lambda function.

StreamManagerClient

Only the following AWS IoT Greengrass Core SDKs can be used for [stream manager](#) operations:

- Java SDK (v1.4.0 or later)
- Python SDK (v1.5.0 or later)
- Node.js SDK (v1.6.0 or later)

To use the AWS IoT Greengrass Core SDK for Python to interact with stream manager, you must install Python 3.7 or later. You must also install dependencies to include in your Python Lambda function deployment packages:

1. Navigate to the SDK directory that contains the `requirements.txt` file. This file lists the dependencies.
2. Install the SDK dependencies. For example, run the following `pip` command to install them in the current directory:

```
pip install --target . -r requirements.txt
```

Install the AWS IoT Greengrass Core SDK for Python on the core device

If you're running Python Lambda functions, you can also use [pip](#) to install the AWS IoT Greengrass Core SDK for Python on the core device. Then you can deploy your functions without including the SDK in the Lambda function deployment package. For more information, see [greengrasssdk](#).

This support is intended for cores with size constraints. We recommend that you include the SDK in your Lambda function deployment packages when possible.

AWS IoT Greengrass Machine Learning SDK

Enables local Lambda functions to consume machine learning (ML) models that are deployed to the Greengrass core as ML resources. Lambda functions can use the SDK to invoke and interact with a local inference service that's deployed to the core as a connector. Lambda functions and ML connectors can also use the SDK to send data to the ML Feedback connector for uploading and publishing. For more information, including code examples that use the SDK, see [the section called “ML Image Classification”](#), [the section called “ML Object Detection”](#), and [the section called “ML Feedback”](#).

The following table lists supported languages or platforms for SDK versions and the versions of AWS IoT Greengrass Core software they can run on.

SDK version	Language or platform	Required GGC version	Changelog
1.1.0	Python 3.7 or 2.7	1.9.3 or later	Added Python 3.7 support and new feedback client.
1.0.0	Python 2.7	1.7 or later	Initial release.

For download information, see [the section called “AWS IoT Greengrass ML SDK software”](#).

AWS SDKs

Enables local Lambda functions to make direct calls to AWS services, such as Amazon S3, DynamoDB, AWS IoT, and AWS IoT Greengrass. To use an AWS SDK in a Greengrass Lambda function, you must include it in your deployment package. When you use the AWS SDK in the same package as the AWS IoT Greengrass Core SDK, make sure that your Lambda functions use the correct namespaces. Greengrass Lambda functions can't communicate with cloud services when the core is offline.

Download the AWS SDKs from the [Getting Started Resource Center](#).

For more information about creating a deployment package, see [the section called "Create and package a Lambda function"](#) in the Getting Started tutorial or [Creating a deployment package](#) in the *AWS Lambda Developer Guide*.

Migrating cloud-based Lambda functions

The AWS IoT Greengrass Core SDK follows the AWS SDK programming model, which makes it easy to port Lambda functions that are developed for the cloud to Lambda functions that run on an AWS IoT Greengrass core.

For example, the following Python Lambda function uses the AWS SDK for Python (Boto3) to publish a message to the topic `some/topic` in the cloud:

```
import boto3

iot_client = boto3.client("iot-data")
response = iot_client.publish(
    topic="some/topic", qos=0, payload="Some payload".encode()
)
```

To port the function for an AWS IoT Greengrass core, in the `import` statement and `client` initialization, change the `boto3` module name to `greengrasssdk`, as shown in the following example:

```
import greengrasssdk

iot_client = greengrasssdk.client("iot-data")
iot_client.publish(topic="some/topic", qos=0, payload="Some payload".encode())
```

Note

The AWS IoT Greengrass Core SDK supports sending MQTT messages with QoS = 0 only. For more information, see [the section called “Message quality of service”](#).

The similarity between programming models also makes it possible for you to develop your Lambda functions in the cloud and then migrate them to AWS IoT Greengrass with minimal effort. [Lambda executables](#) don't run in the cloud, so you can't use the AWS SDK to develop them in the cloud before deployment.

Reference Lambda functions by alias or version

Greengrass groups can reference a Lambda function by alias (recommended) or by version. Using an alias makes it easier to manage code updates because you don't have to change your subscription table or group definition when the function code is updated. Instead, you just point the alias to the new function version. Aliases resolve to version numbers during group deployment. When you use aliases, the resolved version is updated to the version that the alias is pointing to at the time of deployment.

AWS IoT Greengrass doesn't support Lambda aliases for **\$LATEST** versions. **\$LATEST** versions aren't bound to immutable, published function versions and can be changed at any time, which is counter to the AWS IoT Greengrass principle of version immutability.

A common practice for keeping your Greengrass Lambda functions updated with code changes is to use an alias named **PRODUCTION** in your Greengrass group and subscriptions. As you promote new versions of your Lambda function into production, point the alias to the latest stable version and then redeploy the group. You can also use this method to roll back to a previous version.

Controlling execution of Greengrass Lambda functions by using group-specific configuration

AWS IoT Greengrass provides cloud-based management of Greengrass Lambda functions. Although a Lambda function's code and dependencies are managed using AWS Lambda, you can configure how the Lambda function behaves when it runs in a Greengrass group.

Group-specific configuration settings

AWS IoT Greengrass provides the following group-specific configuration settings for Greengrass Lambda functions.

System user and group

The access identity used to run a Lambda function. By default, Lambda functions run as the group's [default access identity](#). Typically, this is the standard AWS IoT Greengrass system accounts (`ggc_user` and `ggc_group`). You can change the setting and choose the user ID and group ID that have the permissions required to run the Lambda function. You can override both UID and GID or just one if you leave the other field blank. This setting gives you more granular control over access to device resources. We recommend that you configure your Greengrass hardware with appropriate resource limits, file permissions, and disk quotas for the users and groups whose permissions are used to run Lambda functions.

This feature is available for AWS IoT Greengrass Core v1.7 and later.

Important

We recommend that you avoid running Lambda functions as root unless absolutely necessary. Running as root increases the following risks:

- The risk of unintended changes, such as accidentally deleting a critical file.
- The risk to your data and device from malicious individuals.
- The risk of container escapes when Docker containers run with `--net=host` and `UID=EUID=0`.

If you do need to run as root, you must update the AWS IoT Greengrass configuration to enable it. For more information, see [the section called “Running a Lambda function as root”](#).

System user ID (number)

The user ID for the user that has the permissions required to run the Lambda function. This setting is only available if you choose to run as **Another user ID/group ID**. You can use the **getent passwd** command on your AWS IoT Greengrass core device to look up the user ID you want to use to run the Lambda function.

If you use the same UID to run processes and the Lambda function on a Greengrass core device, your Greengrass group role can grant the processes temporary credentials. The processes can use the temporary credentials across Greengrass core deployments.

System group ID (number)

The group ID for the group that has the permissions required to run the Lambda function. This setting is only available if you choose to run as **Another user ID/group ID**. You can use the **getent group** command on your AWS IoT Greengrass core device to look up the group ID you want to use to run the Lambda function.

Lambda function containerization

Choose whether the Lambda function runs with the default containerization for the group, or specify the containerization that should always be used for this Lambda function.

A Lambda function's containerization mode determines its level of isolation.

- Containerized Lambda functions run in **Greengrass container** mode. The Lambda function runs in an isolated runtime environment (or namespace) inside the AWS IoT Greengrass container.
- Non-containerized Lambda functions run in **No container** mode. The Lambda functions runs as a regular Linux process without any isolation.

This feature is available for AWS IoT Greengrass Core v1.7 and later.

We recommend that you run Lambda functions in a Greengrass container unless your use case requires them to run without containerization. When your Lambda functions run in a Greengrass container, you can use attached local and device resources and gain the benefits of isolation and increased security. Before you change the containerization, see [the section called "Considerations when choosing Lambda function containerization"](#).

Note

To run without enabling your device kernel namespace and cgroup, all your Lambda functions must run without containerization. You can accomplish this easily by setting the default containerization for the group. For information, see [the section called "Setting default containerization for Lambda functions in a group"](#).

Memory limit

The memory allocation for the function. The default is 16 MB.

Note

The memory limit setting becomes unavailable when you change the Lambda function to run without containerization. Lambda functions that run without containerization have no memory limit. The memory limit setting is discarded when you change the Lambda function or group default containerization setting to run without containerization.

Timeout

The amount of time before the function or request is terminated. The default is 3 seconds.

Pinned

A Lambda function lifecycle can be *on-demand* or *long-lived*. The default is on-demand.

An on-demand Lambda function starts in a new or reused container when invoked. Requests to the function might be processed by any available container. A long-lived—or *pinned*—Lambda function starts automatically after AWS IoT Greengrass starts and keeps running in its own container (or sandbox). All requests to the function are processed by the same container. For more information, see [the section called “Lifecycle configuration”](#).

Read access to /sys directory

Whether the function can access the host's /sys folder. Use this when the function must read device information from /sys. The default is false.

Note

This setting is not available when you run a Lambda function without containerization. The value of this setting is discarded when you change the Lambda function to run without containerization.

Encoding type

The expected encoding type of the input payload for the function, either JSON or binary. The default is JSON.

Support for the binary encoding type is available starting in AWS IoT Greengrass Core Software v1.5.0 and AWS IoT Greengrass Core SDK v1.1.0. Accepting binary input data can be useful for functions that interact with device data, because the restricted hardware capabilities of devices often make it difficult or impossible for them to construct a JSON data type.

Note

[Lambda executables](#) support the binary encoding type only, not JSON.

Process arguments

The command-line arguments are passed to the Lambda function when it runs.

Environment variables

Key-value pairs that can dynamically pass settings to function code and libraries. Local environment variables work the same way as [AWS Lambda function environment variables](#), but are available in the core environment.

Resource access policies

A list of up to 10 [local resources](#), [secret resources](#), and [machine learning resources](#) that the Lambda function is allowed to access, and the corresponding read-only or read-write permission. In the console, these *affiliated* resources are listed on the group configuration page in the **Resources** tab.

The [containerization mode](#) affects how Lambda functions can access local device and volume resources and machine learning resources.

- Non-containerized Lambda functions must access local device and volume resources directly through the file system on the core device.
- To allow non-containerized Lambda functions to access machine learning resources in the Greengrass group, you must set the resource owner and access permissions properties on the machine learning resource. For more information, see [the section called “Access machine learning resources”](#).

For information about using the AWS IoT Greengrass API to set group-specific configuration settings for user-defined Lambda functions, see [CreateFunctionDefinition](#) in the *AWS IoT Greengrass Version 1 API Reference* or [create-function-definition](#) in the *AWS CLI Command Reference*. To deploy Lambda functions to a Greengrass core, create a function definition version that contains your functions, create a group version that references the function definition version and other group components, and then [deploy the group](#).

Running a Lambda function as root

This feature is available for AWS IoT Greengrass Core v1.7 and later.

Before you can run one or more Lambda functions as root, you must first update the AWS IoT Greengrass configuration to enable support. Support for running Lambda functions as root is off by default. The deployment fails if you try to deploy a Lambda function and run it as root (UID and GID of 0) and you haven't updated the AWS IoT Greengrass configuration. An error like the following appears in the runtime log (*greengrass-root*/ggc/var/log/system/runtime.log):

```
lambda(s)
[list of function arns] are configured to run as root while Greengrass is not
configured to run lambdas with root permissions
```

Important

We recommend that you avoid running Lambda functions as root unless absolutely necessary. Running as root increases the following risks:

- The risk of unintended changes, such as accidentally deleting a critical file.
- The risk to your data and device from malicious individuals.
- The risk of container escapes when Docker containers run with `--net=host` and `UID=EUID=0`.

To allow Lambda functions to run as root

1. On your AWS IoT Greengrass device, navigate to the *greengrass-root*/config folder.

Note

By default, *greengrass-root* is the /greengrass directory.

2. Edit the config.json file to add "allowFunctionsToRunAsRoot" : "yes" to the runtime field. For example:

```
{
  "coreThing" : {
    ...
  },
  "runtime" : {
    ...
    "allowFunctionsToRunAsRoot" : "yes"
  },
  ...
}
```

3. Use the following commands to restart AWS IoT Greengrass:

```
cd /greengrass/ggc/core
sudo ./greengrassd restart
```

Now you can set the user ID and group ID (UID/GID) of Lambda functions to 0 to run that Lambda function as root.

You can change the value of "allowFunctionsToRunAsRoot" to "no" and restart AWS IoT Greengrass if you want to disallow Lambda functions to run as root.

Considerations when choosing Lambda function containerization

This feature is available for AWS IoT Greengrass Core v1.7 and later.

By default, Lambda functions run inside an AWS IoT Greengrass container. That container provides isolation between your functions and the host, which offers more security for both the host and the functions in the container.

We recommend that you run Lambda functions in a Greengrass container unless your use case requires them to run without containerization. By running your Lambda functions in a Greengrass container, you have more control over restricting access to resources.

Here are some example use cases for running without containerization:

- You want to run AWS IoT Greengrass on a device that does not support container mode (for example, because you are using a special Linux distribution or have a kernel version that is too old).
- You want to run your Lambda function in another container environment with its own OverlayFS, but encounter OverlayFS conflicts when you run in a Greengrass container.
- You need access to local resources with paths that can't be determined at deployment time or whose paths can change after deployment, such as pluggable devices.
- You have a legacy application that was written as a process and you have encountered issues when running it as a containerized Lambda function.

Containerization differences

Containerization	Notes
Greengrass container	• All AWS IoT Greengrass features are available when you run a Lambda function in a Greengrass container. • Lambda functions that run in a Greengrass container do not have access to the deployed code of other Lambda functions, even if they run with the same group ID. In other words, your Lambda functions run with greater isolation from one another. • Because Lambda functions that run in an AWS IoT Greengrass container have all child processes execute in the same container as the Lambda function, the child processes are terminated when the Lambda function is terminated.

Containerization	Notes
No container	• The following features are not available to non-containerized Lambda functions: • Lambda function memory limits. • Local device and volume resources. You must access these resources on the core device directly instead of accessing them as members of the Greengrass group. • If your non-containerized Lambda function accesses a machine learning resource, you must identify a resource owner and set access permissions on the resource, not on the Lambda function. This requires AWS IoT Greengrass Core software v1.10 or later. For more information, see the section called "Access machine learning resources". • The Lambda function has read-only access to the deployed code of other Lambda functions that are running with the same group ID. • Lambda functions that spawn child processes in a different process session or with an overridden SIGHUP (signal hangup) handler, such as with the nohup utility, are not automatically terminated by AWS IoT Greengrass when the parent Lambda function is terminated.

Note

The default containerization setting for the Greengrass group doesn't apply to [connectors](#).

Changing the containerization for a Lambda function can cause problems when you deploy it. If you had assigned local resources to your Lambda function that are no longer available with your new containerization settings, deployment fails.

- When you change a Lambda function from running in a Greengrass container to running without containerization, memory limits for the function are discarded. You must access the file system directly instead of using attached local resources. You must remove any attached resources before you deploy.
- When you change a Lambda function from running without containerization to running in a container, your Lambda function loses direct access to the file system. You must define a memory limit for each function or accept the default 16 MB. You can configure those settings for each Lambda function before you deploy.

To change containerization settings for a Lambda function

1. In the AWS IoT console navigation pane, under **Manage**, expand **Greengrass devices**, and then choose **Groups (V1)**.
2. Choose the group that contains the Lambda function whose settings you want to change.
3. Choose the **Lambda functions** tab.
4. On the Lambda function that you want to change, choose the ellipsis (...) and then choose **Edit configuration**.
5. Change the containerization settings. If you configure the Lambda function to run in a Greengrass container, you must also set **Memory limit** and **Read access to /sys directory**.
6. Choose **Save** and then **Confirm** to save the changes to your Lambda function.

The changes take effect when the group is deployed.

You can also use the [CreateFunctionDefinition](#) and [CreateFunctionDefinitionVersion](#) in the *AWS IoT Greengrass API Reference*. If you are changing the containerization setting, be sure to update the other parameters too. For example, if you are changing from running a Lambda function in a Greengrass container to running without containerization, be sure to clear the `MemorySize` parameter.

Determine the isolation modes supported by your Greengrass device

You can use the AWS IoT Greengrass dependency checker to determine which isolation modes (Greengrass container/no container) are supported by your Greengrass device.

To run the AWS IoT Greengrass dependency checker

1. Download and run the AWS IoT Greengrass dependency checker from the [GitHub repository](#).

```
wget https://github.com/aws-samples/aws-greengrass-samples/raw/master/greengrass-  
dependency-checker-GGCv1.11.x.zip  
unzip greengrass-dependency-checker-GGCv1.11.x.zip  
cd greengrass-dependency-checker-GGCv1.11.x  
sudo modprobe configs  
sudo ./check_ggc_dependencies | more
```

2. Where more appears, press the **Spacebar** key to display another page of text.

For information about the **modprobe** command, run **man modprobe** in the terminal.

Setting the default access identity for Lambda functions in a group

This feature is available for AWS IoT Greengrass Core v1.8 and later.

For more control over access to device resources, you can configure the default access identity used to run Lambda functions in the group. This setting determines the default permissions given to your Lambda functions when they run on the core device. To override the setting for individual functions in the group, you can use the function's **Run as** property. For more information, see [Run as](#).

This group-level setting is also used for running the underlying AWS IoT Greengrass Core software. This consists of system Lambda functions that manage operations, such as message routing, local shadow sync, and automatic IP address detection.

The default access identity can be configured to run as the standard AWS IoT Greengrass system accounts (ggc_user and ggc_group) or use the permissions of another user or group. We recommend that you configure your Greengrass hardware with appropriate resource limits, file permissions, and disk quotas for any users and groups whose permissions are used to run user-defined or system Lambda functions.

To modify the default access identity for your AWS IoT Greengrass group

1. In the AWS IoT console navigation pane, under **Manage**, expand **Greengrass devices**, and then choose **Groups (V1)**.
2. Choose the group whose settings you want to change.
3. Choose the **Lambda functions** tab and, under the **Default Lambda function runtime environment** section, choose **Edit**.
4. In the **Edit default Lambda function runtime environment** page, under **Default system user and group**, choose **Another user ID/group ID**.

When you choose this option, the **System user ID (number)** and **System group ID (number)** fields are displayed.

5. Enter a user ID, group ID, or both. If you leave a field blank, the respective Greengrass system account (ggc_user or ggc_group) is used.
 - For **System user ID (number)**, enter the user ID for the user who has the permissions you want to use by default to run Lambda functions in the group. You can use the **getent passwd** command on your AWS IoT Greengrass device to look up the user ID.
 - For **System group ID (number)**, enter the group ID for the group that has the permissions you want to use by default to run Lambda functions in the group. You can use the **getent group** command on your AWS IoT Greengrass device to look up the group ID.

Important

Running as the root user increases risks to your data and device. Do not run as root (UID/GID=0) unless your business case requires it. For more information, see [the section called “Running a Lambda function as root”](#).

The changes take effect when the group is deployed.

Setting default containerization for Lambda functions in a group

This feature is available for AWS IoT Greengrass Core v1.7 and later.

The containerization setting for a Greengrass group determines the default containerization for the Lambda functions in the group.

- In **Greengrass container** mode, Lambda functions run in an isolated runtime environment inside the AWS IoT Greengrass container by default.
- In **No container** mode, Lambda functions run as regular Linux processes by default.

You can modify group settings to specify the default containerization for Lambda functions in the group. You can override this setting for one or more Lambda functions in the group if you want the Lambda functions to run with containerization different from the group default. Before you change containerization settings, see [the section called “Considerations when choosing Lambda function containerization”](#).

Important

If you want to change the default containerization for the group, but have one or more functions that use a different containerization, change the settings for the Lambda functions before you change the group setting. If you change the group containerization setting first, the values for the **Memory limit** and **Read access to /sys directory** settings are discarded.

To modify containerization settings for your AWS IoT Greengrass group

1. In the AWS IoT console navigation pane, under **Manage**, expand **Greengrass devices**, and then choose **Groups (V1)**.
2. Choose the group whose settings you want to change.
3. Choose the **Lambda functions** tab.
4. Under **Default Lambda function runtime environment**, choose **Edit**.
5. In the **Edit default Lambda function runtime environment**, page, under **Default Lambda function containerization**, change the containerization setting.
6. Choose **Save**.

The changes take effect when the group is deployed.

Communication flows for Greengrass Lambda functions

Greengrass Lambda functions support several methods of communicating with other members of the AWS IoT Greengrass group, local services, and cloud services (including AWS services).

Communication using MQTT messages

Lambda functions can send and receive MQTT messages using a publish-subscribe pattern that's controlled by subscriptions.

This communication flow allows Lambda functions to exchange messages with the following entities:

- Client devices in the group.
- Connectors in the group.
- Other Lambda functions in the group.
- AWS IoT.
- Local Device Shadow service.

A subscription defines a message source, a message target, and a topic (or subject) that's used to route messages from the source to the target. Messages that are published to a Lambda function are passed to the function's registered handler. Subscriptions enable more security and provide predictable interactions. For more information, see [the section called “Managed subscriptions in the MQTT messaging workflow”](#).

Note

When the core is offline, Greengrass Lambda functions can exchange messages with client devices, connectors, other functions, and local shadows, but messages to AWS IoT are queued. For more information, see [the section called “MQTT message queue”](#).

Other communication flows

- To interact with local device and volume resources and machine learning models on a core device, Greengrass Lambda functions use platform-specific operating system interfaces. For example, you can use the open method in the [os](#) module in Python functions. To allow a function to access a resource, the function must be *affiliated* with the resource and granted read-only or read-write permission. For more information, including AWS IoT Greengrass core version availability, see [Access local resources](#) and [the section called “Accessing machine learning resources from Lambda function code”](#).

Note

If you run your Lambda function without containerization, you cannot use attached local device and volume resources and must access those resources directly.

- Lambda functions can use the Lambda client in the AWS IoT Greengrass Core SDK to invoke other Lambda functions in the Greengrass group.
- Lambda functions can use the AWS SDK to communicate with AWS services. For more information, see [AWS SDK](#).
- Lambda functions can use third-party interfaces to communicate with external cloud services, similar to cloud-based Lambda functions.

Note

Greengrass Lambda functions can't communicate with AWS or other cloud services when the core is offline.

Retrieve the input MQTT topic (or subject)

AWS IoT Greengrass uses subscriptions to control the exchange of MQTT messages between client devices, Lambda functions, and connectors in a group, and with AWS IoT or the local shadow service. Subscriptions define a message source, message target, and an MQTT topic used to route messages. When the target is a Lambda function, the function's handler is invoked when the source publishes a message. For more information, see [the section called "Communication using MQTT messages"](#).

The following example shows how a Lambda function can get the input topic from the context that's passed to the handler. It does this by accessing the `subject` key from the context hierarchy (`context.client_context.custom['subject']`). The example also parses the input JSON message and then publishes the parsed topic and message.

Note

In the AWS IoT Greengrass API, the topic of a [subscription](#) is represented by the `subject` property.

```
import greengrasssdk
import logging

client = greengrasssdk.client('iot-data')

OUTPUT_TOPIC = 'test/topic_results'

def get_input_topic(context):
    try:
        topic = context.client_context.custom['subject']
    except Exception as e:
        logging.error('Topic could not be parsed. ' + repr(e))
    return topic

def get_input_message(event):
    try:
        message = event['test-key']
    except Exception as e:
        logging.error('Message could not be parsed. ' + repr(e))
    return message

def function_handler(event, context):
    try:
        input_topic = get_input_topic(context)
        input_message = get_input_message(event)
        response = 'Invoked on topic "%s" with message "%s"' % (input_topic,
input_message)
        logging.info(response)
    except Exception as e:
        logging.error(e)

    client.publish(topic=OUTPUT_TOPIC, payload=response)

    return
```

To test the function, add it to your group using the default configuration settings. Then, add the following subscriptions and deploy the group. For instructions, see [the section called “Module 3 \(part 1\): Lambda functions on AWS IoT Greengrass”](#).

Topic
filter

test/
input
message

test/
topic
results

After the deployment is completed, invoke the function.

1. In the AWS IoT console, open the **MQTT test client** page.
2. Subscribe to the test/topic_results topic by selecting the **Subscribe to a topic** tab.
3. Publish a message to the test/input_message topic by selecting the **Publish to a topic** tab. For this example, you must include the test-key property in the JSON message.

```
{
  "test-key": "Some string value"
}
```

If successful, the function publishes the input topic and message string to the test/topic_results topic.

Lifecycle configuration for Greengrass Lambda functions

The Greengrass Lambda function lifecycle determines when a function starts and how it creates and uses containers. The lifecycle also determines how variables and preprocessing logic that are outside of the function handler are retained.

AWS IoT Greengrass supports the on-demand (default) or long-lived lifecycles:

- **On-demand** functions start when they are invoked and stop when there are no tasks left to execute. An invocation of the function creates a separate container (or sandbox) to process invocations, unless an existing container is available for reuse. Data that's sent to the function might be pulled by any of the containers.

Multiple invocations of an on-demand function can run in parallel.

Variables and preprocessing logic that are defined outside of the function handler are not retained when new containers are created.

- **Long-lived** (or *pinned*) functions start automatically when the AWS IoT Greengrass core starts and run in a single container. All data that's sent to the function is pulled by the same container.

Multiple invocations are queued until earlier invocations are executed.

Variables and preprocessing logic that are defined outside of the function handler are retained for every invocation of the handler.

Long-lived Lambda functions are useful when you need to start doing work without any initial input. For example, a long-lived function can load and start processing an ML model to be ready when the function starts receiving device data.

Note

Remember that long-lived functions have timeouts that are associated with invocations of their handler. If you want to execute indefinitely running code, you must start it outside the handler. Make sure that there's no blocking code outside the handler that might prevent the function from completing its initialization.

These functions run unless the core stops (for example, during a group deployment or a device reboot) or the function enters an error state (such as a handler timeout, uncaught exception, or when it exceeds its memory limits).

For more information about container reuse, see [Understanding Container Reuse in AWS Lambda](#) in the AWS Compute Blog.

Lambda executables

This feature is available for AWS IoT Greengrass Core v1.6 and later.

A Lambda executable is a type of Greengrass Lambda function that you can use to run binary code in the core environment. It lets you execute device-specific functionality natively and benefit from the smaller footprint of compiled code. Lambda executables can be invoked by events, invoke other functions, and access local resources.

Lambda executables support the binary encoding type only (not JSON), but otherwise you can manage them in your Greengrass group and deploy them like other Greengrass Lambda functions. However, the process of creating Lambda executables is different from creating Python, Java, and Node.js Lambda functions:

- You can't use the AWS Lambda console to create (or manage) a Lambda executable. You can create a Lambda executable only by using the AWS Lambda API.
- You upload the function code to AWS Lambda as a compiled executable that includes the [AWS IoT Greengrass Core SDK for C](#).
- You specify the executable name as the function handler.

Lambda executables must implement certain calls and programming patterns in their function code. For example, the main method must:

- Call `gg_global_init` to initialize Greengrass internal global variables. This function must be called before creating any threads, and before calling any other AWS IoT Greengrass Core SDK functions.
- Call `gg_runtime_start` to register the function handler with the Greengrass Lambda runtime. This function must be called during initialization. Calling this function causes the current thread to be used by the runtime. The optional `GG_RT_OPT_ASYNC` parameter tells this function to not block, but instead to create a new thread for the runtime. This function uses a `SIGTERM` handler.

The following snippet is the main method from the [simple_handler.c](#) code example on GitHub.

```
int main() {
    gg_error err = GGE_SUCCESS;

    err = gg_global_init(0);
    if(err) {
```

```
        gg_log(GG_LOG_ERROR, "gg_global_init failed %d", err);
        goto cleanup;
    }

    gg_runtime_start(handler, 0);

cleanup:
    return -1;
}
```

For more information about requirements, constraints, and other implementation details, see [AWS IoT Greengrass Core SDK for C](#).

Create a Lambda executable

After you compile your code along with the SDK, use the AWS Lambda API to create a Lambda function and upload your compiled executable.

Note

Your function must be compiled with a C89 compatible compiler.

The following example uses the [create-function](#) CLI command to create a Lambda executable. The command specifies:

- The name of the executable for the handler. This must be the exact name of your compiled executable.
- The path to the .zip file that contains the compiled executable.
- `arn:aws:greengrass:::runtime/function/executable` for the runtime. This is the runtime for all Lambda executables.

Note

For `role`, you can specify the ARN of any Lambda execution role. AWS IoT Greengrass doesn't use this role, but the parameter is required to create the function. For more information about Lambda execution roles, see [AWS Lambda permissions model](#) in the *AWS Lambda Developer Guide*.

```
aws lambda create-function \  
--region aws-region \  
--function-name function-name \  
--handler executable-name \  
--role role-arn \  
--zip-file fileb://file-name.zip \  
--runtime arn:aws:greengrass::runtime/function/executable
```

Next, use the AWS Lambda API to publish a version and create an alias.

- Use [publish-version](#) to publish a function version.

```
aws lambda publish-version \  
--function-name function-name \  
--region aws-region
```

- Use [create-alias](#) to create an alias that points to the version you just published. We recommend that you reference Lambda functions by alias when you add them to a Greengrass group.

```
aws lambda create-alias \  
--function-name function-name \  
--name alias-name \  
--function-version version-number \  
--region aws-region
```

Note

The AWS Lambda console doesn't display Lambda executables. To update the function code, you must use the AWS Lambda API.

Then, add the Lambda executable to a Greengrass group, configure it to accept binary input data in its group-specific settings, and deploy the group. You can do this in the AWS IoT Greengrass console or by using the AWS IoT Greengrass API.

Running AWS IoT Greengrass in a Docker container

AWS IoT Greengrass can be configured to run in a [Docker](#) container.

You can download a Dockerfile [through Amazon CloudFront](#) that has the AWS IoT Greengrass Core software and dependencies installed. To modify the Docker image to run on different platform architectures or reduce the size of the Docker image, see the README file in the Docker package download.

To help you get started experimenting with AWS IoT Greengrass, AWS also provides prebuilt Docker images that have the AWS IoT Greengrass Core software and dependencies installed. You can download an image from [Docker Hub](#) or [Amazon Elastic Container Registry](#) (Amazon ECR). These prebuilt images use Amazon Linux 2 (x86_64) and Alpine Linux (x86_64, Armv7l, or AArch64) base images.

 Important

On June 30, 2022, AWS IoT Greengrass ended maintenance for AWS IoT Greengrass Core software v1.x Docker images that are published to Amazon Elastic Container Registry (Amazon ECR) and Docker Hub. You can continue to download these Docker images from Amazon ECR and Docker Hub until June 30, 2023, which is 1 year after maintenance ended. However, the AWS IoT Greengrass Core software v1.x Docker images no longer receive security patches or bug fixes after maintenance ended on June 30, 2022. If you run a production workload that depends on these Docker images, we recommend that you build your own Docker images using the Dockerfiles that AWS IoT Greengrass provides. For more information, see [AWS IoT Greengrass Docker software](#).

This topic describes how to download the AWS IoT Greengrass Docker image from Amazon ECR and run it on a Windows, macOS, or Linux (x86_64) platform. The topic contains the following steps:

1. [Get the AWS IoT Greengrass container image from Amazon ECR](#)
2. [Create and configure the Greengrass group and core](#)
3. [Run AWS IoT Greengrass locally](#)
4. [Configure "No container" containerization for the group](#)
5. [Deploy Lambda functions to the Docker container](#)
6. [\(Optional\) Deploy client devices that interact with Greengrass in the Docker container](#)

The following features aren't supported when you run AWS IoT Greengrass in a Docker container:

- [Connectors](#) that run in **Greengrass container** mode. To run a connector in a Docker container, the connector must run in **No container** mode. To find connectors that support **No container** mode, see [the section called “AWS-provided Greengrass connectors”](#). Some of these connectors have an isolation mode parameter that you must set to **No container**.
- [Local device and volume resources](#). Your user-defined Lambda functions that run in the Docker container must access devices and volumes on the core directly.

These features aren't supported when the Lambda runtime environment for the Greengrass group is set to [No container](#), which is required to run AWS IoT Greengrass in a Docker container.

Prerequisites

Before you start this tutorial, you must do the following.

- You must install the following software and versions on your host computer based on the AWS Command Line Interface (AWS CLI) version that you choose.

AWS CLI version 2

- [Docker](#) version 18.09 or later. Earlier versions might also work, but we recommend 18.09 or later.
- AWS CLI version 2.0.0 or later.
 - To install the AWS CLI version 2, see [Installing the AWS CLI version 2](#).
 - To configure the AWS CLI, see [Configuring the AWS CLI](#).

Note

To upgrade to a later AWS CLI version 2 on a Windows computer, you must repeat the [MSI installation](#) process.

AWS CLI version 1

- [Docker](#) version 18.09 or later. Earlier versions might also work, but we recommend 18.09 or later.
- [Python](#) version 3.6 or later.
- [pip](#) version 18.1 or later.
- AWS CLI version 1.17.10 or later
 - To install the AWS CLI version 1, see [Installing the AWS CLI version 1](#).

- To configure the AWS CLI, see [Configuring the AWS CLI](#).
- To upgrade to the latest version of the AWS CLI version 1, run the following command.

```
pip install awscli --upgrade --user
```

 Note

If you use the [MSI installation](#) of the AWS CLI version 1 on Windows, be aware of the following:

- If the AWS CLI version 1 installation fails to install botocore, try using the [Python and pip installation](#).
- To upgrade to a later AWS CLI version 1, you must repeat the MSI installation process.

- To access Amazon Elastic Container Registry (Amazon ECR) resources, you must grant the following permission.
- Amazon ECR requires users to grant the `ecr:GetAuthorizationToken` permission through an AWS Identity and Access Management (IAM) policy before they can authenticate to a registry and push or pull images from an Amazon ECR repository. For more information, see [Amazon ECR Repository Policy Examples](#) and [Accessing One Amazon ECR Repository](#) in the *Amazon Elastic Container Registry User Guide*.

Step 1: Get the AWS IoT Greengrass container image from Amazon ECR

AWS provides Docker images that have the AWS IoT Greengrass Core software installed.

 Warning

Starting with v1.11.6 of the AWS IoT Greengrass Core software, the Greengrass Docker images no longer include Python 2.7, because Python 2.7 reached end-of-life in 2020 and no longer receives security updates. If you choose to update to these Docker images, we recommend that you validate that your applications work with the new Docker images before you deploy the updates to production devices. If you require Python 2.7 for your application that uses a Greengrass Docker image, you can modify the Greengrass Dockerfile to include Python 2.7 for your application.

For steps that show how to pull the latest image from Amazon ECR, choose your operating system:

Pull the container image (Linux)

Run the following commands in your computer terminal.

1. Log in to the AWS IoT Greengrass registry in Amazon ECR.

```
aws ecr get-login-password --region us-west-2 | docker login --username AWS --password-stdin https://216483018798.dkr.ecr.us-west-2.amazonaws.com
```

If successful, the output prints `Login Succeeded`.

2. Retrieve the AWS IoT Greengrass container image.

```
docker pull 216483018798.dkr.ecr.us-west-2.amazonaws.com/aws-iot-greengrass:latest
```

Note

The latest image contains the latest stable version of the AWS IoT Greengrass Core software installed on an Amazon Linux 2 base image. You can also pull other images from the repository. To find all available images, check the **Tags** page on [Docker Hub](#) or use the **aws ecr list-images** command. For example:

```
aws ecr list-images --region us-west-2 --registry-id 216483018798 --repository-name aws-iot-greengrass
```

3. Enable symlink and hardlink protection. If you're experimenting with running AWS IoT Greengrass in a container, you can enable the settings for the current boot only.

Note

You might need to use **sudo** to run these commands.

- To enable the settings for the current boot only:

```
echo 1 > /proc/sys/fs/protected_hardlinks
```

```
echo 1 > /proc/sys/fs/protected_symlinks
```

- To enable the settings to persist across restarts:

```
echo '# AWS IoT Greengrass' >> /etc/sysctl.conf
echo 'fs.protected_hardlinks = 1' >> /etc/sysctl.conf
echo 'fs.protected_symlinks = 1' >> /etc/sysctl.conf

sysctl -p
```

4. Enable IPv4 network forwarding, which is required for AWS IoT Greengrass cloud deployment and MQTT communications to work on Linux. In the `/etc/sysctl.conf` file, set `net.ipv4.ip_forward` to 1, and then reload `sysctls`.

```
sudo nano /etc/sysctl.conf
# set this net.ipv4.ip_forward = 1
sudo sysctl -p
```

Note

You can use the editor of your choice instead of nano.

Pull the container image (macOS)

Run the following commands in your computer terminal.

1. Log in to the AWS IoT Greengrass registry in Amazon ECR.

```
aws ecr get-login-password --region us-west-2 | docker login --username AWS --password-stdin https://216483018798.dkr.ecr.us-west-2.amazonaws.com
```

If successful, the output prints `Login Succeeded`.

2. Retrieve the AWS IoT Greengrass container image.

```
docker pull 216483018798.dkr.ecr.us-west-2.amazonaws.com/aws-iot-greengrass:latest
```

Note

The latest image contains the latest stable version of the AWS IoT Greengrass Core software installed on an Amazon Linux 2 base image. You can also pull other images from the repository. To find all available images, check the **Tags** page on [Docker Hub](#) or use the **aws ecr list-images** command. For example:

```
aws ecr list-images --region us-west-2 --registry-id 216483018798 --  
repository-name aws-iot-greengrass
```

Pull the container image (Windows)

Run the following commands in a command prompt. Before you can use Docker commands on Windows, Docker Desktop must be running.

1. Log in to the AWS IoT Greengrass registry in Amazon ECR.

```
aws ecr get-login-password --region us-west-2 | docker login --username AWS --  
password-stdin https://216483018798.dkr.ecr.us-west-2.amazonaws.com
```

If successful, the output prints `Login Succeeded`.

2. Retrieve the AWS IoT Greengrass container image.

```
docker pull 216483018798.dkr.ecr.us-west-2.amazonaws.com/aws-iot-greengrass:latest
```

Note

The latest image contains the latest stable version of the AWS IoT Greengrass Core software installed on an Amazon Linux 2 base image. You can also pull other images from the repository. To find all available images, check the **Tags** page on [Docker Hub](#) or use the **aws ecr list-images** command. For example:

```
aws ecr list-images --region us-west-2 --registry-id 216483018798 --  
repository-name aws-iot-greengrass
```

Step 2: Create and configure the Greengrass group and core

The Docker image has the AWS IoT Greengrass Core software installed, but you must create a Greengrass group and core. This includes downloading certificates and the core configuration file.

- Follow the steps in [the section called “Module 2: Installing the AWS IoT Greengrass Core software”](#). Skip the steps where you download and run the AWS IoT Greengrass Core software. The software and its runtime dependencies are already set up in the Docker image.

Step 3: Run AWS IoT Greengrass locally

After your group is configured, you're ready to configure and start the core. For steps that show how to do this, choose your operating system:

Run Greengrass locally (Linux)

Run the following commands in your computer terminal.

- Create a folder for the device's security resources, and move the certificate and keys into that folder. Run the following commands. Replace *path-to-security-files* with the path to the security resources, and replace *certificateId* with the certificate ID in the file names.

```
mkdir /tmp/certs
mv path-to-security-files/certificateId-certificate.pem.crt /tmp/certs
mv path-to-security-files/certificateId-public.pem.key /tmp/certs
mv path-to-security-files/certificateId-private.pem.key /tmp/certs
mv path-to-security-files/AmazonRootCA1.pem /tmp/certs
```

- Create a folder for the device's configuration, and move the AWS IoT Greengrass Core configuration file to that folder. Run the following commands. Replace *path-to-config-file* with the path to the configuration file.

```
mkdir /tmp/config
mv path-to-config-file/config.json /tmp/config
```

- Start AWS IoT Greengrass and bind-mount the certificates and configuration file in the Docker container.

Replace */tmp* with the path where you decompressed your certificates and configuration file.

```
docker run --rm --init -it --name aws-iot-greengrass \
--entrypoint /greengrass-entrypoint.sh \
-v /tmp/certs:/greengrass/certs \
-v /tmp/config:/greengrass/config \
-p 8883:8883 \
216483018798.dkr.ecr.us-west-2.amazonaws.com/aws-iot-greengrass:latest
```

The output should look like this example:

```
Setting up greengrass daemon
Validating hardlink/softlink protection
Waiting for up to 30s for Daemon to start

Greengrass successfully started with PID: 10
```

Run Greengrass locally (macOS)

Run the following commands in your computer terminal.

1. Create a folder for the device's security resources, and move the certificate and keys into that folder. Run the following commands. Replace *path-to-security-files* with the path to the security resources, and replace *certificateId* with the certificate ID in the file names.

```
mkdir /tmp/certs
mv path-to-security-files/certificateId-certificate.pem.crt /tmp/certs
mv path-to-security-files/certificateId-public.pem.key /tmp/certs
mv path-to-security-files/certificateId-private.pem.key /tmp/certs
mv path-to-security-files/AmazonRootCA1.pem /tmp/certs
```

2. Create a folder for the device's configuration, and move the AWS IoT Greengrass Core configuration file to that folder. Run the following commands. Replace *path-to-config-file* with the path to the configuration file.

```
mkdir /tmp/config
mv path-to-config-file/config.json /tmp/config
```

3. Start AWS IoT Greengrass and bind-mount the certificates and configuration file in the Docker container.

Replace `/tmp` with the path where you decompressed your certificates and configuration file.

```
docker run --rm --init -it --name aws-iot-greengrass \  
--entrypoint /greengrass-entrypoint.sh \  
-v /tmp/certs:/greengrass/certs \  
-v /tmp/config:/greengrass/config \  
-p 8883:8883 \  
216483018798.dkr.ecr.us-west-2.amazonaws.com/aws-iot-greengrass:latest
```

The output should look like this example:

```
Setting up greengrass daemon  
Validating hardlink/softlink protection  
Waiting for up to 30s for Daemon to start  
  
Greengrass successfully started with PID: 10
```

Run Greengrass locally (Windows)

1. Create a folder for the device's security resources, and move the certificate and keys into that folder. Run the following commands in a command prompt. Replace *path-to-security-files* with the path to the security resources, and replace *certificateId* with the certificate ID in the file names.

```
mkdir C:\Users\%USERNAME%\Downloads\certs  
move path-to-security-files\certificateId-certificate.pem.crt C:\Users\%USERNAME%  
\Downloads\certs  
move path-to-security-files\certificateId-public.pem.key C:\Users\%USERNAME%  
\Downloads\certs  
move path-to-security-files\certificateId-private.pem.key C:\Users\%USERNAME%  
\Downloads\certs  
move path-to-security-files\AmazonRootCA1.pem C:\Users\%USERNAME%\Downloads\certs
```

2. Create a folder for the device's configuration, and move the AWS IoT Greengrass Core configuration file to that folder. Run the following commands in a command prompt. Replace *path-to-config-file* with the path to the configuration file.

```
mkdir C:\Users\%USERNAME%\Downloads\config
```



```
move path-to-config-file\config.json C:\Users\%USERNAME%\Downloads\config
```

3. Start AWS IoT Greengrass and bind-mount the certificates and configuration file in the Docker container. Run the following commands in your command prompt.

```
docker run --rm --init -it --name aws-iot-greengrass --entrypoint /greengrass-  
entrypoint.sh -v c:/Users/%USERNAME%/Downloads/certs:/greengrass/certs  
-v c:/Users/%USERNAME%/Downloads/config:/greengrass/config -p 8883:8883  
216483018798.dkr.ecr.us-west-2.amazonaws.com/aws-iot-greengrass:latest
```

When Docker prompts you to share your C:\ drive with the Docker daemon, allow it to bind-mount the C:\ directory inside the Docker container. For more information, see [Shared drives](#) in the Docker documentation.

The output should look like this example:

```
Setting up greengrass daemon  
Validating hardlink/softlink protection  
Waiting for up to 30s for Daemon to start  
  
Greengrass successfully started with PID: 10
```

Note

If the container doesn't open the shell and exits immediately, you can debug the issue by bind-mounting the Greengrass runtime logs when you start the image. For more information, see [the section called "To persist Greengrass runtime logs outside of the Docker container"](#).

Step 4: Configure "No container" containerization for the Greengrass group

When you run AWS IoT Greengrass in a Docker container, all Lambda functions must run without containerization. In this step, you set the default containerization for the group to **No container**. You must do this before you deploy the group for the first time.

1. In the AWS IoT console navigation pane, under **Manage**, expand **Greengrass devices**, and then choose **Groups (V1)**.
2. Choose the group whose settings you want to change.
3. Choose the **Lambda functions** tab.
4. Under **Default Lambda function runtime environment**, choose **Edit**.
5. In the **Edit default Lambda function runtime environment**, under **Default Lambda function containerization**, change the containerization settings.
6. Choose **Save**.

The changes take effect when the group is deployed.

For more information, see [the section called “Setting default containerization for Lambda functions in a group”](#).

 Note

By default, Lambda functions use the group containerization setting. If you override the **No container** setting for any Lambda functions when AWS IoT Greengrass is running in a Docker container, the deployment fails.

Step 5: Deploy Lambda functions to the AWS IoT Greengrass Docker container

You can deploy long-lived Lambda functions to the Greengrass Docker container.

- Follow the steps in [the section called “Module 3 \(part 1\): Lambda functions on AWS IoT Greengrass”](#) to deploy a long-lived Hello World Lambda function to the container.

Step 6: (Optional) Deploy client devices that interact with Greengrass running in the Docker container

You can also deploy client devices that interact with AWS IoT Greengrass when it's running in a Docker container.

- Follow the steps in [the section called “Module 4: Interacting with client devices in an AWS IoT Greengrass group”](#) to deploy client devices that connect to the core and send MQTT messages.

Stopping the AWS IoT Greengrass Docker container

To stop the AWS IoT Greengrass Docker container, press Ctrl+C in your terminal or command prompt. This action sends SIGTERM to the Greengrass daemon process to tear down the Greengrass daemon process and all Lambda processes that were started by the daemon process. The Docker container is initialized with /dev/init process as PID 1, which helps in removing any leftover zombie processes. For more information, see the [Docker run reference](#).

Troubleshooting AWS IoT Greengrass in a Docker container

Use the following information to help troubleshoot issues with running AWS IoT Greengrass in a Docker container.

Error: Cannot perform an interactive login from a non TTY device.

Solution: This error can occur when you run the `aws ecr get-login-password` command. Make sure that you installed the latest AWS CLI version 2 or version 1. We recommend that you use the AWS CLI version 2. For more information, see [Installing the AWS CLI](#) in the *AWS Command Line Interface User Guide*.

Error: Unknown options: -no-include-email.

Solution: This error can occur when you run the `aws ecr get-login` command. Make sure that you have the latest AWS CLI version installed (for example, run: `pip install awscli --upgrade --user`). If you're using Windows and you installed the CLI using the MSI installer, you must repeat the installation process. For more information, see [Installing the AWS Command Line Interface on Microsoft Windows](#) in the *AWS Command Line Interface User Guide*.

Warning: IPv4 is disabled. Networking will not work.

Solution: You might receive this warning or a similar message when running AWS IoT Greengrass on a Linux computer. Enable IPv4 network forwarding as described in this [step](#). AWS IoT Greengrass cloud deployment and MQTT communications don't work when IPv4 forwarding isn't enabled. For more information, see [Configure namespaced kernel parameters \(sysctls\) at runtime](#) in the Docker documentation.

Error: A firewall is blocking file Sharing between windows and the containers.

Solution: You might receive this error or a `Firewall Detected` message when running Docker on a Windows computer. This can also occur if you are signed in on a virtual private network (VPN) and your network settings are preventing the shared drive from being mounted. In that situation, turn off VPN and re-run the Docker container.

Error: An error occurred (AccessDeniedException) when calling the GetAuthorizationToken operation: User: arn:aws:iam::<account-id>:user/<user-name> is not authorized to perform: ecr:GetAuthorizationToken on resource: *

You might receive this error when running the `aws ecr get-login-password` command if you don't have sufficient permissions to access an Amazon ECR repository. For more information, see [Amazon ECR Repository Policy Examples](#) and [Accessing One Amazon ECR Repository](#) in the *Amazon ECR User Guide*.

For general AWS IoT Greengrass troubleshooting help, see [Troubleshooting](#).

Debugging AWS IoT Greengrass in a Docker container

To debug issues with a Docker container, you can persist the Greengrass runtime logs or attach an interactive shell to the Docker container.

To persist Greengrass runtime logs outside of the Docker container

You can run the AWS IoT Greengrass Docker container after bind-mounting the `/greengrass/ggc/var/log` directory. The logs persist even after the container exits or is removed.

On Linux or macOS

[Stop any Greengrass Docker containers](#) running on the host, and then run the following command in a terminal. This bind-mounts the Greengrass log directory and starts the Docker image.

Replace `/tmp` with the path where you decompressed your certificates and configuration file.

```
docker run --rm --init -it --name aws-iot-greengrass \
  --entrypoint /greengrass-entrypoint.sh \
  -v /tmp/certs:/greengrass/certs \
  -v /tmp/config:/greengrass/config \
```

```
-v /tmp/log:/greengrass/ggc/var/log \  
-p 8883:8883 \  
216483018798.dkr.ecr.us-west-2.amazonaws.com/aws-iot-greengrass:latest
```

You can then check your logs at `/tmp/log` on your host to see what happened while Greengrass was running inside the Docker container.

On Windows

[Stop any Greengrass Docker containers](#) running on the host, and then run the following command in a command prompt. This bind-mounts the Greengrass log directory and starts the Docker image.

```
cd C:\Users\%USERNAME%\Downloads  
mkdir log  
docker run --rm --init -it --name aws-iot-greengrass --entrypoint /greengrass-  
entrypoint.sh -v c:/Users/%USERNAME%/Downloads/certs:/greengrass/certs -v c:/  
Users/%USERNAME%/Downloads/config:/greengrass/config -v c:/Users/%USERNAME%/  
Downloads/log:/greengrass/ggc/var/log -p 8883:8883 216483018798.dkr.ecr.us-  
west-2.amazonaws.com/aws-iot-greengrass:latest
```

You can then check your logs at `C:/Users/%USERNAME%/Downloads/log` on your host to see what happened while Greengrass was running inside the Docker container.

To attach an interactive shell to the Docker container

You can attach an interactive shell to a running AWS IoT Greengrass Docker container. This can help you investigate the state of the Greengrass Docker container.

On Linux or macOS

While the Greengrass Docker container is running, run the following command in a separate terminal.

```
docker exec -it $(docker ps -a -q -f "name=aws-iot-greengrass") /bin/bash
```

On Windows

While the Greengrass Docker container is running, run the following commands in a separate command prompt.

```
docker ps -a -q -f "name=aws-iot-greengrass"
```

Replace *gg-container-id* with the container_id result from the previous command.

```
docker exec -it gg-container-id /bin/bash
```

Access local resources with Lambda functions and connectors

This feature is available for AWS IoT Greengrass Core v1.3 and later.

With AWS IoT Greengrass, you can author AWS Lambda functions and configure [connectors](#) in the cloud and deploy them to core devices for local execution. On Greengrass cores running Linux, these locally deployed Lambda functions and connectors can access local resources that are physically present on the Greengrass core device. For example, to communicate with devices that are connected through Modbus or CANbus, you can enable your Lambda function to access the serial port on the core device. To configure secure access to local resources, you must guarantee the security of your physical hardware and your Greengrass core device OS.

To get started accessing local resources, see the following tutorials:

- [How to configure local resource access using the AWS command line interface](#)
- [How to configure local resource access using the AWS Management Console](#)

Supported resource types

You can access two types of local resources: volume resources and device resources.

Volume resources

Files or directories on the root file system (except under `/sys`, `/dev`, or `/var`). These include:

- Folders or files used to read or write information across Greengrass Lambda functions (for example, `/usr/lib/python2.x/site-packages/local`).
- Folders or files under the host's `/proc` file system (for example, `/proc/net` or `/proc/stat`). Supported in v1.6 or later. For additional requirements, see [the section called "Volume resources under the `/proc` directory"](#).

Tip

To configure the `/var`, `/var/run`, and `/var/lib` directories as volume resources, first mount the directory in a different folder and then configure the folder as a volume resource.

When you configure volume resources, you specify a *source* path and a *destination* path. The source path is the absolute path of the resource on the host. The destination path is the absolute path of the resource inside the Lambda namespace environment. This is the container that a Greengrass Lambda function or connector runs in. Any changes to the destination path are reflected in the source path on the host file system.

Note

Files in the destination path are visible in the Lambda namespace only. You can't see them in a regular Linux namespace.

Device resources

Files under `/dev`. Only character devices or block devices under `/dev` are allowed for device resources. These include:

- Serial ports used to communicate with devices connected through serial ports (for example, `/dev/ttyS0`, `/dev/ttyS1`).
- USB used to connect USB peripherals (for example, `/dev/ttyUSB0` or `/dev/bus/usb`).
- GPIOs used for sensors and actuators through GPIO (for example, `/dev/gpiomem`).
- GPUs used to accelerate machine learning using on-board GPUs (for example, `/dev/nvidia0`).
- Cameras used to capture images and videos (for example, `/dev/video0`).

Note

`/dev/shm` is an exception. It can be configured as a volume resource only. Resources under `/dev/shm` must be granted `rw` permission.

AWS IoT Greengrass also supports resource types that are used to perform machine learning inference. For more information, see [Perform machine learning inference](#).

Requirements

The following requirements apply to configuring secure access to local resources:

- You must be using AWS IoT Greengrass Core Software v1.3 or later. To create resources for the host's `/proc` directory, you must be using v1.6 or later.
- The local resource (including any required drivers and libraries) must be correctly installed on the Greengrass core device and consistently available during use.
- The desired operation of the resource, and access to the resource, must not require root privileges.
- Only `read` or `read` and `write` permissions are available. Lambda functions cannot perform privileged operations on the resources.
- You must provide the full path of the local resource on the operating system of the Greengrass core device.
- A resource name or ID has a maximum length of 128 characters and must use the pattern `[a-zA-Z0-9:_-]+`.

Volume resources under the `/proc` directory

The following considerations apply to volume resources that are under the host's `/proc` directory.

- You must be using AWS IoT Greengrass Core Software v1.6 or later.
- You can allow read-only access for Lambda functions, but not read-write access. This level of access is managed by AWS IoT Greengrass.
- You might also need to grant OS group permissions to enable read access in the file system. For example, suppose your source directory or file has a 660 file permission, which means that only the owner or user in the group has read (and write) access. In this case, you must add the OS group owner's permissions to the resource. For more information, see [the section called “Group owner file access permission”](#).
- The host environment and the Lambda namespace both contain a `/proc` directory, so be sure to avoid naming conflicts when you specify the destination path. For example, if `/proc` is the source path, you can specify `/host-proc` as the destination path (or any path name other than `"/proc"`).

Group owner file access permission

An AWS IoT Greengrass Lambda function process normally runs as `ggc_user` and `ggc_group`. However, you can give additional file access permissions to the Lambda function process in the local resource definition, as follows:

- To add the permissions of the Linux group that owns the resource, use the `GroupOwnerSetting#AutoAddGroupOwner` parameter or **Automatically add file system permissions of the system group that owns the resource** console option.
- To add the permissions of a different Linux group, use the `GroupOwnerSetting#GroupOwner` parameter or **Specify another system group to add file system permissions** console option. The `GroupOwner` value is ignored if `GroupOwnerSetting#AutoAddGroupOwner` is true.

An AWS IoT Greengrass Lambda function process inherits all of the file system permissions of `ggc_user`, `ggc_group`, and the Linux group (if added). For the Lambda function to access a resource, the Lambda function process must have the required permissions to the resource. You can use the `chmod(1)` command to change the permission of the resource, if necessary.

See also

- [Service Quotas](#) for resources in the *Amazon Web Services General Reference*

How to configure local resource access using the AWS command line interface

This feature is available for AWS IoT Greengrass Core v1.3 and later.

To use a local resource, you must add a resource definition to the group definition that is deployed to your Greengrass core device. The group definition must also contain a Lambda function definition in which you grant access permissions for local resources to your Lambda functions. For more information, including requirements and constraints, see [Access local resources with Lambda functions and connectors](#).

This tutorial describes the process for creating a local resource and configuring access to it using the AWS Command Line Interface (CLI). To follow the steps in the tutorial, you must have already created a Greengrass group as described in [Getting started with AWS IoT Greengrass](#).

For a tutorial that uses the AWS Management Console, see [How to configure local resource access using the AWS Management Console](#).

Create local resources

First, you use the [CreateResourceDefinition](#) command to create a resource definition that specifies the resources to be accessed. In this example, we create two resources, TestDirectory and TestCamera:

```
aws greengrass create-resource-definition --cli-input-json '{
  "Name": "MyLocalVolumeResource",
  "InitialVersion": {
    "Resources": [
      {
        "Id": "data-volume",
        "Name": "TestDirectory",
        "ResourceDataContainer": {
          "LocalVolumeResourceData": {
            "SourcePath": "/src/LRAtest",
            "DestinationPath": "/dest/LRAtest",
            "GroupOwnerSetting": {
              "AutoAddGroupOwner": true,
              "GroupOwner": ""
            }
          }
        }
      },
      {
        "Id": "data-device",
        "Name": "TestCamera",
        "ResourceDataContainer": {
          "LocalDeviceResourceData": {
            "SourcePath": "/dev/video0",
            "GroupOwnerSetting": {
              "AutoAddGroupOwner": true,
              "GroupOwner": ""
            }
          }
        }
      }
    ]
  }
}
```

Resources: A list of Resource objects in the Greengrass group. One Greengrass group can have up to 50 resources.

Resource#Id: The unique identifier of the resource. The ID is used to refer to a resource in the Lambda function configuration. Max length 128 characters. Pattern: [a-zA-Z0-9:_-]+.

Resource#Name: The name of the resource. The resource name is displayed in the Greengrass console. Max length 128 characters. Pattern: [a-zA-Z0-9:_-]+.

LocalDeviceResourceData#SourcePath: The local absolute path of the device resource. The source path for a device resource can refer only to a character device or block device under /dev.

LocalVolumeResourceData#SourcePath: The local absolute path of the volume resource on the Greengrass core device. This location is outside of the [container](#) that the function runs in. The source path for a volume resource type cannot start with /sys.

LocalVolumeResourceData#DestinationPath: The absolute path of the volume resource inside the Lambda environment. This location is inside the container that the function runs in.

GroupOwnerSetting: Allows you to configure additional group privileges for the Lambda process. This field is optional. For more information, see [Group owner file access permission](#).

GroupOwnerSetting#AutoAddGroupOwner: If true, Greengrass automatically adds the specified Linux OS group owner of the resource to the Lambda process privileges. Thus the Lambda process has the file access permissions of the added Linux group.

GroupOwnerSetting#GroupOwner: Specifies the name of the Linux OS group whose privileges are added to the Lambda process. This field is optional.

A resource definition version ARN is returned by [CreateResourceDefinition](#). The ARN should be used when updating a group definition.

```
{
  "LatestVersionArn": "arn:aws:greengrass:us-west-2:012345678901:/greengrass/
definition/resources/ab14d0b5-116e-4951-a322-9cde24a30373/versions/a4d9b882-
d025-4760-9cfe-9d4fada5390d",
  "Name": "MyLocalVolumeResource",
  "LastUpdatedTimestamp": "2017-11-15T01:18:42.153Z",
  "LatestVersion": "a4d9b882-d025-4760-9cfe-9d4fada5390d",
  "CreationTimestamp": "2017-11-15T01:18:42.153Z",
  "Id": "ab14d0b5-116e-4951-a322-9cde24a30373",
```

```
"Arn": "arn:aws:greengrass:us-west-2:123456789012:/greengrass/definition/resources/
ab14d0b5-116e-4951-a322-9cde24a30373"
}
```

Create the Greengrass function

After the resources are created, use the [CreateFunctionDefinition](#) command to create the Greengrass function and grant the function access to the resource:

```
aws greengrass create-function-definition --cli-input-json '{
  "Name": "MyFunctionDefinition",
  "InitialVersion": {
    "Functions": [
      {
        "Id": "greengrassLraTest",
        "FunctionArn": "arn:aws:lambda:us-
west-2:012345678901:function:lraTest:1",
        "FunctionConfiguration": {
          "Pinned": false,
          "MemorySize": 16384,
          "Timeout": 30,
          "Environment": {
            "ResourceAccessPolicies": [
              {
                "ResourceId": "data-volume",
                "Permission": "rw"
              },
              {
                "ResourceId": "data-device",
                "Permission": "ro"
              }
            ],
            "AccessSysfs": true
          }
        }
      }
    ]
  }
}
```

ResourceAccessPolicies: Contains the `resourceId` and `permission` which grant the Lambda function access to the resource. A Lambda function can access a maximum of 20 resources.

ResourceAccessPolicy#Permission: Specifies which permissions the Lambda function has on the resource. The available options are `rw` (read/write) or `ro` (read-only).

AccessSysfs: If true, the Lambda process can have read access to the `/sys` folder on the Greengrass core device. This is used in cases where the Greengrass Lambda function needs to read device information from `/sys`.

Again, [CreateFunctionDefinition](#) returns a function definition version ARN. The ARN should be used in your group definition version.

```
{
  "LatestVersionArn": "arn:aws:greengrass:us-west-2:012345678901:/greengrass/
definition/functions/3c9b1685-634f-4592-8dfd-7ae1183c28ad/versions/37f0d50e-ef50-4faf-
b125-ade8ed12336e",
  "Name": "MyFunctionDefinition",
  "LastUpdatedTimestamp": "2017-11-22T02:28:02.325Z",
  "LatestVersion": "37f0d50e-ef50-4faf-b125-ade8ed12336e",
  "CreationTimestamp": "2017-11-22T02:28:02.325Z",
  "Id": "3c9b1685-634f-4592-8dfd-7ae1183c28ad",
  "Arn": "arn:aws:greengrass:us-west-2:123456789012:/greengrass/definition/
functions/3c9b1685-634f-4592-8dfd-7ae1183c28ad"
}
```

Add the Lambda function to the group

Finally, use [CreateGroupVersion](#) to add the function to the group. For example:

```
aws greengrass create-group-version --group-id "b36a3aeb-3243-47ff-9fa4-7e8d98cd3cf5" \
--resource-definition-version-arn "arn:aws:greengrass:us-west-2:123456789012:/
greengrass/definition/resources/db6bf40b-29d3-4c4e-9574-21ab7d74316c/versions/31d0010f-
e19a-4c4c-8098-68b79906fb87" \
--core-definition-version-arn "arn:aws:greengrass:us-west-2:123456789012:/
greengrass/definition/cores/adbf3475-f6f3-48e1-84d6-502f02729067/
versions/297c419a-9deb-46dd-8ccc-341fc670138b" \
--function-definition-version-arn "arn:aws:greengrass:us-west-2:123456789012:/
greengrass/definition/functions/d1123830-da38-4c4c-a4b7-e92eec7b6d3e/versions/a2e90400-
caae-4ffd-b23a-db1892a33c78" \
--subscription-definition-version-arn "arn:aws:greengrass:us-west-2:123456789012:/
greengrass/definition/subscriptions/7a8ef3d8-1de3-426c-9554-5b55a32fbc6b/
versions/470c858c-7eb3-4abd-9d48-230236bfbf6a"
```

Note

To learn how to get the group ID to use with this command, see [the section called “Getting the group ID”](#).

A new group version is returned:

```
{
  "Arn": "arn:aws:greengrass:us-west-2:012345678901:/greengrass/groups/
b36a3aeb-3243-47ff-9fa4-7e8d98cd3cf5/versions/291917fb-ec54-4895-823e-27b52da25481",
  "Version": "291917fb-ec54-4895-823e-27b52da25481",
  "CreationTimestamp": "2017-11-22T01:47:22.487Z",
  "Id": "b36a3aeb-3243-47ff-9fa4-7e8d98cd3cf5"
}
```

Your Greengrass group now contains the *lraTest* Lambda function that has access to two resources: TestDirectory and TestCamera.

This example Lambda function, `lraTest.py`, written in Python, writes to the local volume resource:

```
# Demonstrates a simple use case of local resource access.
# This Lambda function writes a file test to a volume mounted inside
# the Lambda environment under destLRatest. Then it reads the file and
# publishes the content to the AWS IoT LRatest topic.

import sys
import greengrasssdk
import platform
import os
import logging

# Setup logging to stdout
logger = logging.getLogger(__name__)
logging.basicConfig(stream=sys.stdout, level=logging.DEBUG)

# Create a Greengrass Core SDK client.
client = greengrasssdk.client('iot-data')
volumePath = '/dest/LRatest'

def function_handler(event, context):
```

```
try:
    client.publish(topic='LRA/test', payload='Sent from AWS IoT Greengrass Core.')
    volumeInfo = os.stat(volumePath)
    client.publish(topic='LRA/test', payload=str(volumeInfo))
    with open(volumePath + '/test', 'a') as output:
        output.write('Successfully write to a file.')
    with open(volumePath + '/test', 'r') as myfile:
        data = myfile.read()
    client.publish(topic='LRA/test', payload=data)
except Exception as e:
    logger.error('Failed to publish message: ' + repr(e))
return
```

These commands are provided by the Greengrass API to create and manage resource definitions and resource definition versions:

- [CreateResourceDefinition](#)
- [CreateResourceDefinitionVersion](#)
- [DeleteResourceDefinition](#)
- [GetResourceDefinition](#)
- [GetResourceDefinitionVersion](#)
- [ListResourceDefinitions](#)
- [ListResourceDefinitionVersions](#)
- [UpdateResourceDefinition](#)

Troubleshooting

- **Q:** Why does my Greengrass group deployment fail with an error similar to:

```
group config is invalid:
  ggc_user or [ggc_group root tty] don't have ro permission on the file: /dev/tty0
```

A: This error indicates that the Lambda process doesn't have permission to access the specified resource. The solution is to change the file permission of the resource so that Lambda can access it. (See [Group owner file access permission](#) for details).

- **Q:** When I configure `/var/run` as a volume resource, why does the Lambda function fail to start with an error message in the `runtime.log`:


```
[ERROR]-container_process.go:39,Runtime execution error: unable to start lambda
  container.
container_linux.go:259: starting container process caused "process_linux.go:345:
container init caused \"rootfs_linux.go:62: mounting \"/var/run\" to rootfs \"/
greengrass/ggc/packages/1.3.0/rootfs_sys\" at \"/greengrass/ggc/packages/1.3.0/
rootfs_sys/run\"
caused \"invalid argument\""
```

A: AWS IoT Greengrass core currently doesn't support the configuration of `/var`, `/var/run`, and `/var/lib` as volume resources. One workaround is to first mount `/var`, `/var/run` or `/var/lib` in a different folder and then configure the folder as a volume resource.

- **Q:** When I configure `/dev/shm` as a volume resource with read-only permission, why does the Lambda function fail to start with an error in the `runtime.log`:

```
[ERROR]-container_process.go:39,Runtime execution error: unable to start lambda
  container.
container_linux.go:259: starting container process caused "process_linux.go:345:
container init caused \"rootfs_linux.go:62: mounting \"/dev/shm\" to rootfs \"/
greengrass/ggc/packages/1.3.0/rootfs_sys\" at \"/greengrass/ggc/packages/1.3.0/
rootfs_sys/dev/shm\"
caused \"operation not permitted\""
```

A: `/dev/shm` can only be configured as read/write. Change the resource permission to `rw` to resolve the issue.

How to configure local resource access using the AWS Management Console

This feature is available for AWS IoT Greengrass Core v1.3 and later.

You can configure Lambda functions to securely access local resources on the host Greengrass core device. *Local resources* refer to buses and peripherals that are physically present on the host, or file system volumes on the host OS. For more information, including requirements and constraints, see [Access local resources with Lambda functions and connectors](#).

This tutorial describes how to use the AWS Management Console to configure access to local resources that are present on an AWS IoT Greengrass core device. It contains the following high-level steps:

1. [Create a Lambda function deployment package](#)
2. [Create and publish a Lambda function](#)
3. [Add the Lambda function to the group](#)
4. [Add a local resource to the group](#)
5. [Add subscriptions to the group](#)
6. [Deploy the group](#)

For a tutorial that uses the AWS Command Line Interface, see [How to configure local resource access using the AWS command line interface](#).

Prerequisites

To complete this tutorial, you need:

- A Greengrass group and a Greengrass core (v1.3 or later). To create a Greengrass group or core, see [Getting started with AWS IoT Greengrass](#).
- The following directories on the Greengrass core device:
 - /src/LRAtest
 - /dest/LRAtest

The owner group of these directories must have read and write access to the directories. You might use the following command to grant access:

```
sudo chmod 0775 /src/LRAtest
```

Step 1: Create a Lambda function deployment package

In this step, you create a Lambda function deployment package, which is a ZIP file that contains the function's code and dependencies. You also download the AWS IoT Greengrass Core SDK to include in the package as a dependency.

1. On your computer, copy the following Python script to a local file named `lraTest.py`. This is the app logic for the Lambda function.

```
# Demonstrates a simple use case of local resource access.
# This Lambda function writes a file test to a volume mounted inside
# the Lambda environment under destLRatest. Then it reads the file and
# publishes the content to the AWS IoT LRatest topic.

import sys
import greengrasssdk
import platform
import os
import logging

# Setup logging to stdout
logger = logging.getLogger(__name__)
logging.basicConfig(stream=sys.stdout, level=logging.DEBUG)

# Create a Greengrass Core SDK client.
client = greengrasssdk.client('iot-data')
volumePath = '/dest/LRatest'

def function_handler(event, context):
    try:
        client.publish(topic='LRA/test', payload='Sent from AWS IoT Greengrass
Core.')
        volumeInfo = os.stat(volumePath)
        client.publish(topic='LRA/test', payload=str(volumeInfo))
        with open(volumePath + '/test', 'a') as output:
            output.write('Successfully write to a file.')
        with open(volumePath + '/test', 'r') as myfile:
            data = myfile.read()
            client.publish(topic='LRA/test', payload=data)
    except Exception as e:
        logger.error('Failed to publish message: ' + repr(e))
    return
```

2. From the [AWS IoT Greengrass Core SDK](#) downloads page, download the AWS IoT Greengrass Core SDK for Python to your computer.
3. Unzip the downloaded package to get the SDK. The SDK is the `greengrasssdk` folder.
4. Zip the following items into a file named `lraTestLambda.zip`:

- `lraTest.py`. App logic.
- `greengrasssdk`. Required library for all Python Lambda functions.

The `lraTestLambda.zip` file is your Lambda function deployment package. Now you're ready to create a Lambda function and upload the deployment package.

Step 2: Create and publish a Lambda function

In this step, you use the AWS Lambda console to create a Lambda function and configure it to use your deployment package. Then, you publish a function version and create an alias.

First, create the Lambda function.

1. In the AWS Management Console, choose **Services**, and open the AWS Lambda console.
2. Choose **Functions**.
3. Choose **Create function** and then choose **Author from scratch**.
4. In the **Basic information** section, use the following values.
 - a. For **Function name**, enter **TestLRA**.
 - b. For **Runtime**, choose **Python 3.7**.
 - c. For **Permissions**, keep the default setting. This creates an execution role that grants basic Lambda permissions. This role isn't used by AWS IoT Greengrass.
5. Choose **Create function**.

Basic information

Function name
Enter a name that describes the purpose of your function.

TestLRA

Use only letters, numbers, hyphens, or underscores with no spaces.

Runtime [Info](#)
Choose the language to use to write your function.


Python 3.7 ▼

Permissions [Info](#)
Lambda will create an execution role with permission to upload logs to Amazon CloudWatch Logs. You can configure and modify permissions further when you add triggers.

► Choose or create an execution role

Cancel **Create function**

6. Upload your Lambda function deployment package and register the handler.
 - a. On the **Code** tab, under **Code source**, choose **Upload from**. From the dropdown, choose **.zip file**.



- b. Choose **Upload**, and then choose your `lraTestLambda.zip` deployment package. Then, choose **Save**.
 - c. On the **Code** tab for the function, under **Runtime settings**, choose **Edit**, and then enter the following values.
 - For **Runtime**, choose **Python 3.7**.
 - For **Handler**, enter `lraTest.function_handler`.
 - d. Choose **Save**.

Note

The **Test** button on the AWS Lambda console doesn't work with this function. The AWS IoT Greengrass Core SDK doesn't contain modules that are required to run your Greengrass Lambda functions independently in the AWS Lambda console. These modules (for example, `greengrass_common`) are supplied to the functions after they are deployed to your Greengrass core.

Next, publish the first version of your Lambda function. Then, create an [alias for the version](#).

Greengrass groups can reference a Lambda function by alias (recommended) or by version. Using an alias makes it easier to manage code updates because you don't have to change your subscription table or group definition when the function code is updated. Instead, you just point the alias to the new function version.

7. From **Actions**, choose **Publish new version**.
8. For **Version description**, enter **First version**, and then choose **Publish**.
9. On the **TestLRA: 1** configuration page, from **Actions**, choose **Create alias**.
10. On the **Create alias** page, for **Name**, enter **test**. For **Version**, enter **1**.

Note

AWS IoT Greengrass doesn't support Lambda aliases for **\$LATEST** versions.

11. Choose **Create**.

An alias is a pointer to one or two versions. Choose each version that you want the alias to point to.

Name*

test

Description

Version*

1

You can shift traffic between two versions, based on weights (%) that you assign. Click [here](#) to learn more.

Additional version

Cancel

Create

You can now add the Lambda function to your Greengrass group.

Step 3: Add the Lambda function to the Greengrass group

In this step, you add the function to your group and configure the function's lifecycle.

First, add the Lambda function to your Greengrass group.

1. In the AWS IoT console navigation pane, under **Manage**, expand **Greengrass devices**, and then choose **Groups (V1)**.
2. Choose the Greengrass group where you want to add the Lambda function.

3. On the group configuration page, choose the **Lambda functions** tab.
4. Under **My Lambda functions** section, choose **Add**.
5. On the **Add Lambda function** page, choose the **Lambda function**. Select **TestLRA**.
6. Choose the **Lambda function version**.
7. In the **Lambda function configuration** section, select **System user and group** and **Lambda function containerization**.

Next, configure the lifecycle of the Lambda function.

8. For **Timeout**, choose **30 seconds**.

 Important

Lambda functions that use local resources (as described in this procedure) must run in a Greengrass container. Otherwise, deployment fails if you try to deploy the function. For more information, see [Containerization](#).

9. At the bottom of the page, choose **Add Lambda function**.

Step 4: Add a local resource to the Greengrass group

In this step, you add a local volume resource to the Greengrass group and grant the function read and write access to the resource. A local resource has a group-level scope. You can grant permissions for any Lambda function in the group to access the resource.

1. On the group configuration page, choose the **Resources** tab.
2. Under the **Local resources** section, choose **Add**.
3. On the **Add a local resource** page, use the following values.
 - a. For **Resource name**, enter **testDirectory**.
 - b. For **Resource type**, choose **Volume**.
 - c. For **Local device path**, enter **/src/LRAtest**. This path must exist on the host OS.

The local device path is the local absolute path of the resource on the file system of the core device. This location is outside of the [container](#) that the function runs in. The path can't start with **/sys**.

- d. For **Destination path**, enter **/dest/LRAtest**. This path must exist on the host OS.

The destination path is the absolute path of the resource in the Lambda namespace. This location is inside the container that the function runs in.

- e. Under **System group owner and file access permission**, select **Automatically add file system permissions of the system group that owns the resource**.

The **System group owner and file access permission** option lets you grant additional file access permissions to the Lambda process. For more information, see [Group owner file access permission](#).

4. Choose **Add resource**. The **Resources** page displays the new testDirectory resource.

Step 5: Add subscriptions to the Greengrass group

In this step, you add two subscriptions to the Greengrass group. These subscriptions enable bidirectional communication between the Lambda function and AWS IoT.

First, create a subscription for the Lambda function to send messages to AWS IoT.

1. On the group configuration page, choose the **Subscriptions** tab.
2. Choose **Add**.
3. On the **Create a subscription** page, configure the source and target, as follows:
 - a. For **Source type**, choose **Lambda function**, and then choose **TestLRA**.
 - b. For **Target type**, choose **Service**, and then choose **IoT Cloud**.
 - c. For **Topic filter**, enter **LRA/test**, and then choose **Create subscription**.
4. The **Subscriptions** page displays the new subscription.

Next, configure a subscription that invokes the function from AWS IoT.

5. On the **Subscriptions** page, choose **Add Subscription**.
6. On the **Select your source and target** page, configure the source and target, as follows:
 - a. For **Source type**, choose **Lambda function**, and then choose **IoT Cloud**.
 - b. For **Target type**, choose **Service**, and then choose **TestLRA**.
 - c. Choose **Next**.

7. On the **Filter your data with a topic** page, for **Topic filter**, enter **invoke/LRAFunction**, and then choose **Next**.
8. Choose **Finish**. The **Subscriptions** page displays both subscriptions.

Step 6: Deploy the AWS IoT Greengrass group

In this step, you deploy the current version of the group definition.

1. Make sure that the AWS IoT Greengrass core is running. Run the following commands in your Raspberry Pi terminal, as needed.
 - a. To check whether the daemon is running:

```
ps aux | grep -E 'greengrass.*daemon'
```

If the output contains a root entry for `/greengrass/ggc/packages/1.11.6/bin/daemon`, then the daemon is running.

Note

The version in the path depends on the AWS IoT Greengrass Core software version that's installed on your core device.

- b. To start the daemon:

```
cd /greengrass/ggc/core/  
sudo ./greengrassd start
```

2. On the group configuration page, choose **Deploy**.

Note

Deployment fails if you run your Lambda function without containerization and try to access attached local resources.

3. If prompted, on the **Lambda function** tab, under **System Lambda functions**, select **IP detector**, and then **Edit**, and then **Automatically detect**.

This enables devices to automatically acquire connectivity information for the core, such as IP address, DNS, and port number. Automatic detection is recommended, but AWS IoT Greengrass also supports manually specified endpoints. You're only prompted for the discovery method the first time that the group is deployed.

 Note

If prompted, grant permission to create the [Greengrass service role](#) and associate it with your AWS account in the current AWS Region. This role allows AWS IoT Greengrass to access your resources in AWS services.

The **Deployments** page shows the deployment timestamp, version ID, and status. When completed, the deployment status is **Completed**.

For troubleshooting help, see [Troubleshooting](#).

Test local resource access

Now you can verify whether the local resource access is configured correctly. To test, you subscribe to the LRA/test topic and publish to the invoke/LRAFunction topic. The test is successful if the Lambda function sends the expected payload to AWS IoT.

1. From the AWS IoT console navigation menu, under **Test**, choose **MQTT test client**.
2. Under **Subscribe to a topic**, for **Topic filter**, enter **LRA/test**.
3. Under **Additional information**, for **MQTT payload display**, select **Display payloads as strings**.
4. Choose **Subscribe**. Your Lambda function publishes to the LRA/test topic.

Subscribe to a topic | **Publish to a topic**

Topic filter [Info](#)
The topic filter describes the topic(s) to which you want to subscribe. The topic filter can include MQTT wildcard characters.

▼ **Additional configuration**

Number of messages to keep
The MQTT test client keeps this many of the most recent messages published to a topic that matches this topic filter.

Quality of service
When subscribing to a topic, quality of service 0 will be chosen by default.

☒ Quality of Service 0 - Message will be delivered at most once

☐ Quality of Service 1 - Message will be delivered at least once

MQTT payload display

☐ Auto-format JSON payloads (improves readability)

☒ Display payloads as strings (more accurate)

☐ Display raw payloads (displays binary data as hexadecimal values)

Subscribe

- Under **Publish to a topic**, in the **Topic name** enter **invoke/LRAFunction**, and then choose **Publish** to invoke your Lambda function. The test is successful if the page displays the function's three message payloads.

Subscribe to a topic

Publish to a topic

Topic name

The topic name identifies the message. The message payload will be published to this topic with a Quality of Service (QoS) of 0.

Q

invoke/LRAFunction

X

Message payload

```
{  
  "message": "Hello from AWS IoT console"  
}
```

► Additional configuration

Publish

Subscriptions

lra/test

▼

X

lra/test

▼

May 03, 2021, 12:09:18 (UTC-0400)

Successfully write to a file.

lra/test

▼

May 03, 2021, 12:09:06 (UTC-0400)

```
posix.stat_result(st_mode=16893, st_ino=171142L, st_dev=45831L, st_nlink=2, st_uid=0, st  
_gid=119, st_size=4096L, st_atime=1620054520, st_mtime=1620058120, st_ctime=1620058120)
```

lra/test

▼

May 03, 2021, 12:09:04 (UTC-0400)

Sent from Greengrass Core.

Pause

Clear

Export

Edit

The test file created by the Lambda function is in the `/src/LRAtest` directory on the Greengrass core device. Although the Lambda function writes to a file in the `/dest/LRAtest` directory, that file is visible in the Lambda namespace only. You can't see it in a regular Linux namespace. Any changes to the destination path are reflected in the source path on the file system.

For troubleshooting help, see [Troubleshooting](#).

Perform machine learning inference

This feature is available for AWS IoT Greengrass Core v1.6 or later.


With AWS IoT Greengrass, you can perform machine learning (ML) inference at the edge on locally generated data using cloud-trained models. You benefit from the low latency and cost savings of running local inference, yet still take advantage of cloud computing power for training models and complex processing.

To get started performing local inference, see [the section called “How to configure machine learning inference”](#).

How AWS IoT Greengrass ML inference works

You can train your inference models anywhere, deploy them locally as *machine learning resources* in a Greengrass group, and then access them from Greengrass Lambda functions. For example, you can build and train deep-learning models in [SageMaker AI](#) and deploy them to your Greengrass core. Then, your Lambda functions can use the local models to perform inference on connected devices and send new training data back to the cloud.

The following diagram shows the AWS IoT Greengrass ML inference workflow.



AWS IoT Greengrass ML inference simplifies each step of the ML workflow, including:

- Building and deploying ML framework prototypes.
- Accessing cloud-trained models and deploying them to Greengrass core devices.
- Creating inference apps that can access hardware accelerators (such as GPUs and FPGAs) as [local resources](#).

Machine learning resources

Machine learning resources represent cloud-trained inference models that are deployed to an AWS IoT Greengrass core. To deploy machine learning resources, first you add the resources to a Greengrass group, and then you define how Lambda functions in the group can access them. During group deployment, AWS IoT Greengrass retrieves the source model packages from the cloud and extracts them to directories inside the Lambda runtime namespace. Then, Greengrass Lambda functions use the locally deployed models to perform inference.

To update a locally deployed model, first update the source model (in the cloud) that corresponds to the machine learning resource, and then deploy the group. During deployment, AWS IoT Greengrass checks the source for changes. If changes are detected, then AWS IoT Greengrass updates the local model.

Supported model sources

AWS IoT Greengrass supports SageMaker AI and Amazon S3 model sources for machine learning resources.

The following requirements apply to model sources:

- S3 buckets that store your SageMaker AI and Amazon S3 model sources must not be encrypted using SSE-C. For buckets that use server-side encryption, AWS IoT Greengrass ML inference currently supports the SSE-S3 or SSE-KMS encryption options only. For more information about server-side encryption options, see [Protecting data using server-side encryption](#) in the *Amazon Simple Storage Service User Guide*.
- The names of S3 buckets that store your SageMaker AI and Amazon S3 model sources must not include periods (.). For more information, see the rule about using virtual hosted-style buckets with SSL in [Rules for bucket naming](#) in the *Amazon Simple Storage Service User Guide*.
- Service-level AWS Region support must be available for both [AWS IoT Greengrass](#) and [SageMaker AI](#). Currently, AWS IoT Greengrass supports SageMaker AI models in the following Regions:

- US East (Ohio)
 - US East (N. Virginia)
 - US West (Oregon)
 - Asia Pacific (Mumbai)
 - Asia Pacific (Seoul)
 - Asia Pacific (Singapore)
 - Asia Pacific (Sydney)
 - Asia Pacific (Tokyo)
 - Europe (Frankfurt)
 - Europe (Ireland)
 - Europe (London)
- AWS IoT Greengrass must have read permission to the model source, as described in the following sections.

SageMaker AI

AWS IoT Greengrass supports models that are saved as SageMaker AI training jobs. SageMaker AI is a fully managed ML service that you can use to build and train models using built-in or custom algorithms. For more information, see [What is SageMaker AI?](#) in the *SageMaker AI Developer Guide*.

If you configured your SageMaker AI environment by [creating a bucket](#) whose name contains `sagemaker`, then AWS IoT Greengrass has sufficient permission to access your SageMaker AI training jobs. The `AWSGreengrassResourceAccessRolePolicy` managed policy allows access to buckets whose name contains the string `sagemaker`. This policy is attached to the [Greengrass service role](#).

Otherwise, you must grant AWS IoT Greengrass read permission to the bucket where your training job is stored. To do this, embed the following inline policy in the service role. You can list multiple bucket ARNs.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
```

```

        "Action": [
            "s3:GetObject"
        ],
        "Resource": [
            "arn:aws:s3:::amzn-s3-demo-bucket"
        ]
    }
]
}

```

Amazon S3

AWS IoT Greengrass supports models that are stored in Amazon S3 as `tar.gz` or `.zip` files.

To enable AWS IoT Greengrass to access models that are stored in Amazon S3 buckets, you must grant AWS IoT Greengrass read permission to access the buckets by doing **one** of the following:

- Store your model in a bucket whose name contains `greengrass`.

The `AWSGreengrassResourceAccessRolePolicy` managed policy allows access to buckets whose name contains the string `greengrass`. This policy is attached to the [Greengrass service role](#).

- Embed an inline policy in the Greengrass service role.

If your bucket name doesn't contain `greengrass`, add the following inline policy to the service role. You can list multiple bucket ARNs.

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "s3:GetObject"
      ],
      "Resource": [
        "arn:aws:s3:::amzn-s3-demo-bucket"
      ]
    }
  ]
}

```



```
}
```

For more information, see [Embedding inline policies](#) in the *IAM User Guide*.

Requirements

The following requirements apply for creating and using machine learning resources:

- You must be using AWS IoT Greengrass Core v1.6 or later.
- User-defined Lambda functions can perform `read` or `read` and `write` operations on the resource. Permissions for other operations are not available. The containerization mode of affiliated Lambda functions determines how you set access permissions. For more information, see [the section called "Access machine learning resources"](#).
- You must provide the full path of the resource on the operating system of the core device.
- A resource name or ID has a maximum length of 128 characters and must use the pattern `[a-zA-Z0-9:_-]+`.

Runtimes and libraries for ML inference

You can use the following ML runtimes and libraries with AWS IoT Greengrass.

- [Amazon SageMaker Neo deep learning runtime](#)
- Apache MXNet
- TensorFlow

These runtimes and libraries can be installed on NVIDIA Jetson TX2, Intel Atom, and Raspberry Pi platforms. For download information, see [the section called "Supported machine learning runtimes and libraries"](#). You can install them directly on your core device.

Be sure to read the following information about compatibility and limitations.

SageMaker AI Neo deep learning runtime

You can use the SageMaker AI Neo deep learning runtime to perform inference with optimized machine learning models on your AWS IoT Greengrass devices. These models are optimized using the SageMaker AI Neo deep learning compiler to improve machine learning inference prediction

speeds. For more information about model optimization in SageMaker AI, see the [SageMaker AI Neo documentation](#).

Note

Currently, you can optimize machine learning models using the Neo deep learning compiler in specific Amazon Web Services Regions only. However, you can use the Neo deep learning runtime with optimized models in each AWS Region where AWS IoT Greengrass core is supported. For information, see [How to Configure Optimized Machine Learning Inference](#).

MXNet versioning

Apache MXNet doesn't currently ensure forward compatibility, so models that you train using later versions of the framework might not work properly in earlier versions of the framework. To avoid conflicts between the model-training and model-serving stages, and to provide a consistent end-to-end experience, use the same MXNet framework version in both stages.

MXNet on Raspberry Pi

Greengrass Lambda functions that access local MXNet models must set the following environment variable:

```
MXNET_ENGINE_TYPE=NativeEngine
```

You can set the environment variable in the function code or add it to the function's group-specific configuration. For an example that adds it as a configuration setting, see this [step](#).

Note

For general use of the MXNet framework, such as running a third-party code example, the environment variable must be configured on the Raspberry Pi.

TensorFlow model-serving limitations on Raspberry Pi

The following recommendations for improving inference results are based on our tests with the TensorFlow 32-bit Arm libraries on the Raspberry Pi platform. These recommendations are intended for advanced users for reference only, without guarantees of any kind.

- Models that are trained using the [Checkpoint](#) format should be "frozen" to the protocol buffer format before serving. For an example, see the [TensorFlow-Slim image classification model library](#).
- Don't use the TF-Estimator and TF-Slim libraries in either training or inference code. Instead, use the .pb file model-loading pattern that's shown in the following example.

```
graph = tf.Graph()
graph_def = tf.GraphDef()
graph_def.ParseFromString(pb_file.read())
with graph.as_default():
    tf.import_graph_def(graph_def)
```

Note

For more information about supported platforms for TensorFlow, see [Installing TensorFlow](#) in the TensorFlow documentation.

Access machine learning resources from Lambda functions

User-defined Lambda functions can access machine learning resources to run local inference on the AWS IoT Greengrass core. A machine learning resource consists of the trained model and other artifacts that are downloaded to the core device.

To allow a Lambda function to access a machine learning resource on the core, you must attach the resource to the Lambda function and define access permissions. The [containerization mode](#) of the affiliated (or *attached*) Lambda function determines how you do this.

Access permissions for machine learning resources

Starting in AWS IoT Greengrass Core v1.10.0, you can define a resource owner for a machine learning resource. The resource owner represents the OS group and permissions that AWS IoT Greengrass uses to download the resource artifacts. If a resource owner is not defined, the downloaded resource artifacts are accessible only to root.

- If non-containerized Lambda functions access a machine learning resource, you must define a resource owner because there's no permission control from the container. Non-containerized Lambda functions can inherit resource owner permissions and use them to access the resource.

- If only containerized Lambda functions access the resource, we recommend that you use function-level permissions instead of defining a resource owner.

Resource owner properties

A resource owner specifies a group owner and group owner permissions.

Group owner. The ID of the group (GID) of an existing Linux OS group on the core device. The group's permissions are added to the Lambda process. Specifically, the GID is added to the supplemental group IDs of the Lambda function.

If a Lambda function in the Greengrass group is configured to [run as](#) the same OS group as the resource owner for a machine learning resource, the resource must be attached to the Lambda function. Otherwise, deployment fails because this configuration gives implicit permissions the Lambda function can use to access the resource without AWS IoT Greengrass authorization. The deployment validation check is skipped if the Lambda function runs as root (UID=0).

We recommend that you use an OS group that's not used by other resources, Lambda functions, or files on the Greengrass core. Using a shared OS group gives attached Lambda functions more access permissions than they need. If you use a shared OS group, an attached Lambda function must also be attached to all machine learning resources that use the shared OS group. Otherwise, deployment fails.

Group owner permissions. The read-only or read and write permission to add to the Lambda process.

Non-containerized Lambda functions must inherit these access permissions to the resource. Containerized Lambda functions can inherit these resource-level permissions or define function-level permissions. If they define function-level permissions, the permissions must be the same or more restrictive than the resource-level permissions.

The following table shows supported access permission configurations.

GGC v1.10 or later

Property	If only containerized Lambda functions access the resource	If any non-containerized Lambda functions access the resource
Function-level properties		
Permissions (read/write)	Required unless the resource defines a resource owner. If a resource owner is defined, function-level permissions must be the same or more restrictive than the resource owner permissions. If only containerized Lambda functions access the resource, we recommend that you don't define a resource owner.	Non-containerized Lambda functions: Not supported. Non-containerized Lambda functions must inherit resource-level permissions. Containerized Lambda functions: Optional, but must be the same or more restrictive than resource-level permissions.
Resource-level properties		
Resource owner	Optional (not recommended).	Required.
Permissions (read/write)	Optional (not recommended).	Required.

GGC v1.9 or earlier

Property	If only containerized Lambda functions access the resource	If any non-containerized Lambda functions access the resource
Function-level properties		
Permissions (read/write)	Required.	Not supported.
Resource-level properties		
Resource owner	Not supported.	Not supported.
Permissions (read/write)	Not supported.	Not supported.

 Note

When you use the AWS IoT Greengrass API to configure Lambda functions and resources, the function-level `ResourceId` property is also required. The `ResourceId` property attaches the machine learning resource to the Lambda function.

Defining access permissions for Lambda functions (console)

In the AWS IoT console, you define access permissions when you configure a machine learning resource or attach one to a Lambda function.

Containerized Lambda functions

If only containerized Lambda functions are attached to the machine learning resource:

- Choose **No system group** as the resource owner for the machine learning resource. This is the recommended setting when only containerized Lambda functions access the machine learning resource. Otherwise, you might give attached Lambda functions more access permissions than they need.

Non-containerized Lambda functions (requires GGC v1.10 or later)

If any non-containerized Lambda functions are attached to the machine learning resource:

- Specify the **System group ID (GID)** to use as the resource owner for the machine learning resource. Choose **Specify system group and permissions** and enter the GID. You can use the `getent group` command on your core device to look up the ID of a system group.
- Choose **Read-only access** or **Read and write access** for the **System group permissions**.

Defining access permissions for Lambda functions (API)

In the AWS IoT Greengrass API, you define permissions to machine learning resources in the `ResourceAccessPolicy` property for the Lambda function or the `OwnerSetting` property for the resource.

Containerized Lambda functions

If only containerized Lambda functions are attached to the machine learning resource:

- For containerized Lambda functions, define access permissions in the `Permission` property of the `ResourceAccessPolicies` property. For example:

```
"Functions": [  
  {  
    "Id": "my-containerized-function",  
    "FunctionArn": "arn:aws:lambda:us-west-2:123456789012:function:function-  
name:alias-or-version",  
    "FunctionConfiguration": {  
      "Environment": {  
        "ResourceAccessPolicies": [  
          {  
            "ResourceId": "my-resource-id",  
            "Permission": "ro-or-rw"  
          }  
        ]  
      },  
      "MemorySize": 512,  
      "Pinned": true,  
      "Timeout": 5  
    }  
  }  
]
```

```
}
]
```

- For machine learning resources, omit the `OwnerSetting` property. For example:

```
"Resources": [
  {
    "Id": "my-resource-id",
    "Name": "my-resource-name",
    "ResourceDataContainer": {
      "S3MachineLearningModelResourceData": {
        "DestinationPath": "/local-destination-path",
        "S3Uri": "s3://uri-to-resource-package"
      }
    }
  }
]
```

This is the recommended configuration when only containerized Lambda functions access the machine learning resource. Otherwise, you might give attached Lambda functions more access permissions than they need.

Non-containerized Lambda functions (requires GGC v1.10 or later)

If any non-containerized Lambda functions are attached to the machine learning resource:

- For non-containerized Lambda functions, omit the `Permission` property in `ResourceAccessPolicies`. This configuration is required and allows the function to inherit the resource-level permission. For example:

```
"Functions": [
  {
    "Id": "my-non-containerized-function",
    "FunctionArn": "arn:aws:lambda:us-west-2:123456789012:function:function-  
name:alias-or-version",
    "FunctionConfiguration": {
      "Environment": {
        "Execution": {
          "IsolationMode": "NoContainer",
        },
      },
      "ResourceAccessPolicies": [
```



```

        {
            "ResourceId": "my-resource-id"
        }
    ],
    },
    "Pinned": true,
    "Timeout": 5
}
}
]

```

- For containerized Lambda functions that also access the machine learning resource, omit the `Permission` property in `ResourceAccessPolicies` or define a permission that is the same or more restrictive as the resource-level permission. For example:

```

"Functions": [
    {
        "Id": "my-containerized-function",
        "FunctionArn": "arn:aws:lambda:us-west-2:123456789012:function:function-name:alias-or-version",
        "FunctionConfiguration": {
            "Environment": {
                "ResourceAccessPolicies": [
                    {
                        "ResourceId": "my-resource-id",
                        "Permission": "ro-or-rw" // Optional, but cannot exceed
the GroupPermission defined for the resource.
                    }
                ]
            },
            "MemorySize": 512,
            "Pinned": true,
            "Timeout": 5
        }
    }
]

```

- For machine learning resources, define the `OwnerSetting` property, including the child `GroupOwner` and `GroupPermission` properties. For example:

```

"Resources": [
    {
        "Id": "my-resource-id",

```

```

    "Name": "my-resource-name",
    "ResourceDataContainer": {
      "S3MachineLearningModelResourceData": {
        "DestinationPath": "/local-destination-path",
        "S3Uri": "s3://uri-to-resource-package",
        "OwnerSetting": {
          "GroupOwner": "os-group-id",
          "GroupPermission": "ro-or-rw"
        }
      }
    }
  }
}
]

```

Accessing machine learning resources from Lambda function code

User-defined Lambda functions use platform-specific OS interfaces to access machine learning resources on a core device.

GGC v1.10 or later

For containerized Lambda functions, the resource is mounted inside the Greengrass container and available at the local destination path defined for the resource. For non-containerized Lambda functions, the resource is symlinked to a Lambda-specific working directory and passed to the `AWS_GG_RESOURCE_PREFIX` environment variable in the Lambda process.

To get the path to the downloaded artifacts of a machine learning resource, Lambda functions append the `AWS_GG_RESOURCE_PREFIX` environment variable to the local destination path defined for the resource. For containerized Lambda functions, the returned value is a single forward slash (/).

```

resourcePath = os.getenv("AWS_GG_RESOURCE_PREFIX") + "/destination-path"
with open(resourcePath, 'r') as f:
    # load_model(f)

```

GGC v1.9 or earlier

The downloaded artifacts of a machine learning resource are located in the local destination path defined for the resource. Only containerized Lambda functions can access machine learning resources in AWS IoT Greengrass Core v1.9 and earlier.

```
resourcePath = "/local-destination-path"
with open(resourcePath, 'r') as f:
    # load_model(f)
```

Your model loading implementation depends on your ML library.

Troubleshooting

Use the following information to help troubleshoot issues with accessing machine learning resources.

Topics

- [InvalidMLModelOwner - GroupOwnerSetting is provided in ML model resource, but GroupOwner or GroupPermission is not present](#)
- [NoContainer function cannot configure permission when attaching Machine Learning resources. <function-arn> refers to Machine Learning resource <resource-id> with permission <ro/rw> in resource access policy.](#)
- [Function <function-arn> refers to Machine Learning resource <resource-id> with missing permission in both ResourceAccessPolicy and resource OwnerSetting.](#)
- [Function <function-arn> refers to Machine Learning resource <resource-id> with permission \"rw \", while resource owner setting GroupPermission only allows \"ro\".](#)
- [NoContainer Function <function-arn> refers to resources of nested destination path.](#)
- [Lambda <function-arn> gains access to resource <resource-id> by sharing the same group owner id](#)

InvalidMLModelOwner - GroupOwnerSetting is provided in ML model resource, but GroupOwner or GroupPermission is not present

Solution: You receive this error if a machine learning resource contains the [ResourceDownloadOwnerSetting](#) object but the required GroupOwner or GroupPermission property isn't defined. To resolve this issue, define the missing property.

NoContainer function cannot configure permission when attaching Machine Learning resources. <function-arn> refers to Machine Learning resource <resource-id> with permission <ro/rw> in resource access policy.

Solution: You receive this error if a non-containerized Lambda function specifies function-level permissions to a machine learning resource. Non-containerized functions must inherit permissions from the resource owner permissions defined on the machine learning resource. To resolve this issue, choose to [inherit resource owner permissions](#) (console) or [remove the permissions from the Lambda function's resource access policy](#) (API).

Function <function-arn> refers to Machine Learning resource <resource-id> with missing permission in both ResourceAccessPolicy and resource OwnerSetting.

Solution: You receive this error if permissions to the machine learning resource aren't configured for the attached Lambda function or the resource. To resolve this issue, configure permissions in the [ResourceAccessPolicy](#) property for the Lambda function or the [OwnerSetting](#) property for the resource.

Function <function-arn> refers to Machine Learning resource <resource-id> with permission \"rw\", while resource owner setting GroupPermission only allows \"ro\".

Solution: You receive this error if the access permissions defined for the attached Lambda function exceed the resource owner permissions defined for the machine learning resource. To resolve this issue, set more restrictive permissions for the Lambda function or less restrictive permissions for the resource owner.

NoContainer Function <function-arn> refers to resources of nested destination path.

Solution: You receive this error if multiple machine learning resources attached to a non-containerized Lambda function use the same destination path or a nested destination path. To resolve this issue, specify separate destination paths for the resources.

Lambda <function-arn> gains access to resource <resource-id> by sharing the same group owner id

Solution: You receive this error in `runtime.log` if the same OS group is specified as the Lambda function's [Run as](#) identity and the [resource owner](#) for a machine learning resource, but the resource is not attached to the Lambda function. This configuration gives the Lambda function implicit permissions that it can use to access the resource without AWS IoT Greengrass authorization.

To resolve this issue, use a different OS group for one of the properties or attach the machine learning resource to the Lambda function.

See also

- [Perform machine learning inference](#)
- [the section called “How to configure machine learning inference”](#)
- [the section called “How to configure optimized machine learning inference”](#)
- [AWS IoT Greengrass Version 1 API Reference](#)

How to configure machine learning inference using the AWS Management Console

To follow the steps in this tutorial, you need AWS IoT Greengrass Core v1.10 or later.

You can perform machine learning (ML) inference locally on a Greengrass core device using locally generated data. For information, including requirements and constraints, see [Perform machine learning inference](#).

This tutorial describes how to use the AWS Management Console to configure a Greengrass group to run a Lambda inference app that recognizes images from a camera locally, without sending data to the cloud. The inference app accesses the camera module on a Raspberry Pi and runs inference using the open source [SqueezeNet](#) model.

The tutorial contains the following high-level steps:

1. [Configure the Raspberry Pi](#)

2. [Install the MXNet framework](#)
3. [Create a model package](#)
4. [Create and publish a Lambda function](#)
5. [Add the Lambda function to the group](#)
6. [Add resources to the group](#)
7. [Add a subscription to the group](#)
8. [Deploy the group](#)
9. [Test the app](#)

Prerequisites

To complete this tutorial, you need:

- Raspberry Pi 4 Model B, or Raspberry Pi 3 Model B/B+, set up and configured for use with AWS IoT Greengrass. To set up your Raspberry Pi with AWS IoT Greengrass, run the [Greengrass Device Setup](#) script, or make sure that you have completed [Module 1](#) and [Module 2](#) of [Getting started with AWS IoT Greengrass](#).

Note

The Raspberry Pi might require a 2.5A [power supply](#) to run the deep learning frameworks that are typically used for image classification. A power supply with a lower rating might cause the device to reboot.

- [Raspberry Pi Camera Module V2 - 8 megapixel, 1080p](#). For information about how to set up the camera, see [Connecting the camera](#) in the Raspberry Pi documentation.
- A Greengrass group and a Greengrass core. For information about how to create a Greengrass group or core, see [Getting started with AWS IoT Greengrass](#).

Note

This tutorial uses a Raspberry Pi, but AWS IoT Greengrass supports other platforms, such as [Intel Atom](#) and [NVIDIA Jetson TX2](#). In the example for Jetson TX2, you can use static images instead of images streamed from a camera. If using the Jetson TX2 example, you might need to install Python 3.6 instead of Python 3.7. For information about configuring

your device so that you can install the AWS IoT Greengrass Core software, see [the section called “Setting up other devices”](#).

For third party platforms that AWS IoT Greengrass does not support, you must run your Lambda function in non-containerized mode. To run in non-containerized mode, you must run your Lambda function as root. For more information, see [the section called “Considerations when choosing Lambda function containerization”](#) and [the section called “Setting the default access identity for Lambda functions in a group”](#).

Step 1: Configure the Raspberry Pi

In this step, install updates to the Raspbian operating system, install the camera module software and Python dependencies, and enable the camera interface.

Run the following commands in your Raspberry Pi terminal.

1. Install updates to Raspbian.

```
sudo apt-get update
sudo apt-get dist-upgrade
```

2. Install the picamera interface for the camera module and other Python libraries that are required for this tutorial.

```
sudo apt-get install -y python3-dev python3-setuptools python3-pip python3-picamera
```

Validate the installation:

- Make sure that your Python 3.7 installation includes pip.

```
python3 -m pip
```

If pip isn't installed, download it from the [pip website](#) and then run the following command.

```
python3 get-pip.py
```

- Make sure that your Python version is 3.7 or higher.

```
python3 --version
```

If the output lists an earlier version, run the following command.

```
sudo apt-get install -y python3.7-dev
```

- Make sure that Setuptools and Picamera installed successfully.

```
sudo -u ggc_user bash -c 'python3 -c "import setuptools"'  
sudo -u ggc_user bash -c 'python3 -c "import picamera"'
```

If the output doesn't contain errors, the validation is successful.

 Note

If the Python executable installed on your device is `python3.7`, use `python3.7` instead of `python3` for the commands in this tutorial. Make sure that your pip installation maps to the correct `python3.7` or `python3` version to avoid dependency errors.

3. Reboot the Raspberry Pi.

```
sudo reboot
```

4. Open the Raspberry Pi configuration tool.

```
sudo raspi-config
```

5. Use the arrow keys to open **Interfacing Options** and enable the camera interface. If prompted, allow the device to reboot.
6. Use the following command to test the camera setup.

```
raspistill -v -o test.jpg
```

This opens a preview window on the Raspberry Pi, saves a picture named `test.jpg` to your current directory, and displays information about the camera in the Raspberry Pi terminal.

Step 2: Install the MXNet framework

In this step, install MXNet libraries on your Raspberry Pi.

1. Sign in to your Raspberry Pi remotely.

```
ssh pi@your-device-ip-address
```

2. Open the MXNet documentation, open [Installing MXNet](#), and follow the instructions to install MXNet on the device.

Note

We recommend installing version 1.5.0 and building MXNet from source for this tutorial to avoid device conflicts.

3. After you install MXNet, validate the following configuration:

- Make sure the `ggc_user` system account can use the MXNet framework.

```
sudo -u ggc_user bash -c 'python3 -c "import mxnet"'
```

- Make sure NumPy is installed.

```
sudo -u ggc_user bash -c 'python3 -c "import numpy"'
```

Step 3: Create an MXNet model package

In this step, create a model package that contains a sample pretrained MXNet model to upload to Amazon Simple Storage Service (Amazon S3). AWS IoT Greengrass can use a model package from Amazon S3, provided that you use the `tar.gz` or `zip` format.

1. On your computer, download the MXNet sample for Raspberry Pi from [the section called "Machine learning samples"](#).
2. Unzip the downloaded `mxnet-py3-armv7l.tar.gz` file.
3. Navigate to the `squeezenet` directory.

```
cd path-to-downloaded-sample/mxnet-py3-armv7l/models/squeezenet
```

The `squeezenet.zip` file in this directory is your model package. It contains SqueezeNet open source model artifacts for an image classification model. Later, you upload this model package to Amazon S3.

Step 4: Create and publish a Lambda function

In this step, create a Lambda function deployment package and Lambda function. Then, publish a function version and create an alias.

First, create the Lambda function deployment package.

1. On your computer, navigate to the `examples` directory in the sample package that you unzipped in [the section called “Create a model package”](#).

```
cd path-to-downloaded-sample/mxnet-py3-armv7l/examples
```

The `examples` directory contains function code and dependencies.

- `greengrassObjectClassification.py` is the inference code used in this tutorial. You can use this code as a template to create your own inference function.
- `greengrasssdk` is version 1.5.0 of the AWS IoT Greengrass Core SDK for Python.

Note

If a new version is available, you can download it and upgrade the SDK version in your deployment package. For more information, see [AWS IoT Greengrass Core SDK for Python](#) on GitHub.

2. Compress the contents of the `examples` directory into a file named `greengrassObjectClassification.zip`. This is your deployment package.

```
zip -r greengrassObjectClassification.zip .
```

Note

Make sure the `.py` files and dependencies are in the root of the directory.

Next, create the Lambda function.


3. From the AWS IoT console, choose **Functions** and **Create function**.
4. Choose **Author from scratch** and use the following values to create your function:
 - For **Function name**, enter **greengrassObjectClassification**.
 - For **Runtime**, choose **Python 3.7**.

For **Permissions**, keep the default setting. This creates an execution role that grants basic Lambda permissions. This role isn't used by AWS IoT Greengrass.

5. Choose **Create function**.

Now, upload your Lambda function deployment package and register the handler.

6. Choose your Lambda function and upload your Lambda function deployment package.
 - a. On the **Code** tab, under **Code source**, choose **Upload from**. From the dropdown, choose **.zip file**.



- b. Choose **Upload**, and then choose your `greengrassObjectClassification.zip` deployment package. Then, choose **Save**.
 - c. On the **Code** tab for the function, under **Runtime settings**, choose **Edit**, and then enter the following values.
 - For **Runtime**, choose **Python 3.7**.
 - For **Handler**, enter **greengrassObjectClassification.function_handler**.


Choose **Save**.

Next, publish the first version of your Lambda function. Then, create an [alias for the version](#).


Note

Greengrass groups can reference a Lambda function by alias (recommended) or by version. Using an alias makes it easier to manage code updates because you don't have to change your subscription table or group definition when the function code is updated. Instead, you just point the alias to the new function version.

7. From the **Actions** menu, choose **Publish new version**.



8. For **Version description**, enter **First version**, and then choose **Publish**.
9. On the **greengrassObjectClassification: 1** configuration page, from the **Actions** menu, choose **Create alias**.



10. On the **Create a new alias** page, use the following values:
- For **Name**, enter **m1Test**.
 - For **Version**, enter **1**.

Note

AWS IoT Greengrass doesn't support Lambda aliases for **\$LATEST** versions.

11. Choose **Save**.

Now, add the Lambda function to your Greengrass group.

Step 5: Add the Lambda function to the Greengrass group

In this step, add the Lambda function to the group and then configure its lifecycle and environment variables.

First, add the Lambda function to your Greengrass group.

1. In the AWS IoT console navigation pane, under **Manage**, expand **Greengrass devices**, and then choose **Groups (V1)**.
2. From the group configuration page, choose the **Lambda functions** tab.
3. Under the **My Lambda functions** section, choose **Add**.
4. For the **Lambda function**, choose **greengrassObjectClassification**.
5. For the **Lambda function version**, choose **Alias:mlTest**.

Next, configure the lifecycle and environment variables of the Lambda function.

6. On the **Lambda function configuration** section, make the following updates.

Note

We recommend that you run your Lambda function without containerization unless your business case requires it. This helps enable access to your device GPU and camera without configuring device resources. If you run without containerization, you must also grant root access to your AWS IoT Greengrass Lambda functions.

a. To run without containerization:

- For **System user and group**, choose **Another user ID/group ID**. For **System user ID**, enter **0**. For **System group ID**, enter **0**.

This allows your Lambda function to run as root. For more information about running as root, see [the section called "Setting the default access identity for Lambda functions in a group"](#).

Tip

You also must update your `config.json` file to grant root access to your Lambda function. For the procedure, see [the section called “Running a Lambda function as root”](#).

- For **Lambda function containerization**, choose **No container**.

For more information about running without containerization, see [the section called “Considerations when choosing Lambda function containerization”](#).

- For **Timeout**, enter **10 seconds**.
- For **Pinned**, choose **True**.

For more information, see [the section called “Lifecycle configuration”](#).

b. **To run in containerized mode instead:**

Note

We do not recommend running in containerized mode unless your business case requires it.

- For **System user and group**, choose **Use group default**.
- For **Lambda function containerization**, choose **Use group default**.
- For **Memory limit**, enter **96 MB**.
- For **Timeout**, enter **10 seconds**.
- For **Pinned**, choose **True**.

For more information, see [the section called “Lifecycle configuration”](#).

7. Under **Environment variables**, create a key-value pair. A key-value pair is required by functions that interact with MXNet models on a Raspberry Pi.

For the key, use `MXNET_ENGINE_TYPE`. For the value, use `NaiveEngine`.

Note

In your own user-defined Lambda functions, you can optionally set the environment variable in your function code.

8. Keep the default values for all other properties and choose **Add Lambda function**.

Step 6: Add resources to the Greengrass group

In this step, create resources for the camera module and the ML inference model and affiliate the resources with the Lambda function. This makes it possible for the Lambda function to access the resources on the core device.

Note

If you run in non-containerized mode, AWS IoT Greengrass can access your device GPU and camera without configuring these device resources.

First, create two local device resources for the camera: one for shared memory and one for the device interface. For more information about local resource access, see [Access local resources with Lambda functions and connectors](#).

1. On the group configuration page, choose the **Resources** tab.
2. In the **Local resources** section, choose **Add local resource**.
3. On the **Add a local resource** page, use the following values:
 - For **Resource name**, enter **videoCoreSharedMemory**.
 - For **Resource type**, choose **Device**.
 - For **Local device path**, enter **/dev/vcsm**.

The device path is the local absolute path of the device resource. This path can only refer to a character device or block device under **/dev**.

- For **System group owner and file access permissions**, choose **Automatically add file system permissions of the system group that owns the resource**.

The **System group owner and file access permissions** option lets you grant additional file access permissions to the Lambda process. For more information, see [Group owner file access permission](#).

4. Next, you add a local device resource for the camera interface.
5. Choose **Add local resource**.
6. On the **Add a local resource** page, use the following values:
 - For **Resource name**, enter **videoCoreInterface**.
 - For **Resource type**, choose **Device**.
 - For **Local device path**, enter **/dev/vchiq**.
 - For **System group owner and file access permissions**, choose **Automatically add file system permissions of the system group that owns the resource**.
7. At the bottom of the page, choose **Add resource**.

Now, add the inference model as a machine learning resource. This step includes uploading the `squeezenet.zip` model package to Amazon S3.

1. On the **Resources** tab for your group, under the **Machine Learning** section, choose **Add machine learning resource**.
2. On the **Add a machine learning resource** page, for **Resource name**, enter **squeezenet_model**.
3. For **Model source**, choose **Use a model stored in S3, such as a model optimized through Deep Learning Compiler**.
4. For **S3 URI**, enter a path where the S3 bucket is saved.
5. Choose **Browse S3**. This opens up a new tab to the Amazon S3 console.
6. On the Amazon S3 console tab, upload the `squeezenet.zip` file to an S3 bucket. For information, see [How do I upload files and folders to an S3 Bucket?](#) in the *Amazon Simple Storage Service User Guide*.

 Note

For the S3 bucket to be accessible, your bucket name must contain the string **greengrass** and the bucket must be in the same region that you use for AWS IoT

Greengrass. Choose a unique name (such as **greengrass-bucket-*user-id-epoch-time***). Don't use a period (.) in the bucket name.

7. On the AWS IoT Greengrass console tab, locate and choose your S3 bucket. Locate your uploaded `squeezenet.zip` file, and choose **Select**. You might need to choose **Refresh** to update the list of available buckets and files.
8. For **Destination path**, enter **/greengrass-machine-learning/mxnet/squeezenet**.

This is the destination for the local model in the Lambda runtime namespace. When you deploy the group, AWS IoT Greengrass retrieves the source model package and then extracts the contents to the specified directory. The sample Lambda function for this tutorial is already configured to use this path (in the `model_path` variable).

9. Under **System group owner and file access permissions**, choose **No system group**.
10. Choose **Add resource**.

Using SageMaker AI trained models

This tutorial uses a model that's stored in Amazon S3, but you can easily use SageMaker AI models too. The AWS IoT Greengrass console has built-in SageMaker AI integration, so you don't need to manually upload these models to Amazon S3. For requirements and limitations for using SageMaker AI models, see [the section called "Supported model sources"](#).

To use an SageMaker AI model:

- For **Model source**, choose **Use a model trained in AWS SageMaker AI**, and then choose the name of the model's training job.
- For **Destination path**, enter the path to the directory where your Lambda function looks for the model.

Step 7: Add a subscription to the Greengrass group

In this step, add a subscription to the group. This subscription enables the Lambda function to send prediction results to AWS IoT by publishing to an MQTT topic.

1. On the group configuration page, choose the **Subscriptions** tab, and then choose **Add Subscription**.

2. On the **Subscription details** page, configure the source and target, as follows:
 - a. In **Source type**, choose **Lambda function**, and then choose **greengrassObjectClassification**.
 - b. In **Target type**, choose **Service**, and then choose **IoT Cloud**.
3. In **Topic filter**, enter **hello/world**, and then choose **Create subscription**.

Step 8: Deploy the Greengrass group

In this step, deploy the current version of the group definition to the Greengrass core device. The definition contains the Lambda function, resources, and subscription configurations that you added.

1. Make sure that the AWS IoT Greengrass core is running. Run the following commands in your Raspberry Pi terminal, as needed.
 - a. To check whether the daemon is running:

```
ps aux | grep -E 'greengrass.*daemon'
```

If the output contains a root entry for `/greengrass/ggc/packages/1.11.6/bin/daemon`, then the daemon is running.


Note

The version in the path depends on the AWS IoT Greengrass Core software version that's installed on your core device.

- b. To start the daemon:

```
cd /greengrass/ggc/core/  
sudo ./greengrassd start
```

2. On the group configuration page, choose **Deploy**.



3. In the **Lambda functions** tab, under the **System Lambda functions** section, select **IP detector** and choose **Edit**.
4. In the **Edit IP detector settings** dialog box, select **Automatically detect and override MQTT broker endpoints**.
5. Choose **Save**.

This enables devices to automatically acquire connectivity information for the core, such as IP address, DNS, and port number. Automatic detection is recommended, but AWS IoT Greengrass also supports manually specified endpoints. You're only prompted for the discovery method the first time that the group is deployed.

Note

If prompted, grant permission to create the [Greengrass service role](#) and associate it with your AWS account in the current AWS Region. This role allows AWS IoT Greengrass to access your resources in AWS services.

The **Deployments** page shows the deployment timestamp, version ID, and status. When completed, the status displayed for the deployment should be **Completed**.

For more information about deployments, see [Deploy AWS IoT Greengrass groups](#). For troubleshooting help, see [Troubleshooting](#).

Step 9: Test the inference app

Now you can verify whether the deployment is configured correctly. To test, you subscribe to the hello/world topic and view the prediction results that are published by the Lambda function.

Note

If a monitor is attached to the Raspberry Pi, the live camera feed is displayed in a preview window.

1. In the AWS IoT console, under **Test**, choose **MQTT test client**.
2. For **Subscriptions**, use the following values:
 - For the subscription topic, use `hello/world`.
 - Under **Additional configuration**, for **MQTT payload display**, choose **Display payloads as strings**.
3. Choose **Subscribe**.

If the test is successful, the messages from the Lambda function appear at the bottom of the page. Each message contains the top five prediction results of the image, using the format: probability, predicted class ID, and corresponding class name.

The screenshot shows the AWS IoT console's MQTT test client interface. On the left, there's a sidebar with 'Subscribe to a topic' and 'Publish to a topic' buttons. The 'hello/world' topic is selected. The main area shows a 'Publish' section with a text input containing 'hello/world' and a 'Publish to topic' button. Below this is a code editor showing a JSON message: `{ "message": "Hello from AWS IoT console" }`. At the bottom, a table displays received messages. The first message is a 'New Prediction' with five results. The second and third messages are also 'New Prediction' results. The table has 'Export' and 'Hide' buttons for each row.

Topic	Time	Message	Export	Hide
hello/world	Mar 30, 2018 1:47:07 PM -0700	New Prediction: [(0.31046376, 'n03637318 lampshade, lamp shade'), (0.11445289, 'n04380533 table lamp'), (0.04436367, 'n04254120 soap dispenser'), (0.035816364, 'n04286575 spotlight, spot'), (0.028093718, 'n03201208 dining table, board')]	Export	Hide
hello/world	Mar 30, 2018 1:47:01 PM -0700	New Prediction: [(0.16117829, 'n03876231 paintbrush'), (0.13750333, 'n04442312 toaster'), (0.081819646, 'n03924679 photocopier'), (0.068144165, 'n02783161 ballpoint, ballpoint pen, ballpen, Biro'), (0.044701375, 'n04209239 shower curtain')]	Export	Hide
hello/world	Mar 30, 2018 1:46:55 PM -0700	New Prediction: [(0.46284258, 'n04442312 toaster'), (0.16061385, 'n03908618 pencil box, pencil case'), (0.043834824, 'n03291819 envelope'), (0.027529096, 'n03908714 pencil sharpener'), (0.027273422, 'n04209239 shower curtain')]	Export	Hide

Troubleshooting AWS IoT Greengrass ML inference

If the test is not successful, you can try the following troubleshooting steps. Run the commands in your Raspberry Pi terminal.

Check error logs

1. Switch to the root user and navigate to the log directory. Access to AWS IoT Greengrass logs requires root permissions.

```
sudo su
cd /greengrass/ggc/var/log
```

2. In the system directory, check `runtime.log` or `python_runtime.log`.

In the user/*region*/*account-id* directory, check `greengrassObjectClassification.log`.

For more information, see [the section called “Troubleshooting with logs”](#).

Unpacking error in runtime.log

If `runtime.log` contains an error similar to the following, make sure that your `tar.gz` source model package has a parent directory.

```
Greengrass deployment error: unable to download the artifact model-arn: Error while
processing.
Error while unpacking the file from /tmp/greengrass/artifacts/model-arn/path to /
greengrass/ggc/deployment/path/model-arn,
error: open /greengrass/ggc/deployment/path/model-arn/squeezenet/
squeezenet_v1.1-0000.params: no such file or directory
```

If your package doesn't have a parent directory that contains the model files, use the following command to repack the model:

```
tar -zcvf model.tar.gz ./model
```

For example:

```
#$ tar -zcvf test.tar.gz ./test
./test
./test/some.file
./test/some.file2
./test/some.file3
```

Note

Don't include trailing `/*` characters in this command.

Verify that the Lambda function is successfully deployed

1. List the contents of the deployed Lambda in the `/lambda` directory. Replace the placeholder values before you run the command.

```
cd /greengrass/ggc/deployment/lambda/  
arn:aws:lambda:region:account:function:function-name:function-version  
ls -la
```

2. Verify that the directory contains the same content as the `greengrassObjectClassification.zip` deployment package that you uploaded in [Step 4: Create and publish a Lambda function](#).

Make sure that the `.py` files and dependencies are in the root of the directory.

Verify that the inference model is successfully deployed

1. Find the process identification number (PID) of the Lambda runtime process:

```
ps aux | grep 'lambda-function-name*
```

In the output, the PID appears in the second column of the line for the Lambda runtime process.

2. Enter the Lambda runtime namespace. Be sure to replace the placeholder *pid* value before you run the command.

Note

This directory and its contents are in the Lambda runtime namespace, so they aren't visible in a regular Linux namespace.

```
sudo nsenter -t pid -m /bin/bash
```

3. List the contents of the local directory that you specified for the ML resource.

```
cd /greengrass-machine-learning/mxnet/squeezenet/  
ls -ls
```

You should see the following files:

```
32 -rw-r--r-- 1 ggc_user ggc_group 31675 Nov 18 15:19 synset.txt  
32 -rw-r--r-- 1 ggc_user ggc_group 28707 Nov 18 15:19 squeezenet_v1.1-symbol.json  
4832 -rw-r--r-- 1 ggc_user ggc_group 4945062 Nov 18 15:19  
squeezenet_v1.1-0000.params
```

Next steps

Next, explore other inference apps. AWS IoT Greengrass provides other Lambda functions that you can use to try out local inference. You can find the examples package in the precompiled libraries folder that you downloaded in [the section called “Install the MXNet framework”](#).

Configuring an Intel Atom

To run this tutorial on an Intel Atom device, you must provide source images, configure the Lambda function, and add another local device resource. To use the GPU for inference, make sure the following software is installed on your device:

- OpenCL version 1.0 or later
- Python 3.7 and pip

Note

If your device is prebuilt with Python 3.6, you can create a symlink to Python 3.7 instead. For more information, see [Step 2](#).

- [NumPy](#)
- [OpenCV on Wheels](#)

1. Download static PNG or JPG images for the Lambda function to use for image classification. The example works best with small image files.

Save your image files in the directory that contains the `greengrassObjectClassification.py` file (or in a subdirectory of this directory). This is in the Lambda function deployment package that you upload in [the section called "Create and publish a Lambda function"](#).

Note

If you're using AWS DeepLens, you can use the onboard camera or mount your own camera to perform inference on captured images instead of static images. However, we strongly recommend you start with static images first.

If you use a camera, make sure that the `awscam` APT package is installed and up to date. For more information, see [Update your AWS DeepLens device](#) in the *AWS DeepLens Developer Guide*.

2. If you aren't using Python 3.7, make sure to create a symlink from Python 3.x to Python 3.7. This configures your device to use Python 3 with AWS IoT Greengrass. Run the following command to locate your Python installation:

```
which python3
```

Run the following command to create the symlink:

```
sudo ln -s path-to-python-3.x/python3.x path-to-python-3.7/python3.7
```

Reboot the device.

3. Edit the configuration of the Lambda function. Follow the procedure in [the section called “Add the Lambda function to the group”](#).

 Note

We recommend that you run your Lambda function without containerization unless your business case requires it. This helps enable access to your device GPU and camera without configuring device resources. If you run without containerization, you must also grant root access to your AWS IoT Greengrass Lambda functions.

a. **To run without containerization:**

- For **System user and group**, choose **Another user ID/group ID**. For **System user ID**, enter **0**. For **System group ID**, enter **0**.

This allows your Lambda function to run as root. For more information about running as root, see [the section called “Setting the default access identity for Lambda functions in a group”](#).

 Tip

You also must update your `config.json` file to grant root access to your Lambda function. For the procedure, see [the section called “Running a Lambda function as root”](#).

- For **Lambda function containerization**, choose **No container**.

For more information about running without containerization, see [the section called “Considerations when choosing Lambda function containerization”](#).

- Update the **Timeout** value to 5 seconds. This ensures that the request does not time out too early. It takes a few minutes after setup to run inference.
- Under **Pinned**, choose **True**.
- Under **Additional Parameters**, for **Read access to /sys directory**, choose **Enabled**.
- For **Lambda lifecycle**, choose **Make this function long-lived and keep it running indefinitely**.

b. **To run in containerized mode instead:**

Note

We do not recommend running in containerized mode unless your business case requires it.

- Update the **Timeout** value to 5 seconds. This ensures that the request does not time out too early. It takes a few minutes after setup to run inference.
 - For **Pinned**, choose **True**.
 - Under **Additional Parameters**, for **Read access to /sys directory**, choose **Enabled**.
4. **If running in containerized mode**, add the required local device resource to grant access to your device GPU.

Note

If you run in non-containerized mode, AWS IoT Greengrass can access your device GPU without configuring device resources.

- a. On the group configuration page, choose the **Resources** tab.
- b. Choose **Add local resource**.
- c. Define the resource:
 - For **Resource name**, enter **renderD128**.
 - For **Resource type**, choose **Local device**.
 - For **Device path**, enter **/dev/dri/renderD128**.
 - For **System group owner and file access permissions**, choose **Automatically add file system permissions of the system group that owns the resource**.
 - For **Lambda function affiliations**, grant **Read and write access** to your Lambda function.

Configuring an NVIDIA Jetson TX2

To run this tutorial on an NVIDIA Jetson TX2, provide source images and configure the Lambda function. If you're using the GPU, you must also add local device resources.

1. Make sure your Jetson device is configured so you can install the AWS IoT Greengrass Core software. For more information about configuring your device, see [the section called “Setting up other devices”](#).
2. Open the MXNet documentation, go to [Installing MXNet on a Jetson](#), and follow the instructions to install MXNet on the Jetson device.

Note

If you want to build MXNet from source, follow the instructions to build the shared library. Edit the following settings in your `config.mk` file to work with a Jetson TX2 device:

- Add `-gencode arch=compute-62, code=sm_62` to the `CUDA_ARCH` setting.
- Turn on CUDA.

```
USE_CUDA = 1
```

3. Download static PNG or JPG images for the Lambda function to use for image classification. The app works best with small image files. Alternatively, you can instrument a camera on the Jetson board to capture the source images.

Save your image files in the directory that contains the `greengrassObjectClassification.py` file. You can also save them in a subdirectory of this directory. This directory is in the Lambda function deployment package that you upload in [the section called “Create and publish a Lambda function”](#).

4. Create a symlink from Python 3.7 to Python 3.6 to use Python 3 with AWS IoT Greengrass. Run the following command to locate your Python installation:

```
which python3
```

Run the following command to create the symlink:

```
sudo ln -s path-to-python-3.6/python3.6 path-to-python-3.7/python3.7
```

Reboot the device.

5. Make sure the `ggc_user` system account can use the MXNet framework:

```
"sudo -u ggc_user bash -c 'python3 -c "import mxnet"'
```

6. Edit the configuration of the Lambda function. Follow the procedure in [the section called "Add the Lambda function to the group"](#).

 Note

We recommend that you run your Lambda function without containerization unless your business case requires it. This helps enable access to your device GPU and camera without configuring device resources. If you run without containerization, you must also grant root access to your AWS IoT Greengrass Lambda functions.

a. **To run without containerization:**

- For **System user and group**, choose **Another user ID/group ID**. For **System user ID**, enter `0`. For **System group ID**, enter `0`.

This allows your Lambda function to run as root. For more information about running as root, see [the section called "Setting the default access identity for Lambda functions in a group"](#).

 Tip

You also must update your `config.json` file to grant root access to your Lambda function. For the procedure, see [the section called "Running a Lambda function as root"](#).

- For **Lambda function containerization**, choose **No container**.

For more information about running without containerization, see [the section called "Considerations when choosing Lambda function containerization"](#).

- Under **Additional Parameters**, for **Read access to /sys directory**, choose **Enabled**.
- Under **Environment variables**, add the following key-value pairs to your Lambda function. This configures AWS IoT Greengrass to use the MXNet framework.

Key	Value
PATH	/usr/local/cuda/bin:\$PATH
MXNET_HOME	\$HOME/mxnet/
PYTHONPATH	\$MXNET_HOME/python:\$PYTHONPATH
CUDA_HOME	/usr/local/cuda
LD_LIBRARY_PATH	\$LD_LIBRARY_PATH:\${CUDA_HOME}/lib64

b. **To run in containerized mode instead:**

 Note

We do not recommend running in containerized mode unless your business case requires it.

- Increase the **Memory limit** value. Use 500 MB for CPU, or at least 2000 MB for GPU.
- Under **Additional Parameters**, for **Read access to /sys directory**, choose **Enabled**.
- Under **Environment variables**, add the following key-value pairs to your Lambda function. This configures AWS IoT Greengrass to use the MXNet framework.

Key	Value
PATH	/usr/local/cuda/bin:\$PATH
MXNET_HOME	\$HOME/mxnet/
PYTHONPATH	\$MXNET_HOME/python:\$PYTHONPATH

Key	Value
CUDA_HOME	/usr/local/cuda
LD_LIBRARY_PATH	\$LD_LIBRARY_PATH:\${CUDA_HOME}/lib64

7. **If running in containerized mode**, add the following local device resources to grant access to your device GPU. Follow the procedure in [the section called “Add resources to the group”](#).

 Note

If you run in non-containerized mode, AWS IoT Greengrass can access your device GPU without configuring device resources.

For each resource:

- For **Resource type**, choose **Device**.
- For **System group owner and file access permissions**, choose **Automatically add file system permissions of the system group that owns the resource**.

Name	Device path
nvhost-ctrl	/dev/nvhost-ctrl
nvhost-gpu	/dev/nvhost-gpu
nvhost-ctrl-gpu	/dev/nvhost-ctrl-gpu
nvhost-dbg-gpu	/dev/nvhost-dbg-gpu
nvhost-prof-gpu	/dev/nvhost-prof-gpu
nvmap	/dev/nvmap
nvhost-vic	/dev/nvhost-vic

Name	Device path
tegra_dc_ctrl	/dev/tegra_dc_ctrl

8. **If running in containerized mode**, add the following local volume resource to grant access to your device camera. Follow the procedure in [the section called “Add resources to the group”](#).

 Note

If you run in non-containerized mode, AWS IoT Greengrass can access your device camera without configuring volume resources.

- For **Resource type**, choose **Volume**.
- For **System group owner and file access permissions**, choose **Automatically add file system permissions of the system group that owns the resource**.

Name	Source path	Destination path
shm	/dev/shm	/dev/shm
tmp	/tmp	/tmp

How to configure optimized machine learning inference using the AWS Management Console

To follow the steps in this tutorial, you must be using AWS IoT Greengrass Core v1.10 or later.

You can use the SageMaker AI Neo deep learning compiler to optimize the prediction efficiency of native machine learning inference models in Tensorflow, Apache MXNet, PyTorch, ONNX, and XGBoost frameworks for a smaller footprint and faster performance. You can then download the optimized model and install the SageMaker AI Neo deep learning runtime and deploy them to your AWS IoT Greengrass devices for faster inference.

This tutorial describes how to use the AWS Management Console to configure a Greengrass group to run a Lambda inference example that recognizes images from a camera locally, without sending data to the cloud. The inference example accesses the camera module on a Raspberry Pi. In this tutorial, you download a prepackaged model that is trained by Resnet-50 and optimized in the Neo deep learning compiler. You then use the model to perform local image classification on your AWS IoT Greengrass device.

The tutorial contains the following high-level steps:

1. [Configure the Raspberry Pi](#)
2. [Install the Neo deep learning runtime](#)
3. [Create an inference Lambda function](#)
4. [Add the Lambda function to the group](#)
5. [Add a Neo-optimized model resource to the group](#)
6. [Add your camera device resource to the group](#)
7. [Add subscriptions to the group](#)
8. [Deploy the group](#)
9. [Test the example](#)

Prerequisites

To complete this tutorial, you need:

- Raspberry Pi 4 Model B, or Raspberry Pi 3 Model B/B+, set up and configured for use with AWS IoT Greengrass. To set up your Raspberry Pi with AWS IoT Greengrass, run the [Greengrass Device Setup](#) script, or make sure that you have completed [Module 1](#) and [Module 2](#) of [Getting started with AWS IoT Greengrass](#).

Note

The Raspberry Pi might require a 2.5A [power supply](#) to run the deep learning frameworks that are typically used for image classification. A power supply with a lower rating might cause the device to reboot.

- [Raspberry Pi Camera Module V2 - 8 megapixel, 1080p](#). To learn how to set up the camera, see [Connecting the camera](#) in the Raspberry Pi documentation.

- A Greengrass group and a Greengrass core. To learn how to create a Greengrass group or core, see [Getting started with AWS IoT Greengrass](#).

Note

This tutorial uses a Raspberry Pi, but AWS IoT Greengrass supports other platforms, such as [Intel Atom](#) and [NVIDIA Jetson TX2](#). If using the Intel Atom example, you might need to install Python 3.6 instead of Python 3.7. For information about configuring your device so you can install the AWS IoT Greengrass Core software, see [the section called “Setting up other devices”](#).

For third party platforms that AWS IoT Greengrass does not support, you must run your Lambda function in non-containerized mode. To run in non-containerized mode, you must run your Lambda function as root. For more information, see [the section called “Considerations when choosing Lambda function containerization”](#) and [the section called “Setting the default access identity for Lambda functions in a group”](#).

Step 1: Configure the Raspberry Pi

In this step, install updates to the Raspbian operating system, install the camera module software and Python dependencies, and enable the camera interface.

Run the following commands in your Raspberry Pi terminal.

1. Install updates to Raspbian.

```
sudo apt-get update
sudo apt-get dist-upgrade
```

2. Install the `picamera` interface for the camera module and other Python libraries that are required for this tutorial.

```
sudo apt-get install -y python3-dev python3-setuptools python3-pip python3-picamera
```

Validate the installation:

- Make sure that your Python 3.7 installation includes `pip`.

```
python3 -m pip
```

If pip isn't installed, download it from the [pip website](#) and then run the following command.

```
python3 get-pip.py
```

- Make sure that your Python version is 3.7 or higher.

```
python3 --version
```

If the output lists an earlier version, run the following command.

```
sudo apt-get install -y python3.7-dev
```

- Make sure that Setuptools and Picamera installed successfully.

```
sudo -u ggc_user bash -c 'python3 -c "import setuptools"'  
sudo -u ggc_user bash -c 'python3 -c "import picamera"'
```

If the output doesn't contain errors, the validation is successful.

Note

If the Python executable installed on your device is `python3.7`, use `python3.7` instead of `python3` for the commands in this tutorial. Make sure that your pip installation maps to the correct `python3.7` or `python3` version to avoid dependency errors.

3. Reboot the Raspberry Pi.

```
sudo reboot
```

4. Open the Raspberry Pi configuration tool.

```
sudo raspi-config
```

5. Use the arrow keys to open **Interfacing Options** and enable the camera interface. If prompted, allow the device to reboot.
6. Use the following command to test the camera setup.

```
raspistill -v -o test.jpg
```

This opens a preview window on the Raspberry Pi, saves a picture named `test.jpg` to your current directory, and displays information about the camera in the Raspberry Pi terminal.

Step 2: Install the Amazon SageMaker Neo deep learning runtime

In this step, install the Neo deep learning runtime (DLR) on your Raspberry Pi.

Note

We recommend installing version 1.1.0 for this tutorial.

1. Sign in to your Raspberry Pi remotely.

```
ssh pi@your-device-ip-address
```

2. Open the DLR documentation, open [Installing DLR](#), and locate the wheel URL for Raspberry Pi devices. Then, follow the instructions to install the DLR on your device. For example, you can use pip:

```
pip3 install rasp3b-wheel-url
```

3. After you install the DLR, validate the following configuration:

- Make sure the `ggc_user` system account can use the DLR library.

```
sudo -u ggc_user bash -c 'python3 -c "import dlr"'
```

- Make sure NumPy is installed.

```
sudo -u ggc_user bash -c 'python3 -c "import numpy"'
```

Step 3: Create an inference Lambda function

In this step, create a Lambda function deployment package and Lambda function. Then, publish a function version and create an alias.

1. On your computer, download the DLR sample for Raspberry Pi from [the section called “Machine learning samples”](#).
2. Unzip the downloaded `dlr-py3-armv7l.tar.gz` file.

```
cd path-to-downloaded-sample  
tar -xvzf dlr-py3-armv7l.tar.gz
```

The `examples` directory in the extracted sample package contains function code and dependencies.

- `inference.py` is the inference code used in this tutorial. You can use this code as a template to create your own inference function.
- `greengrasssdk` is version 1.5.0 of the AWS IoT Greengrass Core SDK for Python.

Note

If a new version is available, you can download it and upgrade the SDK version in your deployment package. For more information, see [AWS IoT Greengrass Core SDK for Python](#) on GitHub.

3. Compress the contents of the `examples` directory into a file named `optimizedImageClassification.zip`. This is your deployment package.

```
cd path-to-downloaded-sample/dlr-py3-armv7l/examples  
zip -r optimizedImageClassification.zip .
```

The deployment package contains your function code and dependencies. This includes the code that invokes the Neo deep learning runtime Python APIs to perform inference with the Neo deep learning compiler models.

Note

Make sure the .py files and dependencies are in the root of the directory.

- Now, add the Lambda function to your Greengrass group.

From the Lambda console page, choose **Functions** and choose **Create function**.

- Choose **Author from scratch** and use the following values to create your function:

- For **Function name**, enter **optimizedImageClassification**.
- For **Runtime**, choose **Python 3.7**.

For **Permissions**, keep the default setting. This creates an execution role that grants basic Lambda permissions. This role isn't used by AWS IoT Greengrass.

Basic information

Function name
Enter a name that describes the purpose of your function.

Use only letters, numbers, hyphens, or underscores with no spaces.

Runtime [Info](#)
Choose the language to use to write your function.


Permissions [Info](#)
Lambda will create an execution role with permission to upload logs to Amazon CloudWatch Logs. You can configure and modify permissions further when you add triggers.
[► Choose or create an execution role](#)

[Cancel](#) [Create function](#)

- Choose **Create function**.


Now, upload your Lambda function deployment package and register the handler.

- On the **Code** tab, under **Code source**, choose **Upload from**. From the dropdown, choose **.zip file**.



2. Choose your `optimizedImageClassification.zip` deployment package, and then choose **Save**.
3. On the **Code** tab for the function, under **Runtime settings**, choose **Edit**, and then enter the following values.
 - For **Runtime**, choose **Python 3.7**.
 - For **Handler**, enter **`inference.handler`**.

Choose **Save**.




Next, publish the first version of your Lambda function. Then, create an [alias for the version](#).


Note

Greengrass groups can reference a Lambda function by alias (recommended) or by version. Using an alias makes it easier to manage code updates because you don't have to change your subscription table or group definition when the function code is updated. Instead, you just point the alias to the new function version.

1. From the **Actions** menu, choose **Publish new version**.



2. For **Version description**, enter **First version**, and then choose **Publish**.
3. On the **optimizedImageClassification: 1** configuration page, from the **Actions** menu, choose **Create alias**.



4. On the **Create a new alias** page, use the following values:
 - For **Name**, enter **m1Test0pt**.
 - For **Version**, enter **1**.

Note

AWS IoT Greengrass doesn't support Lambda aliases for **\$LATEST** versions.

5. Choose **Create**.

Now, add the Lambda function to your Greengrass group.

Step 4: Add the Lambda function to the Greengrass group

In this step, add the Lambda function to the group, and then configure its lifecycle.

First, add the Lambda function to your Greengrass group.

1. In the AWS IoT console navigation pane, under **Manage**, expand **Greengrass devices**, and then choose **Groups (V1)**.
2. On the groups configuration page, choose the **Lambda functions** tab, and choose **Add**.
3. Choose the **Lambda function** and select **optimizedImageClassification**.
4. On the **Lambda function version**, choose the alias to the version that you published.

Next, configure the lifecycle of the Lambda function.

1. In the **Lambda function configuration** section, make the following updates.

 Note

We recommend that you run your Lambda function without containerization unless your business case requires it. This helps enable access to your device GPU and camera without configuring device resources. If you run without containerization, you must also grant root access to your AWS IoT Greengrass Lambda functions.

a. **To run without containerization:**

- For **System user and group**, choose **Another user ID/group ID**. For **System user ID**, enter **0**. For **System group ID**, enter **0**.

This allows your Lambda function to run as root. For more information about running as root, see [the section called “Setting the default access identity for Lambda functions in a group”](#).

 Tip

You also must update your `config.json` file to grant root access to your Lambda function. For the procedure, see [the section called “Running a Lambda function as root”](#).

- For **Lambda function containerization**, choose **No container**.

For more information about running without containerization, see [the section called “Considerations when choosing Lambda function containerization”](#).

- For **Timeout**, enter **10 seconds**.
- For **Pinned**, choose **True**.

For more information, see [the section called “Lifecycle configuration”](#).

- Under **Additional Parameter**, for **Read access to /sys directory**, choose **Enabled**.

b. **To run in containerized mode instead:**

Note

We do not recommend running in containerized mode unless your business case requires it.

- For **System user and group**, choose **Use group default**.
- For **Lambda function containerization**, choose **Use group default**.
- For **Memory limit**, enter **1024 MB**.
- For **Timeout**, enter **10 seconds**.
- For **Pinned**, choose **True**.

For more information, see [the section called “Lifecycle configuration”](#).

- Under **Additional Parameters**, for **Read access to /sys directory**, choose **Enabled**.

2. Choose **Add Lambda function**.

Step 5: Add a SageMaker AI Neo-optimized model resource to the Greengrass group

In this step, create a resource for the optimized ML inference model and upload it to an Amazon S3 bucket. Then, locate the Amazon S3 uploaded model in the AWS IoT Greengrass console and affiliate the newly created resource with the Lambda function. This makes it possible for the function to access its resources on the core device.

1. On your computer, navigate to the `resnet50` directory in the sample package that you unzipped in [the section called “Create an inference Lambda function”](#).

Note

If using the NVIDIA Jetson example, you need to use the `resnet18` directory in the sample package instead. For more information, see [the section called “Configuring an NVIDIA Jetson TX2”](#).

```
cd path-to-downloaded-sample/dlr-py3-armv7l/models/resnet50
```

This directory contains precompiled model artifacts for an image classification model trained with Resnet-50.

2. Compress the files inside the `resnet50` directory into a file named `resnet50.zip`.

```
zip -r resnet50.zip .
```

3. On the group configuration page for your AWS IoT Greengrass group, choose the **Resources** tab. Navigate to the **Machine Learning** section and choose **Add machine learning resource**. On the **Create a machine learning resource** page, for **Resource name**, enter **resnet50_model**.
4. For **Model source**, choose **Use a model stored in S3, such as a model optimized through Deep Learning Compiler**.
5. Under **S3 URI**, choose **Browse S3**.

 Note

Currently, optimized SageMaker AI models are stored automatically in Amazon S3. You can find your optimized model in your Amazon S3 bucket using this option. For more information about model optimization in SageMaker AI, see the [SageMaker AI Neo documentation](#).

6. Choose **Upload a model**.
7. On the Amazon S3 console tab, upload your zip file to an Amazon S3 bucket. For information, see [How do I upload files and folders to an S3 bucket?](#) in the *Amazon Simple Storage Service User Guide*.

 Note

Your bucket name must contain the string **greengrass**. Choose a unique name (such as **greengrass-dlr-bucket-*user-id-epoch-time***). Don't use a period (.) in the bucket name.

8. In the AWS IoT Greengrass console tab, locate and choose your Amazon S3 bucket. Locate your uploaded `resnet50.zip` file, and choose **Select**. You might need to refresh the page to update the list of available buckets and files.
9. In **Destination path**, enter `/ml_model`.

Local path

`/ml_model`

This is the destination for the local model in the Lambda runtime namespace. When you deploy the group, AWS IoT Greengrass retrieves the source model package and then extracts the contents to the specified directory.

Note

We strongly recommend that you use the exact path provided for your local path. Using a different local model destination path in this step causes some troubleshooting commands provided in this tutorial to be inaccurate. If you use a different path, you must set up a `MODEL_PATH` environment variable that uses the exact path you provide here. For information about environment variables, see [AWS Lambda environment variables](#).

10. **If running in containerized mode:**
 - a. Under **System group owner and file access permissions**, choose **Specify system group and permissions**.
 - b. Choose **Read-only access** and then choose **Add resource**.

Step 6: Add your camera device resource to the Greengrass group

In this step, create a resource for the camera module and affiliate it with the Lambda function. This makes it possible for the Lambda function to access the resource on the core device.

Note

If you run in non-containerized mode, AWS IoT Greengrass can access your device GPU and camera without configuring this device resource.

1. On the group configuration page, choose the **Resources** tab.
2. On the **Local resources** tab, choose **Add local resource**.
3. On the **Add a local resource** page, use the following values:

- For **Resource name**, enter **videoCoreSharedMemory**.
- For **Resource type**, choose **Device**.
- For **Local device path**, enter **/dev/vcsm**.

The device path is the local absolute path of the device resource. This path can refer only to a character device or block device under **/dev**.

- For **System group owner and file access permissions**, choose **Automatically add file system permissions of the system group that owns the resource**.

The **Group owner file access permission** option lets you grant additional file access permissions to the Lambda process. For more information, see [Group owner file access permission](#).

4. At the bottom of the page, choose **Add resource**.
5. From the **Resources** tab, create another local resource by choosing **Add** and use the following values:
 - For **Resource name**, enter **videoCoreInterface**.
 - For **Resource type**, choose **Device**.
 - For **Local device path**, enter **/dev/vchiq**.
 - For **System group owner and file access permissions**, choose **Automatically add file system permissions of the system group that owns the resource**.
6. Choose **Add resource**.

Step 7: Add subscriptions to the Greengrass group

In this step, add subscriptions to the group. These subscriptions enable the Lambda function to send prediction results to AWS IoT by publishing to an MQTT topic.

1. On the group configuration page, choose the **Subscriptions** tab, and then choose **Add subscription**.
2. On the **Create a subscription** page, configure the source and target, as follows:
 - a. In **Source type**, choose **Lambda function**, and then choose **optimizedImageClassification**.
 - b. In **Target type**, choose **Service**, and then choose **IoT Cloud**.
 - c. In the **Topic filter**, enter **/resnet-50/predictions**, and then choose **Create subscription**.
3. Add a second subscription. Choose the **Subscriptions** tab, choose **Add subscription**, and configure the source and target, as follows:
 - a. In **Source type**, choose **Services**, and then choose **IoT Cloud**.
 - b. In **Target type**, choose **Lambda function**, and then choose **optimizedImageClassification**.
 - c. In the **Topic filter**, enter **/resnet-50/test**, and then choose **Create subscription**.

Step 8: Deploy the Greengrass group

In this step, deploy the current version of the group definition to the Greengrass core device. The definition contains the Lambda function, resources, and subscription configurations that you added.

1. Make sure that the AWS IoT Greengrass core is running. Run the following commands in your Raspberry Pi terminal, as needed:
 - a. To check whether the daemon is running:

```
ps aux | grep -E 'greengrass.*daemon'
```

If the output contains a root entry for `/greengrass/ggc/packages/latest-core-version/bin/daemon`, then the daemon is running.

- b. To start the daemon:

```
cd /greengrass/ggc/core/  
sudo ./greengrassd start
```

2. On the group configuration page, choose **Deploy**.
3. On the **Lambda functions** tab, select **IP detector** and choose **Edit**.
4. From the **Edit IP detector settings** dialog box, select **Automatically detect and override MQTT broker endpoints** and choose **Save**.

This enables devices to automatically acquire connectivity information for the core, such as IP address, DNS, and port number. Automatic detection is recommended, but AWS IoT Greengrass also supports manually specified endpoints. You're only prompted for the discovery method the first time that the group is deployed.

 Note

If prompted, grant permission to create the [Greengrass service role](#) and associate it with your AWS account in the current AWS Region. This role allows AWS IoT Greengrass to access your resources in AWS services.

The **Deployments** page shows the deployment timestamp, version ID, and status. When completed, the status displayed for the deployment should be **Completed**.

For more information about deployments, see [Deploy AWS IoT Greengrass groups](#). For troubleshooting help, see [Troubleshooting](#).

Test the inference example

Now you can verify whether the deployment is configured correctly. To test, you subscribe to the `/resnet-50/predictions` topic and publish any message to the `/resnet-50/test` topic. This triggers the Lambda function to take a photo with your Raspberry Pi and perform inference on the image it captures.

Note

If using the NVIDIA Jetson example, make sure to use the `resnet-18/predictions` and `resnet-18/test` topics instead.

Note

If a monitor is attached to the Raspberry Pi, the live camera feed is displayed in a preview window.

1. On the AWS IoT console home page, under **Test**, choose **MQTT test client**.
2. For **Subscriptions**, choose **Subscribe to a Topic**. Use the following values. Leave the remaining options at their defaults.
 - For **Subscription topic**, enter `/resnet-50/predictions`.
 - Under **Additional configuration**, for **MQTT payload display**, choose **Display payloads as strings**.
3. Choose **Subscribe**.
4. Choose **Publish to a topic**, enter `/resnet-50/test` as the **Topic name**, and choose **Publish**.
5. If the test is successful, the published message causes the Raspberry Pi camera to capture an image. A message from the Lambda function appears at the bottom of the page. This message contains the prediction result of the image, using the format: predicted class name, probability, and peak memory usage.

Configuring an Intel Atom

To run this tutorial on an Intel Atom device, you must provide source images, configure the Lambda function, and add another local device resource. To use the GPU for inference, make sure the following software is installed on your device:

- OpenCL version 1.0 or later
- Python 3.7 and pip
- [NumPy](#)

- [OpenCV on Wheels](#)

1. Download static PNG or JPG images for the Lambda function to use for image classification. The example works best with small image files.

Save your image files in the directory that contains the `inference.py` file (or in a subdirectory of this directory). This is in the Lambda function deployment package that you upload in [the section called “Create an inference Lambda function”](#).

 Note

If you're using AWS DeepLens, you can use the onboard camera or mount your own camera to perform inference on captured images instead of static images. However, we strongly recommend you start with static images first.

If you use a camera, make sure that the `awscam` APT package is installed and up to date. For more information, see [Update your AWS DeepLens device](#) in the *AWS DeepLens Developer Guide*.

2. Edit the configuration of the Lambda function. Follow the procedure in [the section called “Add the Lambda function to the group”](#).

 Note

We recommend that you run your Lambda function without containerization unless your business case requires it. This helps enable access to your device GPU and camera without configuring device resources. If you run without containerization, you must also grant root access to your AWS IoT Greengrass Lambda functions.

a. **To run without containerization:**

- For **System user and group**, choose **Another user ID/group ID**. For **System user ID**, enter `0`. For **System group ID**, enter `0`.

This allows your Lambda function to run as root. For more information about running as root, see [the section called “Setting the default access identity for Lambda functions in a group”](#).

Tip

You also must update your `config.json` file to grant root access to your Lambda function. For the procedure, see [the section called “Running a Lambda function as root”](#).

- For **Lambda function containerization**, choose **No container**.

For more information about running without containerization, see [the section called “Considerations when choosing Lambda function containerization”](#).

- Increase the **Timeout** value to 2 minutes. This ensures that the request does not time out too early. It takes a few minutes after setup to run inference.
- For **Pinned**, choose **True**.
- Under **Additional Parameters**, for **Read access to /sys directory**, choose **Enabled**.

b. **To run in containerized mode instead:**

Note

We do not recommend running in containerized mode unless your business case requires it.

- Increase the **Memory limit** value to 3000 MB.
- Increase the **Timeout** value to 2 minutes. This ensures that the request does not time out too early. It takes a few minutes after setup to run inference.
- For **Pinned**, choose **True**.
- Under **Additional Parameters**, for **Read access to /sys directory**, choose **Enabled**.

3. Add your Neo-optimized model resource to the group. Upload the model resources in the `resnet50` directory of the sample package you unzipped in [the section called “Create an inference Lambda function”](#). This directory contains precompiled model artifacts for an image classification model trained with Resnet-50. Follow the procedure in [the section called “Add a Neo-optimized model resource to the group”](#) with the following updates.

- Compress the files inside the `resnet50` directory into a file named `resnet50.zip`.

- On the **Create a machine learning resource** page, for **Resource name**, enter **resnet50_model**.
 - Upload the `resnet50.zip` file.
4. **If running in containerized mode**, add the required local device resource to grant access to your device GPU.

 Note

If you run in non-containerized mode, AWS IoT Greengrass can access your device GPU without configuring device resources.

- a. On the group configuration page, choose the **Resources** tab.
- b. In the **Local resources** section, choose **Add local resource**.
- c. Define the resource:
 - For **Resource name**, enter **renderD128**.
 - For **Resource type**, choose **Device**.
 - For **Local device path**, enter `/dev/dri/renderD128`.
 - For **System group owner and file access permissions**, choose **Automatically add file system permissions of the system group that owns the resource**.

Configuring an NVIDIA Jetson TX2

To run this tutorial on an NVIDIA Jetson TX2, provide source images, configure the Lambda function, and add more local device resources.

1. Make sure your Jetson device is configured so you can install the AWS IoT Greengrass Core software and use the GPU for inference. For more information about configuring your device, see [the section called “Setting up other devices”](#). To use the GPU for inference on an NVIDIA Jetson TX2, you must install CUDA 10.0 and cuDNN 7.0 on your device when you image your board with Jetpack 4.3.
2. Download static PNG or JPG images for the Lambda function to use for image classification. The example works best with small image files.

Save your image files in the directory that contains the `inference.py` file. You can also save them in a subdirectory of this directory. This directory is in the Lambda function deployment package that you upload in [the section called “Create an inference Lambda function”](#).

Note

You can instead choose to instrument a camera on the Jetson board to capture the source images. However, we strongly recommend you start with static images first.

3. Edit the configuration of the Lambda function. Follow the procedure in [the section called “Add the Lambda function to the group”](#).

Note

We recommend that you run your Lambda function without containerization unless your business case requires it. This helps enable access to your device GPU and camera without configuring device resources. If you run without containerization, you must also grant root access to your AWS IoT Greengrass Lambda functions.

a. **To run without containerization:**

- For **Run as**, choose **Another user ID/group ID**. For **UID**, enter **0**. For **GUID**, enter **0**.

This allows your Lambda function to run as root. For more information about running as root, see [the section called “Setting the default access identity for Lambda functions in a group”](#).

Tip

You also must update your `config.json` file to grant root access to your Lambda function. For the procedure, see [the section called “Running a Lambda function as root”](#).

- For **Lambda function containerization**, choose **No container**.

For more information about running without containerization, see [the section called “Considerations when choosing Lambda function containerization”](#).

- Increase the **Timeout** value to 5 minutes. This ensures that the request does not time out too early. It takes a few minutes after setup to run inference.
- For **Pinned**, choose **True**.
- Under **Additional Parameters**, for **Read access to /sys directory**, choose **Enabled**.

b. **To run in containerized mode instead:**

 Note

We do not recommend running in containerized mode unless your business case requires it.

- Increase the **Memory limit** value. To use the provided model in GPU mode, use at least 2000 MB.
 - Increase the **Timeout** value to 5 minutes. This ensures that the request does not time out too early. It takes a few minutes after setup to run inference.
 - For **Pinned**, choose **True**.
 - Under **Additional Parameters**, for **Read access to /sys directory**, choose **Enabled**.
4. Add your Neo-optimized model resource to the group. Upload the model resources in the `resnet18` directory of the sample package you unzipped in [the section called “Create an inference Lambda function”](#). This directory contains precompiled model artifacts for an image classification model trained with Resnet-18. Follow the procedure in [the section called “Add a Neo-optimized model resource to the group”](#) with the following updates.
- Compress the files inside the `resnet18` directory into a file named `resnet18.zip`.
 - On the **Create a machine learning resource** page, for **Resource name**, enter **resnet18_model**.
 - Upload the `resnet18.zip` file.
5. **If running in containerized mode**, add the required local device resources to grant access to your device GPU.

 Note

If you run in non-containerized mode, AWS IoT Greengrass can access your device GPU without configuring device resources.

- a. On the group configuration page, choose the **Resources** tab.
- b. In the **Local resources** section, choose **Add local resource**.
- c. Define each resource:
 - For **Resource name** and **Device path**, use the values in the following table. Create one device resource for each row in the table.
 - For **Resource type**, choose **Device**.
 - For **System group owner and file access permissions**, choose **Automatically add file system permissions of the system group that owns the resource**.

Name	Device path
nvhost-ctrl	/dev/nvhost-ctrl
nvhost-gpu	/dev/nvhost-gpu
nvhost-ctrl-gpu	/dev/nvhost-ctrl-gpu
nvhost-dbg-gpu	/dev/nvhost-dbg-gpu
nvhost-prof-gpu	/dev/nvhost-prof-gpu
nvmap	/dev/nvmap
nvhost-vic	/dev/nvhost-vic
tegra_dc_ctrl	/dev/tegra_dc_ctrl

6. **If running in containerized mode**, add the following local volume resource to grant access to your device camera. Follow the procedure in [the section called “Add a Neo-optimized model resource to the group”](#).

Note

If you run in non-containerized mode, AWS IoT Greengrass can access your device camera without configuring device resources.

- For **Resource type**, choose **Volume**.
- For **System group owner and file access permissions**, choose **Automatically add file system permissions of the system group that owns the resource**.

Name	Source path	Destination path
shm	/dev/shm	/dev/shm
tmp	/tmp	/tmp

7. Update your group subscriptions to use the correct directory. Follow the procedure in [the section called “Add subscriptions to the group”](#) with the following updates.
 - For your first topic filter, enter **/resnet-18/predictions**.
 - For your second topic filter, enter **/resnet-18/test**.
8. Update your test subscriptions to use the correct directory. Follow the procedure in [the section called “Test the example”](#) with the following updates.
 - For **Subscriptions**, choose **Subscribe to a topic**. For **Subscription topic**, enter **/resnet-18/predictions**.
 - On the **/resnet-18/predictions** page, specify the **/resnet-18/test** topic to publish to.

Troubleshooting AWS IoT Greengrass ML inference

If the test is not successful, you can try the following troubleshooting steps. Run the commands in your Raspberry Pi terminal.

Check error logs

1. Switch to the root user and navigate to the log directory. Access to AWS IoT Greengrass logs requires root permissions.

```
sudo su
cd /greengrass/ggc/var/log
```

2. Check `runtime.log` for any errors.

```
cat system/runtime.log | grep 'ERROR'
```

You can also look in your user-defined Lambda function log for any errors:

```
cat user/your-region/your-account-id/lambda-function-name.log | grep 'ERROR'
```

For more information, see [the section called “Troubleshooting with logs”](#).

Verify the Lambda function is successfully deployed

1. List the contents of the deployed Lambda in the `/lambda` directory. Replace the placeholder values before you run the command.

```
cd /greengrass/ggc/deployment/lambda/
arn:aws:lambda:region:account:function:function-name:function-version
ls -la
```

2. Verify that the directory contains the same content as the `optimizedImageClassification.zip` deployment package that you uploaded in [Step 3: Create an inference Lambda function](#).

Make sure that the `.py` files and dependencies are in the root of the directory.

Verify the inference model is successfully deployed

1. Find the process identification number (PID) of the Lambda runtime process:

```
ps aux | grep lambda-function-name
```

In the output, the PID appears in the second column of the line for the Lambda runtime process.

2. Enter the Lambda runtime namespace. Be sure to replace the placeholder *pid* value before you run the command.

Note

This directory and its contents are in the Lambda runtime namespace, so they aren't visible in a regular Linux namespace.

```
sudo nsenter -t pid -m /bin/bash
```

3. List the contents of the local directory that you specified for the ML resource.

Note

If your ML resource path is something other than `ml_model`, you must substitute that here.

```
cd /ml_model  
ls -ls
```

You should see the following files:

```
56 -rw-r--r-- 1 ggc_user ggc_group 56703 Oct 29 20:07 model.json  
196152 -rw-r--r-- 1 ggc_user ggc_group 200855043 Oct 29 20:08 model.params  
256 -rw-r--r-- 1 ggc_user ggc_group 261848 Oct 29 20:07 model.so  
32 -rw-r--r-- 1 ggc_user ggc_group 30564 Oct 29 20:08 synset.txt
```


Lambda function cannot find `/dev/dri/renderD128`

This can occur if OpenCL cannot connect to the GPU devices it needs. You must create device resources for the necessary devices for your Lambda function.

Next steps

Next, explore other optimized models. For information, see the [SageMaker AI Neo documentation](#).

Manage data streams on the AWS IoT Greengrass core

AWS IoT Greengrass stream manager makes it easier and more reliable to transfer high-volume IoT data to the AWS Cloud. Stream manager processes data streams locally and exports them to the AWS Cloud automatically. This feature integrates with common edge scenarios, such as machine learning (ML) inference, where data is processed and analyzed locally before being exported to the AWS Cloud or local storage destinations.

Stream manager simplifies application development. Your IoT applications can use a standardized mechanism to process high-volume streams and manage local data retention policies instead of building custom stream management functionality. IoT applications can read and write to streams. They can define policies for storage type, size, and data retention on a per-stream basis to control how stream manager processes and exports streams.

Stream manager is designed to work in environments with intermittent or limited connectivity. You can define bandwidth use, timeout behavior, and how stream data is handled when the core is connected or disconnected. For critical data, you can set priorities to control the order in which streams are exported to the AWS Cloud.

You can configure automatic exports to the AWS Cloud for storage or further processing and analysis. Stream manager supports exporting to the following AWS Cloud destinations.


- Channels in AWS IoT Analytics. AWS IoT Analytics lets you perform advanced analysis on your data to help make business decisions and improve machine learning models. For more information, see [What is AWS IoT Analytics?](#) in the *AWS IoT Analytics User Guide*.
- Streams in Kinesis Data Streams. Kinesis Data Streams is commonly used to aggregate high-volume data and load it into a data warehouse or map-reduce cluster. For more information, see [What is Amazon Kinesis Data Streams?](#) in the *Amazon Kinesis Developer Guide*.
- Asset properties in AWS IoT SiteWise. AWS IoT SiteWise lets you collect, organize, and analyze data from industrial equipment at scale. For more information, see [What is AWS IoT SiteWise?](#) in the *AWS IoT SiteWise User Guide*.
- Objects in Amazon S3. You can use Amazon S3 to store and retrieve large amounts of data. For more information, see [What is Amazon S3?](#) in the *Amazon Simple Storage Service Developer Guide*.

Stream management workflow

Your IoT applications interact with stream manager through the AWS IoT Greengrass Core SDK. In a simple workflow, a user-defined Lambda function running on the Greengrass core consumes IoT data, such as time-series temperature and pressure metrics. The Lambda function might filter or compress the data and then call the AWS IoT Greengrass Core SDK to write the data to a stream in stream manager. Stream manager can export the stream to the AWS Cloud automatically, based on the policies defined for the stream. User-defined Lambda functions can also send data directly to local databases or storage repositories.

Your IoT applications can include multiple user-defined Lambda functions that read or write to streams. These local Lambda functions can read and write to streams to filter, aggregate, and analyze data locally. This makes it possible to respond quickly to local events and extract valuable information before the data is transferred from the core to cloud or local destinations.

An example workflow is shown in the following diagram.



To use stream manager, start by configuring stream manager parameters to define group-level runtime settings that apply to all streams on the Greengrass core. These customizable settings allow you to control how stream manager stores, processes, and exports streams based on your business need and environment constraints. For more information, see [the section called “Configure stream manager”](#).

After you configure stream manager, you can create and deploy your IoT applications. These are typically user-defined Lambda functions that use `StreamManagerClient` in the AWS IoT Greengrass Core SDK to create and interact with streams. During stream creation, the Lambda function defines per-stream policies, such as export destinations, priority, and persistence. For more information, including code snippets for `StreamManagerClient` operations, see [the section called “Use StreamManagerClient to work with streams”](#).

For tutorials that configure a simple workflow, see [the section called “Export data streams \(console\)”](#) or [the section called “Export data streams \(CLI\)”](#).

Requirements

The following requirements apply for using stream manager:

- You must use AWS IoT Greengrass Core software v1.10 or later, with stream manager enabled. For more information, see [the section called “Configure stream manager”](#).

Stream manager is not supported on OpenWrt distributions.

- The Java 8 runtime (JDK 8) must be installed on the core.
 - For Debian-based distributions (including Raspbian) or Ubuntu-based distributions, run the following command:

```
sudo apt install openjdk-8-jdk
```

- For Red Hat-based distributions (including Amazon Linux), run the following command:

```
sudo yum install java-1.8.0-openjdk
```

For more information, see [How to download and install prebuilt OpenJDK packages](#) in the OpenJDK documentation.

- Stream manager requires a minimum of 70 MB RAM in addition to your base AWS IoT Greengrass Core software. Your total memory requirement depends on your workload.
- User-defined Lambda functions must use the [AWS IoT Greengrass Core SDK](#) to interact with stream manager. The AWS IoT Greengrass Core SDK is available in several languages, but only the following versions support stream manager operations:
 - Java SDK (v1.4.0 or later)
 - Python SDK (v1.5.0 or later)
 - Node.js SDK (v1.6.0 or later)

Download the version of the SDK that corresponds to your Lambda function runtime and include it in your Lambda function deployment package.

 Note

The AWS IoT Greengrass Core SDK for Python requires Python 3.7 or later and has other package dependencies. For more information, see [Create a Lambda function deployment package \(console\)](#) or [Create a Lambda function deployment package \(CLI\)](#).

- If you define AWS Cloud export destinations for a stream, you must create your export targets and grant access permissions in the Greengrass group role. Depending on the destination, other requirements might also apply. For more information, see:
 - [the section called “AWS IoT Analytics channels”](#)
 - [the section called “Amazon Kinesis data streams”](#)
 - [the section called “AWS IoT SiteWise asset properties”](#)
 - [the section called “Amazon S3 objects”](#)

You are responsible for maintaining these AWS Cloud resources.

Data security

When you use stream manager, be aware of the following security considerations.

Local data security

AWS IoT Greengrass does not encrypt stream data at rest or in transit locally between components on the core device.

- **Data at rest.** Stream data is stored locally in a storage directory on the Greengrass core. For data security, AWS IoT Greengrass relies on Unix file permissions and full-disk encryption, if enabled. You can use the optional [STREAM_MANAGER_STORE_ROOT_DIR](#) parameter to specify the storage directory. If you change this parameter later to use a different storage directory, AWS IoT Greengrass does not delete the previous storage directory or its contents.

- **Data in transit locally.** AWS IoT Greengrass does not encrypt stream data in local transit on the core between data sources, Lambda functions, the AWS IoT Greengrass Core SDK, and stream manager.
- **Data in transit to the AWS Cloud.** Data streams exported by stream manager to the AWS Cloud use standard AWS service client encryption with Transport Layer Security (TLS).

For more information, see [the section called “Data encryption”](#).

Client authentication

Stream manager clients use the AWS IoT Greengrass Core SDK to communicate with stream manager. When client authentication is enabled, only Lambda functions in the Greengrass group can interact with streams in stream manager. When client authentication is disabled, any process running on the Greengrass core (such as [Docker containers](#)) can interact with streams in stream manager. You should disable authentication only if your business case requires it.

You use the [STREAM_MANAGER_AUTHENTICATE_CLIENT](#) parameter to set the client authentication mode. You can configure this parameter from the console or AWS IoT Greengrass API. Changes take effect after the group is deployed.

	Enabled	Disabled
Parameter value	true (default and recommended)	false
Allowed clients	User-defined Lambda functions in the Greengrass group	User-defined Lambda functions in the Greengrass group Other processes running on the Greengrass core device

See also

- [the section called “Configure stream manager”](#)

- [the section called “Use StreamManagerClient to work with streams”](#)
- [the section called “Export configurations for supported AWS Cloud destinations”](#)
- [the section called “Export data streams \(console\)”](#)
- [the section called “Export data streams \(CLI\)”](#)

Configure AWS IoT Greengrass stream manager

On the AWS IoT Greengrass core, stream manager can store, process, and export IoT device data. Stream manager provides parameters that you use to configure group-level runtime settings. These settings apply to all streams on the Greengrass core. You can use the AWS IoT console or AWS IoT Greengrass API to configure stream manager settings. Changes take effect after the group is deployed.

Note

After you configure stream manager, you can create and deploy IoT applications that run on the Greengrass core and interact with stream manager. These IoT applications are typically user-defined Lambda functions. For more information, see [the section called “Use StreamManagerClient to work with streams”](#).

Stream manager parameters

Stream manager provides the following parameters that allow you to define group-level settings. All parameters are optional.

Storage directory

Parameter name: `STREAM_MANAGER_STORE_ROOT_DIR`

The absolute path of the local directory used to store streams. This value must start with a forward slash (for example, `/data`).

For information about securing stream data, see [the section called “Local data security”](#).

Minimum AWS IoT Greengrass Core version: 1.10.0

Server port

Parameter name: `STREAM_MANAGER_SERVER_PORT`

The local port number used to communicate with stream manager. The default is 8088.

Minimum AWS IoT Greengrass Core version: 1.10.0

Authenticate client

Parameter name: `STREAM_MANAGER_AUTHENTICATE_CLIENT`

Indicates whether clients must be authenticated to interact with stream manager. All interaction between clients and stream manager is controlled by the AWS IoT Greengrass Core SDK. This parameter determines which clients can call the AWS IoT Greengrass Core SDK to work with streams. For more information, see [the section called “Client authentication”](#).

Valid values are `true` or `false`. The default is `true` (recommended).

- `true`. Allows only Greengrass Lambda functions as clients. Lambda function clients use internal AWS IoT Greengrass core protocols to authenticate with the AWS IoT Greengrass Core SDK.
- `false`. Allows any process that runs on the AWS IoT Greengrass core to be a client. Do not set to `false` unless your business case requires it. For example, set this value to `false` only if non-Lambda processes on the core device must communicate directly with stream manager, such as [Docker containers](#) running on the core.

Minimum AWS IoT Greengrass Core version: 1.10.0

Maximum bandwidth

Parameter name: `STREAM_MANAGER_EXPORTER_MAX_BANDWIDTH`

The average maximum bandwidth (in kilobits per second) that can be used to export data. The default allows unlimited use of available bandwidth.

Minimum AWS IoT Greengrass Core version: 1.10.0

Thread pool size

Parameter name: `STREAM_MANAGER_EXPORTER_THREAD_POOL_SIZE`

The maximum number of active threads that can be used to export data. The default is 5.

The optimal size depends on your hardware, stream volume, and planned number of export streams. If your export speed is slow, you can adjust this setting to find the optimal size for your hardware and business case. The CPU and memory of your core device hardware are limiting

factors. To start, you might try setting this value equal to the number of processor cores on the device.

Be careful not to set a size that's higher than your hardware can support. Each stream consumes hardware resources, so you should try to limit the number of export streams on constrained devices.

Minimum AWS IoT Greengrass Core version: 1.10.0

JVM arguments

Parameter name: JVM_ARGS

Custom Java Virtual Machine arguments to pass to stream manager at startup. Multiple arguments should be separated by spaces.

Use this parameter only when you must override the default settings used by the JVM. For example, you might need to increase the default heap size if you plan to export a large number of streams.

Minimum AWS IoT Greengrass Core version: 1.10.0

Read-only input file directories

Parameter name: STREAM_MANAGER_READ_ONLY_DIRS

A comma-separated list of absolute paths to the directories outside of the root file system that store input files. Stream manager reads and uploads the files to Amazon S3 and mounts the directories as read-only. For more information about exporting to Amazon S3, see [the section called "Amazon S3 objects"](#).

Use this parameter only if the following conditions are true:

- The input file directory for a stream that exports to Amazon S3 is in one of the following locations:
 - A partition other than the root file system.
 - Under /tmp on the root file system.
- The [default containerization](#) of the Greengrass group is **Greengrass container**.

Example value: /mnt/directory-1,/mnt/directory-2,/tmp

Minimum AWS IoT Greengrass Core version: 1.11.0

Minimum size for multipart upload

Parameter name:

STREAM_MANAGER_EXPORTER_S3_DESTINATION_MULTIPART_UPLOAD_MIN_PART_SIZE_BYTES

The minimum size (in bytes) of a part in a multipart upload to Amazon S3. Stream manager uses this setting and the size of the input file to determine how to batch data in a multipart PUT request. The default and minimum value is 5242880 bytes (5 MB).

Note

Stream manager uses the stream's `sizeThresholdForMultipartUploadBytes` property to determine whether to export to Amazon S3 as a single or multipart upload. User-defined Lambda functions set this threshold when they create a stream that exports to Amazon S3. The default threshold is 5 MB.

Minimum AWS IoT Greengrass Core version: 1.11.0

Configure stream manager settings (console)

You can use the AWS IoT console for the following management tasks:

- [Check if stream manager is enabled](#)
- [Enable or disable stream manager during group creation](#)
- [Enable or disable stream manager for an existing group](#)
- [Change stream manager settings](#)

Changes take effect after the Greengrass group is deployed. For a tutorial that shows how to deploy a Greengrass group that contains a Lambda function that interacts with stream manager, see [the section called “Export data streams \(console\)”](#).

Note

When you use the console to enable stream manager and deploy the group, the memory size for stream manager is set to 4194304 KB (4 GB) by default. We recommend that you set the memory size to at least 128000 KB.

To check if stream manager is enabled (console)

1. In the AWS IoT console navigation pane, under **Manage**, expand **Greengrass devices**, and then choose **Groups (V1)**.
2. Choose the target group.
3. Choose the **Lambda functions** tab.
4. Under **System Lambda functions**, select **Stream manager**, and choose **Edit**.
5. Check the enabled or disabled status. Any custom stream manager settings that are configured are also displayed.

To enable or disable stream manager during group creation (console)

1. In the AWS IoT console navigation pane, under **Manage**, expand **Greengrass devices**, and then choose **Groups (V1)**.
2. Choose **Create Group**. Your choice on the next page determines how you configure stream manager for the group.
3. Proceed through the **Name your Group** and choose a **Greengrass core** pages.
4. Choose **Create group**.
5. On the group configuration page, choose the **Lambda functions** tab, select **Stream manager**, and choose **Edit**.
 - To enable stream manager with default settings, choose **Enable with default settings**.
 - To enable stream manager with custom settings, choose **Customize settings**.
 1. On the **Configure Stream manager** page, choose **Enable with custom settings**.
 2. Under **Custom settings**, enter values for stream manager parameters. For more information, see [the section called "Stream manager parameters"](#). Leave fields empty to allow AWS IoT Greengrass to use their default values.
 - To disable stream manager, choose **Disable**.

1. On the **Configure stream manager** page, choose **Disable**.
6. Choose **Save**.
7. Continue through the remaining pages to create your group.
8. On the **Client devices** page, download your security resources, review the information, and then choose **Finish**.

 Note

When stream manager is enabled, you must [install the Java 8 runtime](#) on the core device before you deploy the group.

To enable or disable stream manager for an existing group (console)

1. In the AWS IoT console navigation pane, under **Manage**, expand **Greengrass devices**, and then choose **Groups (V1)**.
2. Choose the target group.
3. Choose the **Lambda functions** tab.
4. Under **System Lambda functions**, select **Stream manager**, and choose **Edit**.
5. Check the enabled or disabled status. Any custom stream manager settings that are configured are also displayed.

To change stream manager settings (console)

1. In the AWS IoT console navigation pane, under **Manage**, expand **Greengrass devices**, and then choose **Groups (V1)**.
2. Choose the target group.
3. Choose the **Lambda functions** tab.
4. Under **System Lambda functions**, select **Stream manager**, and choose **Edit**.

5. Check the enabled or disabled status. Any custom stream manager settings that are configured are also displayed.
6. Choose **Save**.

Configure stream manager settings (CLI)

In the AWS CLI, use the system `GGStreamManager` Lambda function to configure stream manager. System Lambda functions are components of the AWS IoT Greengrass Core software. For stream manager and some other system Lambda functions, you can configure Greengrass functionality by managing the corresponding `Function` and `FunctionDefinitionVersion` objects in the Greengrass group. For more information, see [the section called “Overview of the group object model”](#).

You can use the API for the following management tasks. The examples in this section show how to use the AWS CLI, but you can also call the AWS IoT Greengrass API directly or use an AWS SDK.

- [Check if stream manager is enabled](#)
- [Enable, disable, or configure stream manager](#)

Changes take effect after the group is deployed. For a tutorial that shows how to deploy a Greengrass group with a Lambda function that interacts with stream manager, see [the section called “Export data streams \(CLI\)”](#).

Tip

To see if stream manager is enabled and running from your core device, you can run the following command in a terminal on the device.

```
ps aux | grep -i 'streammanager'
```

To check if stream manager is enabled (CLI)

Stream manager is enabled if your deployed function definition version includes the system `GGStreamManager` Lambda function. To check, do the following;

1. Get the IDs of the target Greengrass group and group version. This procedure assumes that this is the latest group and group version. The following query returns the most recently created group.

```
aws greengrass list-groups --query "reverse(sort_by(Groups, &CreationTimestamp))
[0]"
```

Or, you can query by name. Group names are not required to be unique, so multiple groups might be returned.

```
aws greengrass list-groups --query "Groups[?Name=='MyGroup']"
```

 Note

You can also find these values in the AWS IoT console. The group ID is displayed on the group's **Settings** page. Group version IDs are displayed on the group's **Deployments** tab.

2. Copy the `Id` and `LatestVersion` values from the target group in the output.
3. Get the latest group version.
 - Replace *group-id* with the `Id` that you copied.
 - Replace *latest-group-version-id* with the `LatestVersion` that you copied.

```
aws greengrass get-group-version \
--group-id group-id \
--group-version-id latest-group-version-id
```

4. From the `FunctionDefinitionVersionArn` in the output, get the IDs of the function definition and function definition version.
 - The function definition ID is the GUID that follows the `functions` segment in the Amazon Resource Name (ARN).
 - The function definition version ID is the GUID that follows the `versions` segment in the ARN.

```
arn:aws:greengrass:us-west-2:123456789012:/greengrass/definition/  
functions/function-definition-id/versions/function-definition-version-id
```

5. Get the function definition version.

- Replace *function-definition-id* with the function definition ID.
- Replace *function-definition-version-id* with the function definition version ID.

```
aws greengrass get-function-definition-version \  
--function-definition-id function-definition-id \  
--function-definition-version-id function-definition-version-id
```

If the functions array in the output includes the GGStreamManager function, then stream manager is enabled. Any environment variables defined for the function represent custom settings for stream manager.

To enable, disable, or configure stream manager (CLI)

In the AWS CLI, use the system GGStreamManager Lambda function to configure stream manager. Changes take effect after you deploy the group.

- To enable stream manager, include GGStreamManager in the functions array of your function definition version. To configure custom settings, define environment variables for the corresponding [stream manager parameters](#).
- To disable stream manager, remove GGStreamManager from the functions array of your function definition version.

Stream manager with default settings

The following example configuration enables stream manager with default settings. It sets the arbitrary function ID to streamManager.

```
{  
  "FunctionArn": "arn:aws:lambda::function:GGStreamManager:1",  
  "FunctionConfiguration": {  
    "MemorySize": 4194304,
```



```
    "Pinned": true,  
    "Timeout": 3  
  },  
  "Id": "streamManager"  
}
```

Note

For the FunctionConfiguration properties, you might know the following:

- MemorySize is set to 4194304 KB (4 GB) with default settings. You can always change this value. We recommend that you set MemorySize to at least 128000 KB.
- Pinned must be set to true.
- Timeout is required by the function definition version, but GGStreamManager doesn't use it.

Stream manager with custom settings

The following example configuration enables stream manager with custom values for the storage directory, server port, and thread pool size parameters.

```
{  
  "FunctionArn": "arn:aws:lambda::function:GGStreamManager:1",  
  "FunctionConfiguration": {  
    "Environment": {  
      "Variables": {  
        "STREAM_MANAGER_STORE_ROOT_DIR": "/data",  
        "STREAM_MANAGER_SERVER_PORT": "1234",  
        "STREAM_MANAGER_EXPORTER_THREAD_POOL_SIZE": "4"  
      }  
    },  
    "MemorySize": 4194304,  
    "Pinned": true,  
    "Timeout": 3  
  },  
  "Id": "streamManager"  
}
```

AWS IoT Greengrass uses default values for [stream manager parameters](#) that aren't specified as environment variables.

Stream manager with custom settings for Amazon S3 exports

The following example configuration enables stream manager with custom values for the upload directory and minimum multipart upload size parameters.

```
{
  "FunctionArn": "arn:aws:lambda::function:GGStreamManager:1",
  "FunctionConfiguration": {
    "Environment": {
      "Variables": {
        "STREAM_MANAGER_READ_ONLY_DIRS": "/mnt/directory-1,/mnt/
directory-2,/tmp",
        "STREAM_MANAGER_EXPORTER_S3_DESTINATION_MULTIPART_UPLOAD_MIN_PART_SIZE_BYTES":
        "10485760"
      }
    },
    "MemorySize": 4194304,
    "Pinned": true,
    "Timeout": 3
  },
  "Id": "streamManager"
}
```

To enable, disable, or configure stream manager (CLI)

1. Get the IDs of the target Greengrass group and group version. This procedure assumes that this is the latest group and group version. The following query returns the most recently created group.

```
aws greengrass list-groups --query "reverse(sort_by(Groups, &CreationTimestamp))
[0]"
```

Or, you can query by name. Group names are not required to be unique, so multiple groups might be returned.

```
aws greengrass list-groups --query "Groups[?Name=='MyGroup']"
```

Note

You can also find these values in the AWS IoT console. The group ID is displayed on the group's **Settings** page. Group version IDs are displayed on the group's **Deployments** tab.

2. Copy the `Id` and `LatestVersion` values from the target group in the output.
3. Get the latest group version.
 - Replace *group-id* with the `Id` that you copied.
 - Replace *latest-group-version-id* with the `LatestVersion` that you copied.

```
aws greengrass get-group-version \  
--group-id group-id \  
--group-version-id latest-group-version-id
```

4. Copy the `CoreDefinitionVersionArn` and all other version ARNs from the output, except `FunctionDefinitionVersionArn`. You use these values later when you create a group version.
5. From the `FunctionDefinitionVersionArn` in the output, copy the ID of the function definition. The ID is the GUID that follows the `functions` segment in the ARN, as shown in the following example.

```
arn:aws:greengrass:us-west-2:123456789012:/greengrass/  
definition/functions/bcfc6b49-beb0-4396-b703-6dEXAMPLEcu5/  
versions/0f7337b4-922b-45c5-856f-1aEXAMPLEsf6
```

Note

Or, you can create a function definition by running the [create-function-definition](#) command, and then copying the ID from the output.

6. Add a function definition version to the function definition.
 - Replace *function-definition-id* with the ID that you copied for the function definition.

- In the functions array, include all other functions that you want to make available on the Greengrass core. You can use the `get-function-definition-version` command to get the list of existing functions.

Enable stream manager with default settings

The following example enables stream manager, by including the `GGStreamManager` function in the functions array. This example uses default values for [stream manager parameters](#).

```
aws greengrass create-function-definition-version \  
--function-definition-id function-definition-id \  
--functions '[  
    {  
        "FunctionArn": "arn:aws:lambda::function:GGStreamManager:1",  
        "FunctionConfiguration": {  
            "MemorySize": 4194304,  
            "Pinned": true,  
            "Timeout": 3  
        },  
        "Id": "streamManager"  
    },  
    {  
        "FunctionArn": "arn:aws:lambda:us-  
west-2:123456789012:function:MyLambdaFunction:MyAlias",  
        "FunctionConfiguration": {  
            "Executable": "myLambdaFunction.function_handler",  
            "MemorySize": 16000,  
            "Pinned": true,  
            "Timeout": 5  
        },  
        "Id": "myLambdaFunction"  
    },  
    ... more user-defined functions  
]  
'
```

Note

The myLambdaFunction function in the examples represents one of your user-defined Lambda functions.

Enable stream manager with custom settings

The following example enables stream manager by including the GGStreamManager function in the functions array. All stream manager settings are optional, unless you want to change the default values. This example shows how to use environment variables to set custom values.

```
aws greengrass create-function-definition-version \
--function-definition-id function-definition-id \
--functions '[
    {
        "FunctionArn": "arn:aws:lambda::function:GGStreamManager:1",
        "FunctionConfiguration": {
            "Environment": {
                "Variables": {
                    "STREAM_MANAGER_STORE_ROOT_DIR": "/data",
                    "STREAM_MANAGER_SERVER_PORT": "1234",
                    "STREAM_MANAGER_EXPORTER_THREAD_POOL_SIZE": "4"
                }
            },
            "MemorySize": 4194304,
            "Pinned": true,
            "Timeout": 3
        },
        "Id": "streamManager"
    },
    {
        "FunctionArn": "arn:aws:lambda:us-
west-2:123456789012:function:MyLambdaFunction:MyAlias",
        "FunctionConfiguration": {
            "Executable": "myLambdaFunction.function_handler",
            "MemorySize": 16000,
            "Pinned": true,
            "Timeout": 5
        }
    },
]
```

```

        "Id": "myLambdaFunction"
      },
      ... more user-defined functions
    ]
  }'

```

Note

For the FunctionConfiguration properties, you might know the following:

- MemorySize is set to 4194304 KB (4 GB) with default settings. You can always change this value. We recommend that you set MemorySize to at least 128000 KB.
- Pinned must be set to true.
- Timeout is required by the function definition version, but GGStreamManager doesn't use it.

Disable stream manager

The following example omits the GGStreamManager function, which disables stream manager.

```

aws greengrass create-function-definition-version \
--function-definition-id function-definition-id \
--functions '[
  {
    "FunctionArn": "arn:aws:lambda:us-
west-2:123456789012:function:MyLambdaFunction:MyAlias",
    "FunctionConfiguration": {
      "Executable": "myLambdaFunction.function_handler",
      "MemorySize": 16000,
      "Pinned": true,
      "Timeout": 5
    },
    "Id": "myLambdaFunction"
  },
  ... more user-defined functions
]'

```

Note

If you don't want to deploy any Lambda functions, you can omit the function definition version entirely.

7. Copy the Arn of the function definition version from the output.
8. Create a group version that contains the system Lambda function.
 - Replace *group-id* with the Id for the group.
 - Replace *core-definition-version-arn* with the CoreDefinitionVersionArn that you copied from the latest group version.
 - Replace *function-definition-version-arn* with the Arn that you copied for the new function definition version.
 - Replace the ARNs for other group components (for example, SubscriptionDefinitionVersionArn or DeviceDefinitionVersionArn) that you copied from the latest group version.
 - Remove any unused parameters. For example, remove the `--resource-definition-version-arn` if your group version doesn't contain any resources.

```
aws greengrass create-group-version \  
--group-id group-id \  
--core-definition-version-arn core-definition-version-arn \  
--function-definition-version-arn function-definition-version-arn \  
--device-definition-version-arn device-definition-version-arn \  
--logger-definition-version-arn logger-definition-version-arn \  
--resource-definition-version-arn resource-definition-version-arn \  
--subscription-definition-version-arn subscription-definition-version-arn
```

9. Copy the Version from the output. This is the ID of the new group version.
10. Deploy the group with the new group version.
 - Replace *group-id* with the Id that you copied for the group.
 - Replace *group-version-id* with the Version that you copied for the new group version.

```
aws greengrass create-deployment \  
--group-id group-id \  
--group-version-id group-version-id
```

```
--group-version-id group-version-id \  
--deployment-type NewDeployment
```

Follow this procedure if you want to edit stream manager settings again later. Make sure to create a function definition version that includes the `GGStreamManager` function with the updated configuration. The group version must reference all component version ARNs that you want to deploy to the core. Changes take effect after the group is deployed.

See also

- [Manage data streams](#)
- [the section called “Use StreamManagerClient to work with streams”](#)
- [the section called “Export configurations for supported AWS Cloud destinations”](#)
- [the section called “Export data streams \(console\)”](#)
- [the section called “Export data streams \(CLI\)”](#)

Use StreamManagerClient to work with streams

User-defined Lambda functions running on the AWS IoT Greengrass core can use the `StreamManagerClient` object in the [AWS IoT Greengrass Core SDK](#) to create streams in [stream manager](#) and then interact with the streams. When a Lambda function creates a stream, it defines the AWS Cloud destinations, prioritization, and other export and data retention policies for the stream. To send data to stream manager, Lambda functions append the data to the stream. If an export destination is defined for the stream, stream manager exports the stream automatically.

Note

Typically, clients of stream manager are user-defined Lambda functions. If your business case requires it, you can also allow non-Lambda processes running on the Greengrass core (for example, a Docker container) to interact with stream manager. For more information, see [the section called “Client authentication”](#).

The snippets in this topic show you how clients call `StreamManagerClient` methods to work with streams. For implementation details about the methods and their arguments, use the links to the SDK reference listed after each snippet. For tutorials that include a complete Python Lambda function, see [the section called “Export data streams \(console\)”](#) or [the section called “Export data streams \(CLI\)”](#).

Your Lambda function should instantiate `StreamManagerClient` outside of the function handler. If instantiated in the handler, the function creates a client and connection to stream manager every time that it's invoked.

Note

If you do instantiate `StreamManagerClient` in the handler, you must explicitly call the `close()` method when the client completes its work. Otherwise, the client keeps the connection open and another thread running until the script exits.

`StreamManagerClient` supports the following operations:

- [the section called “Create message stream”](#)
- [the section called “Append message”](#)
- [the section called “Read messages”](#)
- [the section called “List streams”](#)
- [the section called “Describe message stream”](#)
- [the section called “Update message stream”](#)
- [the section called “Delete message stream”](#)

Create message stream

To create a stream, a user-defined Lambda function calls the `create` method and passes in a `MessageStreamDefinition` object. This object specifies the unique name for the stream and defines how stream manager should handle new data when the maximum stream size is reached. You can use `MessageStreamDefinition` and its data types (such as `ExportDefinition`, `StrategyOnFull`, and `Persistence`) to define other stream properties. These include:

- The target AWS IoT Analytics, Kinesis Data Streams, AWS IoT SiteWise, and Amazon S3 destinations for automatic exports. For more information, see [the section called “Export configurations for supported AWS Cloud destinations”](#).
- Export priority. Stream manager exports higher priority streams before lower priority streams.
- Maximum batch size and batch interval for AWS IoT Analytics, Kinesis Data Streams, and AWS IoT SiteWise destinations. Stream manager exports messages when either condition is met.
- Time-to-live (TTL). The amount of time to guarantee that the stream data is available for processing. You should make sure that the data can be consumed within this time period. This is not a deletion policy. The data might not be deleted immediately after TTL period.
- Stream persistence. Choose to save streams to the file system to persist data across core restarts or save streams in memory.
- Starting sequence number. Specify the sequence number of the message to use as the starting message in the export.

For more information about `MessageStreamDefinition`, see the SDK reference for your target language:

- [MessageStreamDefinition](#) in the Java SDK
- [MessageStreamDefinition](#) in the Node.js SDK
- [MessageStreamDefinition](#) in the Python SDK

Note

`StreamManagerClient` also provides a target destination you can use to export streams to an HTTP server. This target is intended for testing purposes only. It is not stable or supported for use in production environments.

After a stream is created, your Lambda functions can [append messages](#) to the stream to send data for export and [read messages](#) from the stream for local processing. The number of streams that you create depends on your hardware capabilities and business case. One strategy is to create a stream for each target channel in AWS IoT Analytics or Kinesis data stream, though you can define multiple targets for a stream. A stream has a durable lifespan.

Requirements

This operation has the following requirements:

- Minimum AWS IoT Greengrass Core version: 1.10.0
- Minimum AWS IoT Greengrass Core SDK version: Python: 1.5.0 | Java: 1.4.0 | Node.js: 1.6.0

Note

Creating streams with an AWS IoT SiteWise or Amazon S3 export destination has the following requirements:

- Minimum AWS IoT Greengrass Core version: 1.11.0
- Minimum AWS IoT Greengrass Core SDK version: Python: 1.6.0 | Java: 1.5.0 | Node.js: 1.7.0

Examples

The following snippet creates a stream named `StreamName`. It defines stream properties in the `MessageStreamDefinition` and subordinate data types.

Python

```
client = StreamManagerClient()

try:
    client.create_message_stream(MessageStreamDefinition(
        name="StreamName", # Required.
        max_size=268435456, # Default is 256 MB.
        stream_segment_size=16777216, # Default is 16 MB.
        time_to_live_millis=None, # By default, no TTL is enabled.
        strategy_on_full=StrategyOnFull.OverwriteOldestData, # Required.
        persistence=Persistence.File, # Default is File.
        flush_on_write=False, # Default is false.
        export_definition=ExportDefinition( # Optional. Choose where/how the stream
            is exported to the AWS Cloud.
            kinesis=None,
            iot_analytics=None,
            iot_sitewise=None,
```

```

        s3_task_executor=None
    )
))
except StreamManagerException:
    pass
    # Properly handle errors.
except ConnectionError or asyncio.TimeoutError:
    pass
    # Properly handle errors.

```

Python SDK reference: [create_message_stream](#) | [MessageStreamDefinition](#)

Java

```

try (final StreamManagerClient client =
    GreengrassClientBuilder.streamManagerClient().build()) {
    client.createMessageStream(
        new MessageStreamDefinition()
            .withName("StreamName") // Required.
            .withMaxSize(268435456L) // Default is 256 MB.
            .withStreamSegmentSize(16777216L) // Default is 16 MB.
            .withTimeToLiveMillis(null) // By default, no TTL is enabled.
            .withStrategyOnFull(StrategyOnFull.OverwriteOldestData) //
Required.

            .withPersistence(Persistence.File) // Default is File.
            .withFlushOnWrite(false) // Default is false.
            .withExportDefinition( // Optional. Choose where/how the stream
is exported to the AWS Cloud.
                new ExportDefinition()
                    .withKinesis(null)
                    .withIotAnalytics(null)
                    .withIotSitewise(null)
                    .withS3TaskExecutor(null)
                )
            );
} catch (StreamManagerException e) {
    // Properly handle exception.
}

```

Java SDK reference: [createMessageStream](#) | [MessageStreamDefinition](#)

Node.js

```
const client = new StreamManagerClient();
client.onConnected(async () => {
  try {
    await client.createMessageStream(
      new MessageStreamDefinition()
        .WithName("StreamName") // Required.
        .withMaxSize(268435456) // Default is 256 MB.
        .withStreamSegmentSize(16777216) // Default is 16 MB.
        .withTimeToLiveMillis(null) // By default, no TTL is enabled.
        .withStrategyOnFull(StrategyOnFull.OverwriteOldestData) //
Required.
        .withPersistence(Persistence.File) // Default is File.
        .withFlushOnWrite(false) // Default is false.
        .withExportDefinition( // Optional. Choose where/how the stream is
exported to the AWS Cloud.
          new ExportDefinition()
            .withKinesis(null)
            .withIotAnalytics(null)
            .withIotSitewise(null)
            .withS3TaskExecutor(null)
          )
        );
  } catch (e) {
    // Properly handle errors.
  }
});
client.onError((err) => {
  // Properly handle connection errors.
  // This is called only when the connection to the StreamManager server fails.
});
```

Node.js SDK reference: [createMessageStream](#) | [MessageStreamDefinition](#)

For more information about configuring export destinations, see [the section called “Export configurations for supported AWS Cloud destinations”](#).

Append message

To send data to stream manager for export, your Lambda functions append the data to the target stream. The export destination determines the data type to pass to this method.

Requirements

This operation has the following requirements:

- Minimum AWS IoT Greengrass Core version: 1.10.0
- Minimum AWS IoT Greengrass Core SDK version: Python: 1.5.0 | Java: 1.4.0 | Node.js: 1.6.0

Note

Appending messages with an AWS IoT SiteWise or Amazon S3 export destination has the following requirements:

- Minimum AWS IoT Greengrass Core version: 1.11.0
- Minimum AWS IoT Greengrass Core SDK version: Python: 1.6.0 | Java: 1.5.0 | Node.js: 1.7.0

Examples

AWS IoT Analytics or Kinesis Data Streams export destinations

The following snippet appends a message to the stream named `StreamName`. For AWS IoT Analytics or Kinesis Data Streams destinations, your Lambda functions append a blob of data.

This snippet has the following requirements:

- Minimum AWS IoT Greengrass Core version: 1.10.0
- Minimum AWS IoT Greengrass Core SDK version: Python: 1.5.0 | Java: 1.4.0 | Node.js: 1.6.0

Python

```
client = StreamManagerClient()
```

```
try:
    sequence_number = client.append_message(stream_name="StreamName",
    data=b'Arbitrary bytes data')
except StreamManagerException:
    pass
    # Properly handle errors.
except ConnectionError or asyncio.TimeoutError:
    pass
    # Properly handle errors.
```

Python SDK reference: [append_message](#)

Java

```
try (final StreamManagerClient client =
    GreengrassClientBuilder.streamManagerClient().build()) {
    long sequenceNumber = client.appendMessage("StreamName", "Arbitrary byte
    array".getBytes());
} catch (StreamManagerException e) {
    // Properly handle exception.
}
```

Java SDK reference: [appendMessage](#)

Node.js

```
const client = new StreamManagerClient();
client.onConnected(async () => {
    try {
        const sequenceNumber = await client.appendMessage("StreamName",
        Buffer.from("Arbitrary byte array"));
    } catch (e) {
        // Properly handle errors.
    }
});
client.onError((err) => {
    // Properly handle connection errors.
    // This is called only when the connection to the StreamManager server fails.
});
```

Node.js SDK reference: [appendMessage](#)

AWS IoT SiteWise export destinations

The following snippet appends a message to the stream named `StreamName`. For AWS IoT SiteWise destinations, your Lambda functions append a serialized `PutAssetPropertyValueEntry` object. For more information, see [the section called “Exporting to AWS IoT SiteWise”](#).

Note

When you send data to AWS IoT SiteWise, your data must meet the requirements of the `BatchPutAssetPropertyValue` action. For more information, see [BatchPutAssetPropertyValue](#) in the *AWS IoT SiteWise API Reference*.

This snippet has the following requirements:

- Minimum AWS IoT Greengrass Core version: 1.11.0
- Minimum AWS IoT Greengrass Core SDK version: Python: 1.6.0 | Java: 1.5.0 | Node.js: 1.7.0

Python

```
client = StreamManagerClient()

try:
    # SiteWise requires unique timestamps in all messages. Add some randomness to
    # time and offset.

    # Note: To create a new asset property data, you should use the classes defined
    # in the
    # greengrasssdk.stream_manager module.

    time_in_nanos = TimeInNanos(
        time_in_seconds=calendar.timegm(time.gmtime()) - random.randint(0, 60),
        offset_in_nanos=random.randint(0, 10000)
    )
    variant = Variant(double_value=random.random())
    asset = [AssetPropertyValue(value=variant, quality=Quality.GOOD,
        timestamp=time_in_nanos)]
    putAssetPropertyValueEntry =
    PutAssetPropertyValueEntry(entry_id=str(uuid.uuid4()),
        property_alias="PropertyAlias", property_values=asset)
```



```

        sequence_number = client.append_message(stream_name="StreamName",
        data=Util.validate_and_serialize_to_json_bytes(putAssetPropertyValueEntry))
except StreamManagerException:
    pass
    # Properly handle errors.
except ConnectionError or asyncio.TimeoutError:
    pass
    # Properly handle errors.

```

Python SDK reference: [append_message](#) | [PutAssetPropertyValueEntry](#)

Java

```

try (final StreamManagerClient client =
GreengrassClientBuilder.streamManagerClient().build()) {
    Random rand = new Random();
    // Note: To create a new asset property data, you should use the classes defined
in the
    // com.amazonaws.greengrass.streammanager.model.sitewise package.
    List<AssetPropertyValue> entries = new ArrayList<>() ;

    // IoTSiteWise requires unique timestamps in all messages. Add some randomness
to time and offset.
    final int maxTimeRandomness = 60;
    final int maxOffsetRandomness = 10000;
    double randomValue = rand.nextDouble();
    TimeInNanos timestamp = new TimeInNanos()
        .withTimeInSeconds(Instant.now().getEpochSecond() -
rand.nextInt(maxTimeRandomness))
        .withOffsetInNanos((long) (rand.nextInt(maxOffsetRandomness)));
    AssetPropertyValue entry = new AssetPropertyValue()
        .withValue(new Variant().withDoubleValue(randomValue))
        .withQuality(Quality.GOOD)
        .withTimestamp(timestamp);
    entries.add(entry);

    PutAssetPropertyValueEntry putAssetPropertyValueEntry = new
PutAssetPropertyValueEntry()
        .withEntryId(UUID.randomUUID().toString())
        .withPropertyAlias("PropertyAlias")
        .withPropertyValues(entries);
    long sequenceNumber = client.appendMessage("StreamName",
ValidateAndSerialize.validateAndSerializeToJsonBytes(putAssetPropertyValueEntry));
} catch (StreamManagerException e) {

```

```
// Properly handle exception.
}
```

Java SDK reference: [appendMessage](#) | [PutAssetPropertyValueEntry](#)

Node.js

```
const client = new StreamManagerClient();
client.onConnected(async () => {
  try {
    const maxTimeRandomness = 60;
    const maxOffsetRandomness = 10000;
    const randomValue = Math.random();
    // Note: To create a new asset property data, you should use the classes
    defined in the
    // aws-greengrass-core-sdk StreamManager module.
    const timestamp = new TimeInNanos()
      .withTimeInSeconds(Math.round(Date.now() / 1000) -
        Math.floor(Math.random() * maxTimeRandomness))
      .withOffsetInNanos(Math.floor(Math.random() * maxOffsetRandomness));
    const entry = new AssetPropertyValue()
      .withValue(new Variant().withDoubleValue(randomValue))
      .withQuality(Quality.GOOD)
      .withTimestamp(timestamp);

    const putAssetPropertyValueEntry = new PutAssetPropertyValueEntry()
      .withEntryId(`${ENTRY_ID_PREFIX}${i}`)
      .withPropertyAlias("PropertyAlias")
      .withPropertyValues([entry]);
    const sequenceNumber = await client.appendMessage("StreamName",
      util.validateAndSerializeToJsonBytes(putAssetPropertyValueEntry));
  } catch (e) {
    // Properly handle errors.
  }
});
client.onError((err) => {
  // Properly handle connection errors.
  // This is called only when the connection to the StreamManager server fails.
});
```

Node.js SDK reference: [appendMessage](#) | [PutAssetPropertyValueEntry](#)

Amazon S3 export destinations

The following snippet appends an export task to the stream named `StreamName`. For Amazon S3 destinations, your Lambda functions append a serialized `S3ExportTaskDefinition` object that contains information about the source input file and target Amazon S3 object. If the specified object doesn't exist, Stream Manager creates it for you. For more information, see [the section called "Exporting to Amazon S3"](#).

This snippet has the following requirements:

- Minimum AWS IoT Greengrass Core version: 1.11.0
- Minimum AWS IoT Greengrass Core SDK version: Python: 1.6.0 | Java: 1.5.0 | Node.js: 1.7.0

Python

```
client = StreamManagerClient()

try:
    # Append an Amazon S3 Task definition and print the sequence number.
    s3_export_task_definition = S3ExportTaskDefinition(input_url="URLToFile",
    bucket="BucketName", key="KeyName")
    sequence_number = client.append_message(stream_name="StreamName",
    data=Util.validate_and_serialize_to_json_bytes(s3_export_task_definition))
except StreamManagerException:
    pass
    # Properly handle errors.
except ConnectionError or asyncio.TimeoutError:
    pass
    # Properly handle errors.
```

Python SDK reference: [append_message](#) | [S3ExportTaskDefinition](#)

Java

```
try (final StreamManagerClient client =
GreengrassClientBuilder.streamManagerClient().build()) {
    // Append an Amazon S3 export task definition and print the sequence number.
    S3ExportTaskDefinition s3ExportTaskDefinition = new S3ExportTaskDefinition()
        .withBucket("BucketName")
        .withKey("KeyName")
        .withInputUrl("URLToFile");
```

```
    long sequenceNumber = client.appendMessage("StreamName",
        ValidateAndSerialize.validateAndSerializeToJsonBytes(s3ExportTaskDefinition));
} catch (StreamManagerException e) {
    // Properly handle exception.
}
```

Java SDK reference: [appendMessage](#) | [S3ExportTaskDefinition](#)

Node.js

```
const client = new StreamManagerClient();
client.onConnected(async () => {
    try {
        // Append an Amazon S3 export task definition and print the sequence number.
        const taskDefinition = new S3ExportTaskDefinition()
            .withBucket("BucketName")
            .withKey("KeyName")
            .withInputUrl("URLToFile");
        const sequenceNumber = await client.appendMessage("StreamName",
            util.validateAndSerializeToJsonBytes(taskDefinition));
    } catch (e) {
        // Properly handle errors.
    }
});
client.onError((err) => {
    // Properly handle connection errors.
    // This is called only when the connection to the StreamManager server fails.
});
```

Node.js SDK reference: [appendMessage](#) | [S3ExportTaskDefinition](#)

Read messages

Read messages from a stream.

Requirements

This operation has the following requirements:

- Minimum AWS IoT Greengrass Core version: 1.10.0

- Minimum AWS IoT Greengrass Core SDK version: Python: 1.5.0 | Java: 1.4.0 | Node.js: 1.6.0

Examples

The following snippet reads messages from the stream named `StreamName`. The `read` method takes an optional `ReadMessagesOptions` object that specifies the sequence number to start reading from, the minimum and maximum numbers to read, and a timeout for reading messages.

Python

```
client = StreamManagerClient()

try:
    message_list = client.read_messages(
        stream_name="StreamName",
        # By default, if no options are specified, it tries to read one message from
        the beginning of the stream.
        options=ReadMessagesOptions(
            desired_start_sequence_number=100,
            # Try to read from sequence number 100 or greater. By default, this is
            0.
            min_message_count=10,
            # Try to read 10 messages. If 10 messages are not available, then
            NotEnoughMessagesException is raised. By default, this is 1.
            max_message_count=100, # Accept up to 100 messages. By default this is
            1.
            read_timeout_millis=5000
            # Try to wait at most 5 seconds for the min_message_count to be
            fulfilled. By default, this is 0, which immediately returns the messages or an
            exception.
        )
    )
except StreamManagerException:
    pass
    # Properly handle errors.
except ConnectionError or asyncio.TimeoutError:
    pass
    # Properly handle errors.
```

Python SDK reference: [read_messages](#) | [ReadMessagesOptions](#)

Java

```
try (final StreamManagerClient client =
    GreengrassClientBuilder.streamManagerClient().build()) {
    List<Message> messages = client.readMessages("StreamName",
        // By default, if no options are specified, it tries to read one message
        from the beginning of the stream.
        new ReadMessagesOptions()
            // Try to read from sequence number 100 or greater. By default
            this is 0.
            .withDesiredStartSequenceNumber(100L)
            // Try to read 10 messages. If 10 messages are not available,
            then NotEnoughMessagesException is raised. By default, this is 1.
            .withMinMessageCount(10L)
            // Accept up to 100 messages. By default this is 1.
            .withMaxMessageCount(100L)
            // Try to wait at most 5 seconds for the min_message_count to
            be fulfilled. By default, this is 0, which immediately returns the messages or an
            exception.
            .withReadTimeoutMillis(Duration.ofSeconds(5L).toMillis()));
} catch (StreamManagerException e) {
    // Properly handle exception.
}
```

Java SDK reference: [readMessages](#) | [ReadMessagesOptions](#)

Node.js

```
const client = new StreamManagerClient();
client.onConnected(async () => {
    try {
        const messages = await client.readMessages("StreamName",
            // By default, if no options are specified, it tries to read one message
            from the beginning of the stream.
            new ReadMessagesOptions()
                // Try to read from sequence number 100 or greater. By default this
            is 0.
                .withDesiredStartSequenceNumber(100)
                // Try to read 10 messages. If 10 messages are not available, then
            NotEnoughMessagesException is thrown. By default, this is 1.
                .withMinMessageCount(10)
                // Accept up to 100 messages. By default this is 1.
                .withMaxMessageCount(100)
```

```
        // Try to wait at most 5 seconds for the minMessageCount to be
        fulfilled. By default, this is 0, which immediately returns the messages or an
        exception.
        .withReadTimeoutMillis(5 * 1000)
    );
} catch (e) {
    // Properly handle errors.
}
});
client.onError((err) => {
    // Properly handle connection errors.
    // This is called only when the connection to the StreamManager server fails.
});
```

Node.js SDK reference: [readMessages](#) | [ReadMessagesOptions](#)

List streams

Get the list of streams in stream manager.

Requirements

This operation has the following requirements:

- Minimum AWS IoT Greengrass Core version: 1.10.0
- Minimum AWS IoT Greengrass Core SDK version: Python: 1.5.0 | Java: 1.4.0 | Node.js: 1.6.0

Examples

The following snippet gets a list of the streams (by name) in stream manager.

Python

```
client = StreamManagerClient()

try:
    stream_names = client.list_streams()
except StreamManagerException:
```

```
pass
# Properly handle errors.
except ConnectionError or asyncio.TimeoutError:
    pass
# Properly handle errors.
```

Python SDK reference: [list_streams](#)

Java

```
try (final StreamManagerClient client =
    GreengrassClientBuilder.streamManagerClient().build()) {
    List<String> streamNames = client.listStreams();
} catch (StreamManagerException e) {
    // Properly handle exception.
}
```

Java SDK reference: [listStreams](#)

Node.js

```
const client = new StreamManagerClient();
client.onConnected(async () => {
    try {
        const streams = await client.listStreams();
    } catch (e) {
        // Properly handle errors.
    }
});
client.onError((err) => {
    // Properly handle connection errors.
    // This is called only when the connection to the StreamManager server fails.
});
```

Node.js SDK reference: [listStreams](#)

Describe message stream

Get metadata about a stream, including the stream definition, size, and export status.

Requirements

This operation has the following requirements:

- Minimum AWS IoT Greengrass Core version: 1.10.0
- Minimum AWS IoT Greengrass Core SDK version: Python: 1.5.0 | Java: 1.4.0 | Node.js: 1.6.0

Examples

The following snippet gets metadata about the stream named `StreamName`, including the stream's definition, size, and exporter statuses.

Python

```
client = StreamManagerClient()

try:
    stream_description = client.describe_message_stream(stream_name="StreamName")
    if stream_description.export_statuses[0].error_message:
        # The last export of export destination 0 failed with some error
        # Here is the last sequence number that was successfully exported
        stream_description.export_statuses[0].last_exported_sequence_number

    if (stream_description.storage_status.newest_sequence_number >
        stream_description.export_statuses[0].last_exported_sequence_number):
        pass
        # The end of the stream is ahead of the last exported sequence number
except StreamManagerException:
    pass
    # Properly handle errors.
except ConnectionError or asyncio.TimeoutError:
    pass
    # Properly handle errors.
```

Python SDK reference: [describe_message_stream](#)

Java

```
try (final StreamManagerClient client =
    GreengrassClientBuilder.streamManagerClient().build()) {
    MessageStreamInfo description = client.describeMessageStream("StreamName");
```

```

String lastErrorMessage =
description.getExportStatuses().get(0).getErrorMessage();
if (lastErrorMessage != null && !lastErrorMessage.equals("")) {
    // The last export of export destination 0 failed with some error.
    // Here is the last sequence number that was successfully exported.
    description.getExportStatuses().get(0).getLastExportedSequenceNumber();
}

if (description.getStorageStatus().getNewestSequenceNumber() >
    description.getExportStatuses().get(0).getLastExportedSequenceNumber())
{
    // The end of the stream is ahead of the last exported sequence number.
}
} catch (StreamManagerException e) {
    // Properly handle exception.
}

```

Java SDK reference: [describeMessageStream](#)

Node.js

```

const client = new StreamManagerClient();
client.onConnected(async () => {
    try {
        const description = await client.describeMessageStream("StreamName");
        const lastErrorMessage = description.exportStatuses[0].errorMessage;
        if (lastErrorMessage) {
            // The last export of export destination 0 failed with some error.
            // Here is the last sequence number that was successfully exported.
            description.exportStatuses[0].lastExportedSequenceNumber;
        }

        if (description.storageStatus.newestSequenceNumber >
            description.exportStatuses[0].lastExportedSequenceNumber) {
            // The end of the stream is ahead of the last exported sequence number.
        }
    } catch (e) {
        // Properly handle errors.
    }
});
client.onError((err) => {
    // Properly handle connection errors.
    // This is called only when the connection to the StreamManager server fails.

```

```
});
```

Node.js SDK reference: [describeMessageStream](#)

Update message stream

Update properties of an existing stream. You might want to update a stream if your requirements change after the stream was created. For example:

- Add a new [export configuration](#) for an AWS Cloud destination.
- Increase the maximum size of a stream to change how data is exported or retained. For example, the stream size in combination with your strategy on full settings might result in data being deleted or rejected before stream manager can process it.
- Pause and resume exports; for example, if export tasks are long running and you want to ration your upload data.

Your Lambda functions follow this high-level process to update a stream:

1. [Get the description of the stream.](#)
2. Update the target properties on the corresponding `MessageStreamDefinition` and subordinate objects.
3. Pass in the updated `MessageStreamDefinition`. Make sure to include the complete object definitions for the updated stream. Undefined properties revert to the default values.

You can specify the sequence number of the message to use as the starting message in the export.

Requirements

This operation has the following requirements:

- Minimum AWS IoT Greengrass Core version: 1.11.0
- Minimum AWS IoT Greengrass Core SDK version: Python: 1.6.0 | Java: 1.5.0 | Node.js: 1.7.0

Examples

The following snippet updates the stream named `StreamName`. It updates multiple properties of a stream that exports to Kinesis Data Streams.

Python

```
client = StreamManagerClient()

try:
    message_stream_info = client.describe_message_stream(STREAM_NAME)
    message_stream_info.definition.max_size=536870912
    message_stream_info.definition.stream_segment_size=33554432
    message_stream_info.definition.time_to_live_millis=3600000
    message_stream_info.definition.strategy_on_full=StrategyOnFull.RejectNewData
    message_stream_info.definition.persistence=Persistence.Memory
    message_stream_info.definition.flush_on_write=False
    message_stream_info.definition.export_definition.kinesis=
        [KinesisConfig(
            # Updating Export definition to add a Kinesis Stream configuration.
            identifier=str(uuid.uuid4()), kinesis_stream_name=str(uuid.uuid4()))]
    client.update_message_stream(message_stream_info.definition)
except StreamManagerException:
    pass
    # Properly handle errors.
except ConnectionError or asyncio.TimeoutError:
    pass
    # Properly handle errors.
```

Python SDK reference: [updateMessageStream](#) | [MessageStreamDefinition](#)

Java

```
try (final StreamManagerClient client =
    GreengrassClientBuilder.streamManagerClient().build()) {
    MessageStreamInfo messageStreamInfo = client.describeMessageStream(STREAM_NAME);
    // Update the message stream with new values.
    client.updateMessageStream(
        messageStreamInfo.getDefinition()
            .withStrategyOnFull(StrategyOnFull.RejectNewData) // Required. Updating
            Strategy on full to reject new data.
            // Max Size update should be greater than initial Max Size defined in
            Create Message Stream request
```

```

        .withMaxSize(536870912L) // Update Max Size to 512 MB.
        .withStreamSegmentSize(33554432L) // Update Segment Size to 32 MB.
        .withFlushOnWrite(true) // Update flush on write to true.
        .withPersistence(Persistence.Memory) // Update the persistence to
Memory.

        .withTimeToLiveMillis(3600000L) // Update TTL to 1 hour.
        .withExportDefinition(
            // Optional. Choose where/how the stream is exported to the AWS
Cloud.
            messageStreamInfo.getDefinition().getExportDefinition().
            // Updating Export definition to add a Kinesis Stream
configuration.

            .withKinesis(new ArrayList<KinesisConfig>() {{
                add(new KinesisConfig()
                    .withIdentifier(EXPORT_IDENTIFIER)
                    .withKinesisStreamName("test"));
            }}))
        );
    } catch (StreamManagerException e) {
        // Properly handle exception.
    }

```

Java SDK reference: [update_message_stream](#) | [MessageStreamDefinition](#)

Node.js

```

const client = new StreamManagerClient();
client.onConnected(async () => {
    try {
        const messageStreamInfo = await c.describeMessageStream(STREAM_NAME);
        await client.updateMessageStream(
            messageStreamInfo.definition
            // Max Size update should be greater than initial Max Size defined
in Create Message Stream request
            .withMaxSize(536870912) // Default is 256 MB. Updating Max Size to
512 MB.

            .withStreamSegmentSize(33554432) // Default is 16 MB. Updating
Segment Size to 32 MB.

            .withTimeToLiveMillis(3600000) // By default, no TTL is enabled.
Update TTL to 1 hour.

            .withStrategyOnFull(StrategyOnFull.RejectNewData) // Required.
Updating Strategy on full to reject new data.

            .withPersistence(Persistence.Memory) // Default is File. Update the
persistence to Memory

```

```
        .withFlushOnWrite(true) // Default is false. Updating to true.
        .withExportDefinition(
            // Optional. Choose where/how the stream is exported to the AWS
Cloud.
            messageStreamInfo.definition.exportDefinition
            // Updating Export definition to add a Kinesis Stream
configuration.
            .withKinesis([new
KinesisConfig().withIdentifier(uuidv4()).withKinesisStreamName(uuidv4())])
        )
    );
} catch (e) {
    // Properly handle errors.
}
});
client.onError((err) => {
    // Properly handle connection errors.
    // This is called only when the connection to the StreamManager server fails.
});
```

Node.js SDK reference: [updateMessageStream](#) | [MessageStreamDefinition](#)

Constraints for updating streams

The following constraints apply when updating streams. Unless noted in the following list, updates take effect immediately.

- You can't update a stream's persistence. To change this behavior, [delete the stream](#) and [create a stream](#) that defines the new persistence policy.
- You can update the maximum size of a stream only under the following conditions:
 - The maximum size must be greater or equal to the current size of the stream. To find this information, [describe the stream](#) and then check the storage status of the returned `MessageStreamInfo` object.
 - The maximum size must be greater than or equal to the stream's segment size.
- You can update the stream segment size to a value less than the maximum size of the stream. The updated setting applies to new segments.
- Updates to the time to live (TTL) property apply to new append operations. If you decrease this value, stream manager might also delete existing segments that exceed the TTL.

- Updates to the strategy on full property apply to new append operations. If you set the strategy to overwrite the oldest data, stream manager might also overwrite existing segments based on the new setting.
- Updates to the flush on write property apply to new messages.
- Updates to export configurations apply to new exports. The update request must include all export configurations that you want to support. Otherwise, stream manager deletes them.
 - When you update an export configuration, specify the identifier of the target export configuration.
 - To add an export configuration, specify a unique identifier for the new export configuration.
 - To delete an export configuration, omit the export configuration.
- To [update](#) the starting sequence number of an export configuration in a stream, you must specify a value that's less than the latest sequence number. To find this information, [describe the stream](#) and then check the storage status of the returned `MessageStreamInfo` object.

Delete message stream

Deletes a stream. When you delete a stream, all of the stored data for the stream is deleted from the disk.

Requirements

This operation has the following requirements:

- Minimum AWS IoT Greengrass Core version: 1.10.0
- Minimum AWS IoT Greengrass Core SDK version: Python: 1.5.0 | Java: 1.4.0 | Node.js: 1.6.0

Examples

The following snippet deletes the stream named `StreamName`.

Python

```
client = StreamManagerClient()

try:
```

```
client.delete_message_stream(stream_name="StreamName")
except StreamManagerException:
    pass
    # Properly handle errors.
except ConnectionError or asyncio.TimeoutError:
    pass
    # Properly handle errors.
```

Python SDK reference: [deleteMessageStream](#)

Java

```
try (final StreamManagerClient client =
    GreengrassClientBuilder.streamManagerClient().build()) {
    client.deleteMessageStream("StreamName");
} catch (StreamManagerException e) {
    // Properly handle exception.
}
```

Java SDK reference: [delete_message_stream](#)

Node.js

```
const client = new StreamManagerClient();
client.onConnected(async () => {
    try {
        await client.deleteMessageStream("StreamName");
    } catch (e) {
        // Properly handle errors.
    }
});
client.onError((err) => {
    // Properly handle connection errors.
    // This is called only when the connection to the StreamManager server fails.
});
```

Node.js SDK reference: [deleteMessageStream](#)


See also

- [Manage data streams](#)
- [the section called "Configure stream manager"](#)

- [the section called “Export configurations for supported AWS Cloud destinations”](#)
- [the section called “Export data streams \(console\)”](#)
- [the section called “Export data streams \(CLI\)”](#)
- StreamManagerClient in the AWS IoT Greengrass Core SDK reference:
 - [Python](#)
 - [Java](#)
 - [Node.js](#)

Export configurations for supported AWS Cloud destinations

User-defined Lambda functions use `StreamManagerClient` in the AWS IoT Greengrass Core SDK to interact with stream manager. When a Lambda function [creates a stream](#) or [updates a stream](#), it passes a `MessageStreamDefinition` object that represents stream properties, including the export definition. The `ExportDefinition` object contains the export configurations defined for the stream. Stream manager uses these export configurations to determine where and how to export the stream.



You can define zero or more export configurations on a stream, including multiple export configurations for a single destination type. For example, you can export a stream to two AWS IoT Analytics channels and one Kinesis data stream.

For failed export attempts, stream manager continually retries exporting data to the AWS Cloud at intervals of up to five minutes. The number of retry attempts doesn't have a maximum limit.

Note

`StreamManagerClient` also provides a target destination you can use to export streams to an HTTP server. This target is intended for testing purposes only. It is not stable or supported for use in production environments.

Supported AWS Cloud destinations

- [AWS IoT Analytics channels](#)
- [Amazon Kinesis data streams](#)
- [AWS IoT SiteWise asset properties](#)
- [Amazon S3 objects](#)

You are responsible for maintaining these AWS Cloud resources.

AWS IoT Analytics channels

Stream manager supports automatic exports to AWS IoT Analytics. AWS IoT Analytics lets you perform advanced analysis on your data to help make business decisions and improve machine learning models. For more information, see [What is AWS IoT Analytics?](#) in the *AWS IoT Analytics User Guide*.

In the AWS IoT Greengrass Core SDK, your Lambda functions use the `IoTAnalyticsConfig` to define the export configuration for this destination type. For more information, see the SDK reference for your target language:

- [IoTAnalyticsConfig](#) in the Python SDK
- [IoTAnalyticsConfig](#) in the Java SDK
- [IoTAnalyticsConfig](#) in the Node.js SDK

Requirements

This export destination has the following requirements:

- Target channels in AWS IoT Analytics must be in the same AWS account and AWS Region as the Greengrass group.
- The [the section called “Greengrass group role”](#) must allow the `iotanalytics:BatchPutMessage` permission to target channels. For example:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "iotanalytics:BatchPutMessage"
      ],
      "Resource": [
        "arn:aws:iotanalytics:region:account-id:channel/channel_1_name",
        "arn:aws:iotanalytics:region:account-id:channel/channel_2_name"
      ]
    }
  ]
}
```

You can grant granular or conditional access to resources, for example, by using a wildcard * naming scheme. For more information, see [Adding and removing IAM policies](#) in the *IAM User Guide*.

Exporting to AWS IoT Analytics

To create a stream that exports to AWS IoT Analytics, your Lambda functions [create a stream](#) with an export definition that includes one or more `IoTAnalyticsConfig` objects. This object defines export settings, such as the target channel, batch size, batch interval, and priority.

When your Lambda functions receive data from devices, they [append messages](#) that contain a blob of data to the target stream.

Then, stream manager exports the data based on the batch settings and priority defined in the stream's export configurations.

Amazon Kinesis data streams

Stream manager supports automatic exports to Amazon Kinesis Data Streams. Kinesis Data Streams is commonly used to aggregate high-volume data and load it into a data warehouse or map-reduce cluster. For more information, see [What is Amazon Kinesis Data Streams?](#) in the *Amazon Kinesis Developer Guide*.

In the AWS IoT Greengrass Core SDK, your Lambda functions use the `KinesisConfig` to define the export configuration for this destination type. For more information, see the SDK reference for your target language:

- [KinesisConfig](#) in the Python SDK
- [KinesisConfig](#) in the Java SDK
- [KinesisConfig](#) in the Node.js SDK

Requirements

This export destination has the following requirements:

- Target streams in Kinesis Data Streams must be in the same AWS account and AWS Region as the Greengrass group.
- The [the section called “Greengrass group role”](#) must allow the `kinesis:PutRecords` permission to target data streams. For example:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "kinesis:PutRecords"
      ],
      "Resource": [
        "arn:aws:kinesis:region:account-id:stream/stream_1_name",
        "arn:aws:kinesis:region:account-id:stream/stream_2_name"
      ]
    }
  ]
}
```

You can grant granular or conditional access to resources, for example, by using a wildcard * naming scheme. For more information, see [Adding and removing IAM policies](#) in the *IAM User Guide*.

Exporting to Kinesis Data Streams

To create a stream that exports to Kinesis Data Streams, your Lambda functions [create a stream](#) with an export definition that includes one or more `KinesisConfig` objects. This object defines export settings, such as the target data stream, batch size, batch interval, and priority.

When your Lambda functions receive data from devices, they [append messages](#) that contain a blob of data to the target stream. Then, stream manager exports the data based on the batch settings and priority defined in the stream's export configurations.

Stream manager generates a unique, random UUID as a partition key for each record uploaded to Amazon Kinesis.

AWS IoT SiteWise asset properties

Stream manager supports automatic exports to AWS IoT SiteWise. AWS IoT SiteWise lets you collect, organize, and analyze data from industrial equipment at scale. For more information, see [What is AWS IoT SiteWise?](#) in the *AWS IoT SiteWise User Guide*.

In the AWS IoT Greengrass Core SDK, your Lambda functions use the `IoTSiteWiseConfig` to define the export configuration for this destination type. For more information, see the SDK reference for your target language:

- [IoTSiteWiseConfig](#) in the Python SDK
- [IoTSiteWiseConfig](#) in the Java SDK
- [IoTSiteWiseConfig](#) in the Node.js SDK

Note

AWS also provides the [the section called "IoT SiteWise"](#), which is a pre-built solution that you can use with OPC-UA sources.

Requirements

This export destination has the following requirements:

- Target asset properties in AWS IoT SiteWise must be in the same AWS account and AWS Region as the Greengrass group.

Note

For the list of Regions that AWS IoT SiteWise supports, see [AWS IoT SiteWise endpoints and quotas](#) in the *AWS General Reference*.

- The [the section called “Greengrass group role”](#) must allow the `iotsitewise:BatchPutAssetPropertyValue` permission to target asset properties. The following example policy uses the `iotsitewise:assetHierarchyPath` condition key to grant access to a target root asset and its children. You can remove the Condition from the policy to allow access to all of your AWS IoT SiteWise assets or specify ARNs of individual assets.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": "iotsitewise:BatchPutAssetPropertyValue",
      "Resource": "*",
      "Condition": {
        "StringLike": {
          "iotsitewise:assetHierarchyPath": [
            "/root node asset ID",
            "/root node asset ID/*"
          ]
        }
      }
    }
  ]
}
```

You can grant granular or conditional access to resources, for example, by using a wildcard * naming scheme. For more information, see [Adding and removing IAM policies](#) in the *IAM User Guide*.

For important security information, see [BatchPutAssetPropertyValue authorization](#) in the *AWS IoT SiteWise User Guide*.

Exporting to AWS IoT SiteWise

To create a stream that exports to AWS IoT SiteWise, your Lambda functions [create a stream](#) with an export definition that includes one or more `IoTSiteWiseConfig` objects. This object defines export settings, such as the batch size, batch interval, and priority.

When your Lambda functions receive asset property data from devices, they append messages that contain the data to the target stream. Messages are JSON-serialized `PutAssetPropertyValueEntry` objects that contain property values for one or more asset properties. For more information, see [Append message](#) for AWS IoT SiteWise export destinations.

Note

When you send data to AWS IoT SiteWise, your data must meet the requirements of the `BatchPutAssetPropertyValue` action. For more information, see [BatchPutAssetPropertyValue](#) in the *AWS IoT SiteWise API Reference*.

Then, stream manager exports the data based on the batch settings and priority defined in the stream's export configurations.

You can adjust your stream manager settings and Lambda function logic to design your export strategy. For example:

- For near real time exports, set low batch size and interval settings and append the data to the stream when it's received.
- To optimize batching, mitigate bandwidth constraints, or minimize cost, your Lambda functions can pool the timestamp-quality-value (TQV) data points received for a single asset property before appending the data to the stream. One strategy is to batch entries for up to 10 different property-asset combinations, or property aliases, in one message instead of sending more than one entry for the same property. This helps stream manager to remain within [AWS IoT SiteWise quotas](#).

Amazon S3 objects

Stream manager supports automatic exports to Amazon S3. You can use Amazon S3 to store and retrieve large amounts of data. For more information, see [What is Amazon S3?](#) in the *Amazon Simple Storage Service Developer Guide*.

In the AWS IoT Greengrass Core SDK, your Lambda functions use the `S3ExportTaskExecutorConfig` to define the export configuration for this destination type. For more information, see the SDK reference for your target language:

- [S3ExportTaskExecutorConfig](#) in the Python SDK
- [S3ExportTaskExecutorConfig](#) in the Java SDK
- [S3ExportTaskExecutorConfig](#) in the Node.js SDK

Requirements

This export destination has the following requirements:

- Target Amazon S3 buckets must be in the same AWS account as the Greengrass group.
- If the [default containerization](#) for the Greengrass group is **Greengrass container**, you must set the [STREAM_MANAGER_READ_ONLY_DIRS](#) parameter to use an input file directory that's under `/tmp` or isn't on the root file system.
- If a Lambda function running in **Greengrass container** mode writes input files to the input file directory, you must create a local volume resource for the directory and mount the directory to the container with write permissions. This ensures that the files are written to the root file system and visible outside the container. For more information, see [Access local resources](#).
- The [the section called "Greengrass group role"](#) must allow the following permissions to the target buckets. For example:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "s3:PutObject",
```



```
        "s3:AbortMultipartUpload",
        "s3:ListMultipartUploadParts"
    ],
    "Resource": [
        "arn:aws:s3:::bucket-1-name/*",
        "arn:aws:s3:::bucket-2-name/*"
    ]
}
]
```

You can grant granular or conditional access to resources, for example, by using a wildcard * naming scheme. For more information, see [Adding and removing IAM policies](#) in the *IAM User Guide*.


Exporting to Amazon S3

To create a stream that exports to Amazon S3, your Lambda functions use the `S3ExportTaskExecutorConfig` object to configure the export policy. The policy defines export settings, such as the multipart upload threshold and priority. For Amazon S3 exports, stream manager uploads data that it reads from local files on the core device. To initiate an upload, your Lambda functions append an export task to the target stream. The export task contains information about the input file and target Amazon S3 object. Stream manager executes tasks in the sequence that they are appended to the stream.

Note

The target bucket must already exist in your AWS account. If an object for the specified key doesn't exist, stream manager creates the object for you.

This high-level workflow is shown in the following diagram.



Stream manager uses the multipart upload threshold property, [minimum part size](#) setting, and size of the input file to determine how to upload data. The multipart upload threshold must be greater or equal to the minimum part size. If you want to upload data in parallel, you can create multiple streams.

The keys that specify your target Amazon S3 objects can include valid [Java DateTimeFormatter](#) strings in `!{timestamp:value}` placeholders. You can use these timestamp placeholders to partition data in Amazon S3 based on the time that the input file data was uploaded. For example, the following key name resolves to a value such as `my-key/2020/12/31/data.txt`.

```
my-key/!{timestamp:YYYY}/!{timestamp:MM}/!{timestamp:dd}/data.txt
```

Note

If you want to monitor the export status for a stream, first create a status stream and then configure the export stream to use it. For more information, see [the section called “Monitor export tasks”](#).

Manage input data

You can author code that IoT applications use to manage the lifecycle of the input data. The following example workflow shows how you might use Lambda functions to manage this data.

1. A local process receives data from devices or peripherals, and then writes the data to files in a directory on the core device. These are the input files for stream manager.

Note

To determine if you must configure access to the input file directory, see the [STREAM_MANAGER_READ_ONLY_DIRS](#) parameter.

The process that stream manager runs in inherits all of the file system permissions of the [default access identity](#) for the group. Stream manager must have permission to access the input files. You can use the `chmod(1)` command to change the permission of the files, if necessary.

2. A Lambda function scans the directory and [appends an export task](#) to the target stream when a new file is created. The task is a JSON-serialized `S3ExportTaskDefinition` object that specifies the URL of the input file, the target Amazon S3 bucket and key, and optional user metadata.
3. Stream manager reads the input file and exports the data to Amazon S3 in the order of appended tasks. The target bucket must already exist in your AWS account. If an object for the specified key doesn't exist, stream manager creates the object for you.
4. The Lambda function [reads messages](#) from a status stream to monitor the export status. After export tasks are completed, the Lambda function can delete the corresponding input files. For more information, see [the section called "Monitor export tasks"](#).

Monitor export tasks

You can author code that IoT applications use to monitor the status of your Amazon S3 exports. Your Lambda functions must create a status stream and then configure the export stream to write status updates to the status stream. A single status stream can receive status updates from multiple streams that export to Amazon S3.

First, [create a stream](#) to use as the status stream. You can configure the size and retention policies for the stream to control the lifespan of the status messages. For example:

- Set `Persistence to Memory` if you don't want to store the status messages.
- Set `StrategyOnFull` to `OverwriteOldestData` so that new status messages are not lost.

Then, create or update the export stream to use the status stream. Specifically, set the status configuration property of the stream's `S3ExportTaskExecutorConfig` export configuration. This tells stream manager to write status messages about the export tasks to the status stream. In the `StatusConfig` object, specify the name of the status stream and the level of verbosity. The following supported values range from least verbose (`ERROR`) to most verbose (`TRACE`). The default is `INFO`.

- `ERROR`
- `WARN`
- `INFO`
- `DEBUG`
- `TRACE`

The following example workflow shows how Lambda functions might use a status stream to monitor export status.

1. As described in the previous workflow, a Lambda function [appends an export task](#) to a stream that's configured to write status messages about export tasks to a status stream. The append operation return a sequence number that represents the task ID.
2. A Lambda function [reads messages](#) sequentially from the status stream, and then filters the messages based on the stream name and task ID or based on an export task property from the message context. For example, the Lambda function can filter by the input file URL of the export task, which is represented by the `S3ExportTaskDefinition` object in the message context.

The following status codes indicate that an export task has reached a completed state:

- **Success.** The upload was completed successfully.
- **Failure.** Stream manager encountered an error, for example, the specified bucket does not exist. After resolving the issue, you can append the export task to the stream again.
- **Canceled.** The task was aborted because the stream or export definition was deleted, or the time-to-live (TTL) period of the task expired.

Note

The task might also have a status of `InProgress` or `Warning`. Stream manager issues warnings when an event returns an error that doesn't affect the execution of the task. For example, a failure to clean up an aborted partial upload returns a warning.

3. After export tasks are completed, the Lambda function can delete the corresponding input files.

The following example shows how a Lambda function might read and process status messages.

Python

```
import time
from greengrasssdk.stream_manager import (
    ReadMessagesOptions,
    Status,
    StatusConfig,
    StatusLevel,
    StatusMessage,
    StreamManagerClient,
)
from greengrasssdk.stream_manager.util import Util

client = StreamManagerClient()

try:
    # Read the statuses from the export status stream
    is_file_uploaded_to_s3 = False
    while not is_file_uploaded_to_s3:
        try:
            messages_list = client.read_messages(
                "StatusStreamName", ReadMessagesOptions(min_message_count=1,
read_timeout_millis=1000)
            )
            for message in messages_list:
                # Deserialize the status message first.
                status_message = Util.deserialize_json_bytes_to_obj(message.payload,
StatusMessage)

                # Check the status of the status message. If the status is
                "Success",
```

```

        # the file was successfully uploaded to S3.
        # If the status was either "Failure" or "Cancelled", the server was
        # unable to upload the file to S3.
        # We will print the message for why the upload to S3 failed from the
        # status message.
        # If the status was "InProgress", the status indicates that the
        # server has started uploading
        # the S3 task.
        if status_message.status == Status.Success:
            logger.info("Successfully uploaded file at path " + file_url + "
to S3.")

            is_file_uploaded_to_s3 = True
        elif status_message.status == Status.Failure or
status_message.status == Status.Canceled:
            logger.info(
                "Unable to upload file at path " + file_url + " to S3.
Message: " + status_message.message
            )
            is_file_uploaded_to_s3 = True
            time.sleep(5)
        except StreamManagerException:
            logger.exception("Exception while running")
except StreamManagerException:
    pass
    # Properly handle errors.
except ConnectionError or asyncio.TimeoutError:
    pass
    # Properly handle errors.

```

Python SDK reference: [read_messages](#) | [StatusMessage](#)

Java

```

import com.amazonaws.greengrass.streammanager.client.StreamManagerClient;
import com.amazonaws.greengrass.streammanager.client.utils.ValidateAndSerialize;
import com.amazonaws.greengrass.streammanager.model.ReadMessagesOptions;
import com.amazonaws.greengrass.streammanager.model.Status;
import com.amazonaws.greengrass.streammanager.model.StatusConfig;
import com.amazonaws.greengrass.streammanager.model.StatusLevel;
import com.amazonaws.greengrass.streammanager.model.StatusMessage;

try (final StreamManagerClient client =
    GreengrassClientBuilder.streamManagerClient().build()) {
    try {

```

```

        boolean isS3UploadComplete = false;
        while (!isS3UploadComplete) {
            try {
                // Read the statuses from the export status stream
                List<Message> messages = client.readMessages("StatusStreamName",
                    new
ReadMessagesOptions().withMinMessageCount(1L).withReadTimeoutMillis(1000L));
                for (Message message : messages) {
                    // Deserialize the status message first.
                    StatusMessage statusMessage =
ValidateAndSerialize.deserializeJsonBytesToObj(message.getPayload(),
StatusMessage.class);
                    // Check the status of the status message. If the status is
                    "Success", the file was successfully uploaded to S3.
                    // If the status was either "Failure" or "Canceled", the server
                    was unable to upload the file to S3.
                    // We will print the message for why the upload to S3 failed
                    from the status message.
                    // If the status was "InProgress", the status indicates that the
                    server has started uploading the S3 task.
                    if (Status.Success.equals(statusMessage.getStatus())) {
                        System.out.println("Successfully uploaded file at path " +
FILE_URL + " to S3.");
                        isS3UploadComplete = true;
                    } else if (Status.Failure.equals(statusMessage.getStatus()) ||
Status.Canceled.equals(statusMessage.getStatus())) {
                        System.out.println(String.format("Unable to upload file at
path %s to S3. Message %s",
statusMessage.getStatusContext().getS3ExportTaskDefinition().getInputUrl(),
statusMessage.getMessage()));
                        sS3UploadComplete = true;
                    }
                }
            } catch (StreamManagerException ignored) {
            } finally {
                // Sleep for sometime for the S3 upload task to complete before
                trying to read the status message.
                Thread.sleep(5000);
            }
        } catch (e) {
            // Properly handle errors.
        }
    } catch (StreamManagerException e) {

```

```
// Properly handle exception.
}
```

Java SDK reference: [readMessages](#) | [StatusMessage](#)

Node.js

```
const {
  StreamManagerClient, ReadMessagesOptions,
  Status, StatusConfig, StatusLevel, StatusMessage,
  util,
} = require('aws-greengrass-core-sdk').StreamManager;

const client = new StreamManagerClient();
client.onConnected(async () => {
  try {
    let isS3UploadComplete = false;
    while (!isS3UploadComplete) {
      try {
        // Read the statuses from the export status stream
        const messages = await c.readMessages("StatusStreamName",
          new ReadMessagesOptions()
            .withMinMessageCount(1)
            .withReadTimeoutMillis(1000));

        messages.forEach((message) => {
          // Deserialize the status message first.
          const statusMessage =
            util.deserializeJsonBytesToObj(message.payload, StatusMessage);
          // Check the status of the status message. If the status is
          'Success', the file was successfully uploaded to S3.
          // If the status was either 'Failure' or 'Cancelled', the server
          was unable to upload the file to S3.
          // We will print the message for why the upload to S3 failed
          from the status message.
          // If the status was "InProgress", the status indicates that the
          server has started uploading the S3 task.
          if (statusMessage.status === Status.Success) {
            console.log(`Successfully uploaded file at path ${FILE_URL}
            to S3.`);

            isS3UploadComplete = true;
          } else if (statusMessage.status === Status.Failure ||
            statusMessage.status === Status.Canceled) {
```




```
        console.log(`Unable to upload file at path ${FILE_URL} to
S3. Message: ${statusMessage.message}`);
        isS3UploadComplete = true;
    }
    });
    // Sleep for sometime for the S3 upload task to complete before
    trying to read the status message.
    await new Promise((r) => setTimeout(r, 5000));
    } catch (e) {
        // Ignored
    }
    } catch (e) {
        // Properly handle errors.
    }
    });
    client.onError((err) => {
        // Properly handle connection errors.
        // This is called only when the connection to the StreamManager server fails.
    });
```

Node.js SDK reference: [readMessages](#) | [StatusMessage](#)

Export data streams to the AWS Cloud (console)

This tutorial shows you how to use the AWS IoT console to configure and deploy an AWS IoT Greengrass group with stream manager enabled. The group contains a user-defined Lambda function that writes to a stream in stream manager, which is then exported automatically to the AWS Cloud.

Stream manager makes ingesting, processing, and exporting high-volume data streams more efficient and reliable. In this tutorial, you create a `TransferStream` Lambda function that consumes IoT data. The Lambda function uses the AWS IoT Greengrass Core SDK to create a stream in stream manager and then read and write to it. Stream manager then exports the stream to Kinesis Data Streams. The following diagram shows this workflow.



The focus of this tutorial is to show how user-defined Lambda functions use the `StreamManagerClient` object in the AWS IoT Greengrass Core SDK to interact with stream manager. For simplicity, the Python Lambda function that you create for this tutorial generates simulated device data.

Prerequisites

To complete this tutorial, you need:

- A Greengrass group and a Greengrass core (v1.10 or later). For information about how to create a Greengrass group and core, see [Getting started with AWS IoT Greengrass](#). The Getting Started tutorial also includes steps for installing the AWS IoT Greengrass Core software.

Note

Stream manager is not supported on OpenWrt distributions.

- The Java 8 runtime (JDK 8) installed on the core device.
 - For Debian-based distributions (including Raspbian) or Ubuntu-based distributions, run the following command:

```
sudo apt install openjdk-8-jdk
```

- For Red Hat-based distributions (including Amazon Linux), run the following command:

```
sudo yum install java-1.8.0-openjdk
```

For more information, see [How to download and install prebuilt OpenJDK packages](#) in the OpenJDK documentation.

- AWS IoT Greengrass Core SDK for Python v1.5.0 or later. To use `StreamManagerClient` in the AWS IoT Greengrass Core SDK for Python, you must:
 - Install Python 3.7 or later on the core device.
 - Include the SDK and its dependencies in your Lambda function deployment package. Instructions are provided in this tutorial.

Tip

You can use `StreamManagerClient` with Java or NodeJS. For example code, see the [AWS IoT Greengrass Core SDK for Java](#) and [AWS IoT Greengrass Core SDK for Node.js](#) on GitHub.

- A destination stream named **MyKinesisStream** created in Amazon Kinesis Data Streams in the same AWS Region as your Greengrass group. For more information, see [Create a stream](#) in the *Amazon Kinesis Developer Guide*.

Note

In this tutorial, stream manager exports data to Kinesis Data Streams, which results in charges to your AWS account. For information about pricing, see [Kinesis Data Streams pricing](#).

To avoid incurring charges, you can run this tutorial without creating a Kinesis data stream. In this case, you check the logs to see that stream manager attempted to export the stream to Kinesis Data Streams.

- An IAM policy added to the [the section called “Greengrass group role”](#) that allows the `kinesis:PutRecords` action on the target data stream, as shown in the following example:

```
{
  "Version": "2012-10-17",
  "Statement": [
```

```
{
  "Effect": "Allow",
  "Action": [
    "kinesis:PutRecords"
  ],
  "Resource": [
    "arn:aws:kinesis:region:account-id:stream/MyKinesisStream"
  ]
}
```

The tutorial contains the following high-level steps:

1. [Create a Lambda function deployment package](#)
2. [Create a Lambda function](#)
3. [Add a function to the group](#)
4. [Enable stream manager](#)
5. [Configure local logging](#)
6. [Deploy the group](#)
7. [Test the application](#)

The tutorial should take about 20 minutes to complete.

Step 1: Create a Lambda function deployment package

In this step, you create a Lambda function deployment package that contains Python function code and dependencies. You upload this package later when you create the Lambda function in AWS Lambda. The Lambda function uses the AWS IoT Greengrass Core SDK to create and interact with local streams.

Note

Your user-defined Lambda functions must use the [AWS IoT Greengrass Core SDK](#) to interact with stream manager. For more information about requirements for the Greengrass stream manager, see [Greengrass stream manager requirements](#).

1. Download the [AWS IoT Greengrass Core SDK for Python](#) v1.5.0 or later.
2. Unzip the downloaded package to get the SDK. The SDK is the greengrasssdk folder.
3. Install package dependencies to include with the SDK in your Lambda function deployment package.
 1. Navigate to the SDK directory that contains the requirements.txt file. This file lists the dependencies.
 2. Install the SDK dependencies. For example, run the following pip command to install them in the current directory:

```
pip install --target . -r requirements.txt
```

4. Save the following Python code function in a local file named transfer_stream.py.

 Tip

For example code that uses Java and NodeJS, see the [AWS IoT Greengrass Core SDK for Java](#) and [AWS IoT Greengrass Core SDK for Node.js](#) on GitHub.

```
import asyncio
import logging
import random
import time

from greengrasssdk.stream_manager import (
    ExportDefinition,
    KinesisConfig,
    MessageStreamDefinition,
    ReadMessagesOptions,
    ResourceNotFoundException,
    StrategyOnFull,
    StreamManagerClient,
)

# This example creates a local stream named "SomeStream".
# It starts writing data into that stream and then stream manager automatically
# exports
# the data to a customer-created Kinesis data stream named "MyKinesisStream".
```

```
# This example runs forever until the program is stopped.

# The size of the local stream on disk will not exceed the default (which is 256
  MB).
# Any data appended after the stream reaches the size limit continues to be
  appended, and
# stream manager deletes the oldest data until the total stream size is back under
  256 MB.
# The Kinesis data stream in the cloud has no such bound, so all the data from this
  script is
# uploaded to Kinesis and you will be charged for that usage.

def main(logger):
    try:
        stream_name = "SomeStream"
        kinesis_stream_name = "MyKinesisStream"

        # Create a client for the StreamManager
        client = StreamManagerClient()

        # Try deleting the stream (if it exists) so that we have a fresh start
        try:
            client.delete_message_stream(stream_name=stream_name)
        except ResourceNotFoundException:
            pass

        exports = ExportDefinition(
            kinesis=[KinesisConfig(identifier="KinesisExport" + stream_name,
            kinesis_stream_name=kinesis_stream_name)]
        )
        client.create_message_stream(
            MessageStreamDefinition(
                name=stream_name,
            strategy_on_full=StrategyOnFull.OverwriteOldestData, export_definition=exports
            )
        )

        # Append two messages and print their sequence numbers
        logger.info(
            "Successfully appended message to stream with sequence number %d",
            client.append_message(stream_name, "ABCDEFGHJKLMNOP".encode("utf-8")),
        )
        logger.info(
```

```

        "Successfully appended message to stream with sequence number %d",
        client.append_message(stream_name, "PQRSTUVWXYZ".encode("utf-8")),
    )

    # Try reading the two messages we just appended and print them out
    logger.info(
        "Successfully read 2 messages: %s",
        client.read_messages(stream_name,
ReadMessagesOptions(min_message_count=2, read_timeout_millis=1000)),
    )

    logger.info("Now going to start writing random integers between 0 and 1000
to the stream")
    # Now start putting in random data between 0 and 1000 to emulate device
sensor input
    while True:
        logger.debug("Appending new random integer to stream")
        client.append_message(stream_name, random.randint(0,
1000).to_bytes(length=4, signed=True, byteorder="big"))
        time.sleep(1)

    except asyncio.TimeoutError:
        logger.exception("Timed out while executing")
    except Exception:
        logger.exception("Exception while running")

def function_handler(event, context):
    return

logging.basicConfig(level=logging.INFO)
# Start up this sample code
main(logger=logging.getLogger())

```

5. Zip the following items into a file named `transfer_stream_python.zip`. This is your Lambda function deployment package.

- **transfer_stream.py**. App logic.
- **greengrasssdk**. Required library for Python Greengrass Lambda functions that publish MQTT messages.

[Stream manager operations](#) are available in version 1.5.0 or later of the AWS IoT Greengrass Core SDK for Python.


- The dependencies you installed for the AWS IoT Greengrass Core SDK for Python (for example, the `cbor2` directories).

When you create the zip file, include only these items, not the containing folder.

Step 2: Create a Lambda function

In this step, you use the AWS Lambda console to create a Lambda function and configure it to use your deployment package. Then, you publish a function version and create an alias.

1. First, create the Lambda function.
 - a. In the AWS Management Console, choose **Services**, and open the AWS Lambda console.
 - b. Choose **Create function** and then choose **Author from scratch**.
 - c. In the **Basic information** section, use the following values:
 - For **Function name**, enter **TransferStream**.
 - For **Runtime**, choose **Python 3.7**.
 - For **Permissions**, keep the default setting. This creates an execution role that grants basic Lambda permissions. This role isn't used by AWS IoT Greengrass.
 - d. At the bottom of the page, choose **Create function**.
2. Next, register the handler and upload your Lambda function deployment package.
 - a. On the **Code** tab, under **Code source**, choose **Upload from**. From the dropdown, choose **.zip file**.



- b. Choose **Upload**, and then choose your `transfer_stream_python.zip` deployment package. Then, choose **Save**.
- c. On the **Code** tab for the function, under **Runtime settings**, choose **Edit**, and then enter the following values.
 - For **Runtime**, choose **Python 3.7**.
 - For **Handler**, enter **`transfer_stream.function_handler`**
- d. Choose **Save**.

 Note

The **Test** button on the AWS Lambda console doesn't work with this function. The AWS IoT Greengrass Core SDK doesn't contain modules that are required to run your Greengrass Lambda functions independently in the AWS Lambda console. These modules (for example, `greengrass_common`) are supplied to the functions after they are deployed to your Greengrass core.

3. Now, publish the first version of your Lambda function and create an [alias for the version](#).

 Note

Greengrass groups can reference a Lambda function by alias (recommended) or by version. Using an alias makes it easier to manage code updates because you don't have to change your subscription table or group definition when the function code is updated. Instead, you just point the alias to the new function version.

- a. From the **Actions** menu, choose **Publish new version**.
- b. For **Version description**, enter **First version**, and then choose **Publish**.
- c. On the **TransferStream: 1** configuration page, from the **Actions** menu, choose **Create alias**.
- d. On the **Create a new alias** page, use the following values:
 - For **Name**, enter **`GG_TransferStream`**.
 - For **Version**, choose **1**.

Note

AWS IoT Greengrass doesn't support Lambda aliases for **\$LATEST** versions.

- e. Choose **Create**.

Now you're ready to add the Lambda function to your Greengrass group.

Step 3: Add a Lambda function to the Greengrass group

In this step, you add the Lambda function to the group and then configure its lifecycle and environment variables. For more information, see [the section called "Controlling Greengrass Lambda function execution"](#).

1. In the AWS IoT console navigation pane, under **Manage**, expand **Greengrass devices**, and then choose **Groups (V1)**.
2. Choose the target group.
3. On the group configuration page, choose the **Lambda functions** tab.
4. Under **My Lambda functions**, choose **Add**.
5. On the **Add Lambda function** page, choose the **Lambda function** for your Lambda function.
6. For the **Lambda version**, choose **Alias:GG_TransferStream**.

Now, configure properties that determine the behavior of the Lambda function in the Greengrass group.

7. In the **Lambda function configuration** section, make the following changes:
 - Set **Memory limit** to 32 MB.
 - For **Pinned**, choose **True**.

Note

A *long-lived* (or *pinned*) Lambda function starts automatically after AWS IoT Greengrass starts and keeps running in its own container. This is in contrast to an *on-demand*

Lambda function, which starts when invoked and stops when there are no tasks left to run. For more information, see [the section called "Lifecycle configuration"](#).

8. Choose **Add Lambda function**.

Step 4: Enable stream manager

In this step, you make sure that stream manager is enabled.

1. On the group configuration page, choose the **Lambda functions** tab.
2. Under **System Lambda functions**, select **Stream manager**, and check the status. If disabled, choose **Edit**. Then, choose **Enable** and **Save**. You can use the default parameter settings for this tutorial. For more information, see [the section called "Configure stream manager"](#).

Note

When you use the console to enable stream manager and deploy the group, the memory size for stream manager is set to 4194304 KB (4 GB) by default. We recommend that you set the memory size to at least 128000 KB.

Step 5: Configure local logging

In this step, you configure AWS IoT Greengrass system components, user-defined Lambda functions, and connectors in the group to write logs to the file system of the core device. You can use logs to troubleshoot any issues you might encounter. For more information, see [the section called "Monitoring with AWS IoT Greengrass logs"](#).

1. Under **Local logs configuration**, check if local logging is configured.
2. If logs aren't configured for Greengrass system components or user-defined Lambda functions, choose **Edit**.
3. Choose **User Lambda functions log level** and **Greengrass system log level**.
4. Keep the default values for logging level and disk space limit, and then choose **Save**.

Step 6: Deploy the Greengrass group

Deploy the group to the core device.

1. Make sure that the AWS IoT Greengrass core is running. Run the following commands in your Raspberry Pi terminal, as needed.

- a. To check whether the daemon is running:

```
ps aux | grep -E 'greengrass.*daemon'
```

If the output contains a root entry for `/greengrass/ggc/packages/ggc-version/bin/daemon`, then the daemon is running.

Note

The version in the path depends on the AWS IoT Greengrass Core software version that's installed on your core device.

- b. To start the daemon:

```
cd /greengrass/ggc/core/  
sudo ./greengrassd start
```

2. On the group configuration page, choose **Deploy**.
3.
 - a. In the **Lambda functions** tab, under the **System Lambda functions** section, select **IP detector** and choose **Edit**.
 - b. In the **Edit IP detector settings** dialog box, select **Automatically detect and override MQTT broker endpoints**.
 - c. Choose **Save**.

This enables devices to automatically acquire connectivity information for the core, such as IP address, DNS, and port number. Automatic detection is recommended, but AWS IoT Greengrass also supports manually specified endpoints. You're only prompted for the discovery method the first time that the group is deployed.

Note

If prompted, grant permission to create the [Greengrass service role](#) and associate it with your AWS account in the current AWS Region. This role allows AWS IoT Greengrass to access your resources in AWS services.

The **Deployments** page shows the deployment timestamp, version ID, and status. When completed, the status displayed for the deployment should be **Completed**.

For troubleshooting help, see [Troubleshooting](#).

Step 7: Test the application

The `TransferStream` Lambda function generates simulated device data. It writes data to a stream that stream manager exports to the target Kinesis data stream.

1. In the Amazon Kinesis console, under **Kinesis data streams**, choose **MyKinesisStream**.

Note

If you ran the tutorial without a target Kinesis data stream, [check the log file](#) for the stream manager (`GGStreamManager`). If it contains `export stream MyKinesisStream doesn't exist` in an error message, then the test is successful. This error means that the service tried to export to the stream but the stream doesn't exist.

2. On the **MyKinesisStream** page, choose **Monitoring**. If the test is successful, you should see data in the **Put Records** charts. Depending on your connection, it might take a minute before the data is displayed.

Important

When you're finished testing, delete the Kinesis data stream to avoid incurring more charges.

Or, run the following commands to stop the Greengrass daemon. This prevents the core from sending messages until you're ready to continue testing.

```
cd /greengrass/ggc/core/  
sudo ./greengrassd stop
```

3. Remove the **TransferStream** Lambda function from the core.
 - a. In the AWS IoT console navigation pane, under **Manage**, expand **Greengrass devices**, and then choose **Groups (V1)**.
 - b. Under **Greengrass groups**, choose your group.
 - c. On the **Lambdas** page, choose the ellipses (...) for the **TransferStream** function, and then choose **Remove function**.
 - d. From **Actions**, choose **Deploy**.

To view logging information or troubleshoot issues with streams, check the logs for the **TransferStream** and **GGStreamManager** functions. You must have root permissions to read AWS IoT Greengrass logs on the file system.

- **TransferStream** writes log entries to *greengrass-root*/ggc/var/log/user/*region*/*account-id*/TransferStream.log.
- **GGStreamManager** writes log entries to *greengrass-root*/ggc/var/log/system/GGStreamManager.log.

If you need more troubleshooting information, you can [set the logging level](#) for **User Lambda logs** to **Debug logs** and then deploy the group again.


See also

- [Manage data streams](#)
- [the section called "Configure stream manager"](#)
- [the section called "Use StreamManagerClient to work with streams"](#)
- [the section called "Export configurations for supported AWS Cloud destinations"](#)
- [the section called "Export data streams \(CLI\)"](#)

Export data streams to the AWS Cloud (CLI)

This tutorial shows you how to use the AWS CLI to configure and deploy an AWS IoT Greengrass group with stream manager enabled. The group contains a user-defined Lambda function that writes to a stream in stream manager, which is then exported automatically to the AWS Cloud.

Stream manager makes ingesting, processing, and exporting high-volume data streams more efficient and reliable. In this tutorial, you create a `TransferStream` Lambda function that consumes IoT data. The Lambda function uses the AWS IoT Greengrass Core SDK to create a stream in stream manager and then read and write to it. Stream manager then exports the stream to Kinesis Data Streams. The following diagram shows this workflow.



The focus of this tutorial is to show how user-defined Lambda functions use the `StreamManagerClient` object in the AWS IoT Greengrass Core SDK to interact with stream manager. For simplicity, the Python Lambda function that you create for this tutorial generates simulated device data.

When you use the AWS IoT Greengrass API, which includes the Greengrass commands in the AWS CLI, to create a group, stream manager is disabled by default. To enable stream manager on your core, you [create a function definition version](#) that includes the system `GGStreamManager` Lambda function and a group version that references the new function definition version. Then you deploy the group.

Prerequisites

To complete this tutorial, you need:

- A Greengrass group and a Greengrass core (v1.10 or later). For information about how to create a Greengrass group and core, see [Getting started with AWS IoT Greengrass](#). The Getting Started tutorial also includes steps for installing the AWS IoT Greengrass Core software.

Note

Stream manager is not supported on OpenWrt distributions.

- The Java 8 runtime (JDK 8) installed on the core device.
 - For Debian-based distributions (including Raspbian) or Ubuntu-based distributions, run the following command:

```
sudo apt install openjdk-8-jdk
```

- For Red Hat-based distributions (including Amazon Linux), run the following command:

```
sudo yum install java-1.8.0-openjdk
```

For more information, see [How to download and install prebuilt OpenJDK packages](#) in the OpenJDK documentation.

- AWS IoT Greengrass Core SDK for Python v1.5.0 or later. To use `StreamManagerClient` in the AWS IoT Greengrass Core SDK for Python, you must:
 - Install Python 3.7 or later on the core device.
 - Include the SDK and its dependencies in your Lambda function deployment package. Instructions are provided in this tutorial.

Tip

You can use `StreamManagerClient` with Java or NodeJS. For example code, see the [AWS IoT Greengrass Core SDK for Java](#) and [AWS IoT Greengrass Core SDK for Node.js](#) on GitHub.

- A destination stream named **MyKinesisStream** created in Amazon Kinesis Data Streams in the same AWS Region as your Greengrass group. For more information, see [Create a stream](#) in the *Amazon Kinesis Developer Guide*.

 Note

In this tutorial, stream manager exports data to Kinesis Data Streams, which results in charges to your AWS account. For information about pricing, see [Kinesis Data Streams pricing](#).

To avoid incurring charges, you can run this tutorial without creating a Kinesis data stream. In this case, you check the logs to see that stream manager attempted to export the stream to Kinesis Data Streams.

- An IAM policy added to the [the section called “Greengrass group role”](#) that allows the `kinesis:PutRecords` action on the target data stream, as shown in the following example:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "kinesis:PutRecords"
      ],
      "Resource": [
        "arn:aws:kinesis:region:account-id:stream/MyKinesisStream"
      ]
    }
  ]
}
```

- The AWS CLI installed and configured on your computer. For more information, see [Installing the AWS Command Line Interface](#) and [Configuring the AWS CLI](#) in the *AWS Command Line Interface User Guide*.

The example commands in this tutorial are written for Linux and other Unix-based systems. If you're using Windows, see [Specifying parameter values for the AWS command line interface](#) for more information about differences in syntax.

If the command contains a JSON string, the tutorial provides an example that has the JSON on a single line. On some systems, it might be more efficient to edit and run commands using this format.

The tutorial contains the following high-level steps:

1. [Create a Lambda function deployment package](#)
2. [Create a Lambda function](#)
3. [Create a function definition and version](#)
4. [Create a logger definition and version](#)
5. [Get the ARN of your core definition version](#)
6. [Create a group version](#)
7. [Create a deployment](#)
8. [Test the application](#)

The tutorial should take about 30 minutes to complete.

Step 1: Create a Lambda function deployment package

In this step, you create a Lambda function deployment package that contains Python function code and dependencies. You upload this package later when you create the Lambda function in AWS Lambda. The Lambda function uses the AWS IoT Greengrass Core SDK to create and interact with local streams.

Note

Your user-defined Lambda functions must use the [AWS IoT Greengrass Core SDK](#) to interact with stream manager. For more information about requirements for the Greengrass stream manager, see [Greengrass stream manager requirements](#).

1. Download the [AWS IoT Greengrass Core SDK for Python](#) v1.5.0 or later.
2. Unzip the downloaded package to get the SDK. The SDK is the `greengrasssdk` folder.
3. Install package dependencies to include with the SDK in your Lambda function deployment package.
 1. Navigate to the SDK directory that contains the `requirements.txt` file. This file lists the dependencies.
 2. Install the SDK dependencies. For example, run the following `pip` command to install them in the current directory:

```
pip install --target . -r requirements.txt
```

4. Save the following Python code function in a local file named `transfer_stream.py`.

 Tip

For example code that uses Java and NodeJS, see the [AWS IoT Greengrass Core SDK for Java](#) and [AWS IoT Greengrass Core SDK for Node.js](#) on GitHub.

```
import asyncio
import logging
import random
import time

from greengrasssdk.stream_manager import (
    ExportDefinition,
    KinesisConfig,
    MessageStreamDefinition,
    ReadMessagesOptions,
    ResourceNotFoundException,
    StrategyOnFull,
    StreamManagerClient,
)

# This example creates a local stream named "SomeStream".
# It starts writing data into that stream and then stream manager automatically
# exports
# the data to a customer-created Kinesis data stream named "MyKinesisStream".
```

```
# This example runs forever until the program is stopped.

# The size of the local stream on disk will not exceed the default (which is 256
  MB).
# Any data appended after the stream reaches the size limit continues to be
  appended, and
# stream manager deletes the oldest data until the total stream size is back under
  256 MB.
# The Kinesis data stream in the cloud has no such bound, so all the data from this
  script is
# uploaded to Kinesis and you will be charged for that usage.

def main(logger):
    try:
        stream_name = "SomeStream"
        kinesis_stream_name = "MyKinesisStream"

        # Create a client for the StreamManager
        client = StreamManagerClient()

        # Try deleting the stream (if it exists) so that we have a fresh start
        try:
            client.delete_message_stream(stream_name=stream_name)
        except ResourceNotFoundException:
            pass

        exports = ExportDefinition(
            kinesis=[KinesisConfig(identifier="KinesisExport" + stream_name,
            kinesis_stream_name=kinesis_stream_name)]
        )
        client.create_message_stream(
            MessageStreamDefinition(
                name=stream_name,
            strategy_on_full=StrategyOnFull.OverwriteOldestData, export_definition=exports
            )
        )

        # Append two messages and print their sequence numbers
        logger.info(
            "Successfully appended message to stream with sequence number %d",
            client.append_message(stream_name, "ABCDEFGHJKLMNOP".encode("utf-8")),
        )
        logger.info(
```

```

        "Successfully appended message to stream with sequence number %d",
        client.append_message(stream_name, "PQRSTUVWXYZ".encode("utf-8")),
    )

    # Try reading the two messages we just appended and print them out
    logger.info(
        "Successfully read 2 messages: %s",
        client.read_messages(stream_name,
            ReadMessagesOptions(min_message_count=2, read_timeout_millis=1000)),
    )

    logger.info("Now going to start writing random integers between 0 and 1000
to the stream")
    # Now start putting in random data between 0 and 1000 to emulate device
    sensor input
    while True:
        logger.debug("Appending new random integer to stream")
        client.append_message(stream_name, random.randint(0,
1000).to_bytes(length=4, signed=True, byteorder="big"))
        time.sleep(1)

    except asyncio.TimeoutError:
        logger.exception("Timed out while executing")
    except Exception:
        logger.exception("Exception while running")

def function_handler(event, context):
    return

logging.basicConfig(level=logging.INFO)
# Start up this sample code
main(logger=logging.getLogger())

```

5. Zip the following items into a file named `transfer_stream_python.zip`. This is your Lambda function deployment package.

- **transfer_stream.py**. App logic.
- **greengrasssdk**. Required library for Python Greengrass Lambda functions that publish MQTT messages.

[Stream manager operations](#) are available in version 1.5.0 or later of the AWS IoT Greengrass Core SDK for Python.

- The dependencies you installed for the AWS IoT Greengrass Core SDK for Python (for example, the cbor2 directories).

When you create the zip file, include only these items, not the containing folder.

Step 2: Create a Lambda function

1. Create an IAM role so you can pass in the role ARN when you create the function.

JSON Expanded

```
aws iam create-role --role-name Lambda_empty --assume-role-policy '{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Principal": {
        "Service": "lambda.amazonaws.com"
      },
      "Action": "sts:AssumeRole"
    }
  ]
}'
```

JSON Single-line

```
aws iam create-role --role-name Lambda_empty --assume-role-policy '{"Version":
"2012-10-17", "Statement": [{"Effect": "Allow", "Principal": {"Service":
"lambda.amazonaws.com"}, "Action": "sts:AssumeRole"}]}'
```

Note

AWS IoT Greengrass doesn't use this role because permissions for your Greengrass Lambda functions are specified in the Greengrass group role. For this tutorial, you create an empty role.

2. Copy the Arn from the output.
3. Use the AWS Lambda API to create the `TransferStream` function. The following command assumes that the zip file is in the current directory.
 - Replace *role-arn* with the Arn that you copied.

```
aws lambda create-function \  
--function-name TransferStream \  
--zip-file fileb://transfer_stream_python.zip \  
--role role-arn \  
--handler transfer_stream.function_handler \  
--runtime python3.7
```

4. Publish a version of the function.

```
aws lambda publish-version --function-name TransferStream --description 'First  
version'
```

5. Create an alias for the published version.

Greengrass groups can reference a Lambda function by alias (recommended) or by version. Using an alias makes it easier to manage code updates because you don't have to change your subscription table or group definition when the function code is updated. Instead, you just point the alias to the new function version.

```
aws lambda create-alias --function-name TransferStream --name GG_TransferStream --  
function-version 1
```

Note

AWS IoT Greengrass doesn't support Lambda aliases for **\$LATEST** versions.

6. Copy the `AliasArn` from the output. You use this value when you configure the function for AWS IoT Greengrass.

Now you're ready to configure the function for AWS IoT Greengrass.

Step 3: Create a function definition and version

This step creates a function definition version that references the system `GGStreamManager` Lambda function and your user-defined `TransferStream` Lambda function. To enable stream manager when you use the AWS IoT Greengrass API, your function definition version must include the `GGStreamManager` function.

1. Create a function definition with an initial version that contains the system and user-defined Lambda functions.

The following definition version enables stream manager with default [parameter settings](#). To configure custom settings, you must define environment variables for corresponding stream manager parameters. For an example, see [the section called "Enable, disable, or configure stream manager"](#). AWS IoT Greengrass uses default settings for parameters that are omitted. `MemorySize` should be at least 128000. `Pinned` must be set to `true`.

Note

A *long-lived* (or *pinned*) Lambda function starts automatically after AWS IoT Greengrass starts and keeps running in its own container. This is in contrast to an *on-demand* Lambda function, which starts when invoked and stops when there are no tasks left to run. For more information, see [the section called "Lifecycle configuration"](#).

- Replace *arbitrary-function-id* with a name for the function, such as **stream-manager**.
- Replace *alias-arn* with the `AliasArn` that you copied when you created the alias for the `TransferStream` Lambda function.

JSON expanded

```
aws greengrass create-function-definition --name MyGreengrassFunctions --
initial-version '{
  "Functions": [
    {
      "Id": "arbitrary-function-id",
      "FunctionArn": "arn:aws:lambda:::function:GGStreamManager:1",
      "FunctionConfiguration": {
        "MemorySize": 128000,
        "Pinned": true,
        "Timeout": 3
      }
    },
    {
      "Id": "TransferStreamFunction",
      "FunctionArn": "alias-arn",
      "FunctionConfiguration": {
        "Executable": "transfer_stream.function_handler",
        "MemorySize": 16000,
        "Pinned": true,
        "Timeout": 5
      }
    }
  ]
}'
```

JSON single

```
aws greengrass create-function-definition \
--name MyGreengrassFunctions \
--initial-version '{"Functions": [{"Id": "arbitrary-function-  
id", "FunctionArn": "arn:aws:lambda:::function:GGStreamManager:1",  
  "FunctionConfiguration": {"Environment": {"Variables":  
{"STREAM_MANAGER_STORE_ROOT_DIR": "/data", "STREAM_MANAGER_SERVER_PORT":  
  "1234", "STREAM_MANAGER_EXPORTER_MAX_BANDWIDTH": "20000"}}}, "MemorySize":  
  128000, "Pinned": true, "Timeout": 3}}, {"Id": "TransferStreamFunction",  
  "FunctionArn": "alias-arn", "FunctionConfiguration": {"Executable":
```

```
"transfer_stream.function_handler", "MemorySize": 16000, "Pinned": true, "Timeout": 5}}]}'
```

Note

Timeout is required by the function definition version, but GGStreamManager doesn't use it. For more information about Timeout and other group-level settings, see [the section called "Controlling Greengrass Lambda function execution"](#).

2. Copy the LatestVersionArn from the output. You use this value to add the function definition version to the group version that you deploy to the core.

Step 4: Create a logger definition and version

Configure the group's logging settings. For this tutorial, you configure AWS IoT Greengrass system components, user-defined Lambda functions, and connectors to write logs to the file system of the core device. You can use logs to troubleshoot any issues you might encounter. For more information, see [the section called "Monitoring with AWS IoT Greengrass logs"](#).

1. Create a logger definition that includes an initial version.

JSON Expanded

```
aws greengrass create-logger-definition --name "LoggingConfigs" --initial-version '{
  "Loggers": [
    {
      "Id": "1",
      "Component": "GreengrassSystem",
      "Level": "INFO",
      "Space": 10240,
      "Type": "FileSystem"
    },
    {
      "Id": "2",
      "Component": "Lambda",
      "Level": "INFO",
      "Space": 10240,
      "Type": "FileSystem"
    }
  ]
}
```

```
}
  ]
}'
```

JSON Single-line

```
aws greengrass create-logger-definition \
  --name "LoggingConfigs" \
  --initial-version '{"Loggers":
[{"Id":"1","Component":"GreengrassSystem","Level":"INFO","Space":10240,"Type":"FileSystem"},
{"Id":"2","Component":"Lambda","Level":"INFO","Space":10240,"Type":"FileSystem"}]}'
```

2. Copy the `LatestVersionArn` of the logger definition from the output. You use this value to add the logger definition version to the group version that you deploy to the core.

Step 5: Get the ARN of your core definition version

Get the ARN of the core definition version to add to your new group version. To deploy a group version, it must reference a core definition version that contains exactly one core.

1. Get the IDs of the target Greengrass group and group version. This procedure assumes that this is the latest group and group version. The following query returns the most recently created group.

```
aws greengrass list-groups --query "reverse(sort_by(Groups, &CreationTimestamp))
[0]"
```

Or, you can query by name. Group names are not required to be unique, so multiple groups might be returned.

```
aws greengrass list-groups --query "Groups[?Name=='MyGroup']"
```

Note

You can also find these values in the AWS IoT console. The group ID is displayed on the group's **Settings** page. Group version IDs are displayed on the group's **Deployments** tab.

2. Copy the Id of the target group from the output. You use this to get the core definition version and when you deploy the group.
3. Copy the LatestVersion from the output, which is the ID of the last version added to the group. You use this to get the core definition version.
4. Get the ARN of the core definition version:
 - a. Get the group version.
 - Replace *group-id* with the Id that you copied for the group.
 - Replace *group-version-id* with the LatestVersion that you copied for the group.

```
aws greengrass get-group-version \  
--group-id group-id \  
--group-version-id group-version-id
```

- b. Copy the CoreDefinitionVersionArn from the output. You use this value to add the core definition version to the group version that you deploy to the core.

Step 6: Create a group version

Now, you're ready to create a group version that contains the entities that you want to deploy. You do this by creating a group version that references the target version of each component type. For this tutorial, you include a core definition version, a function definition version, and a logger definition version.

1. Create a group version.
 - Replace *group-id* with the Id that you copied for the group.
 - Replace *core-definition-version-arn* with the CoreDefinitionVersionArn that you copied for the core definition version.
 - Replace *function-definition-version-arn* with the LatestVersionArn that you copied for your new function definition version.
 - Replace *logger-definition-version-arn* with the LatestVersionArn that you copied for your new logger definition version.

```
aws greengrass create-group-version \  

```

```
--group-id group-id \  
--core-definition-version-arn core-definition-version-arn \  
--function-definition-version-arn function-definition-version-arn \  
--logger-definition-version-arn logger-definition-version-arn
```

2. Copy the Version from the output. This is the ID of the new group version.

Step 7: Create a deployment

Deploy the group to the core device.

1. Make sure that the AWS IoT Greengrass core is running. Run the following commands in your Raspberry Pi terminal, as needed.

- a. To check whether the daemon is running:

```
ps aux | grep -E 'greengrass.*daemon'
```

If the output contains a root entry for `/greengrass/ggc/packages/ggc-version/bin/daemon`, then the daemon is running.

Note

The version in the path depends on the AWS IoT Greengrass Core software version that's installed on your core device.

- b. To start the daemon:

```
cd /greengrass/ggc/core/  
sudo ./greengrassd start
```

2. Create a deployment.

- Replace *group-id* with the Id that you copied for the group.
- Replace *group-version-id* with the Version that you copied for the new group version.

```
aws greengrass create-deployment \  
--deployment-type NewDeployment \  
--group-id group-id \  

```

```
--group-version-id group-version-id
```

3. Copy the DeploymentId from the output.
4. Get the deployment status.
 - Replace *group-id* with the Id that you copied for the group.
 - Replace *deployment-id* with the DeploymentId that you copied for the deployment.

```
aws greengrass get-deployment-status \  
--group-id group-id \  
--deployment-id deployment-id
```

If the status is Success, the deployment was successful. For troubleshooting help, see [Troubleshooting](#).

Step 8: Test the application

The TransferStream Lambda function generates simulated device data. It writes data to a stream that stream manager exports to the target Kinesis data stream.

1. In the Amazon Kinesis console, under **Kinesis data streams**, choose **MyKinesisStream**.

Note

If you ran the tutorial without a target Kinesis data stream, [check the log file](#) for the stream manager (GGStreamManager). If it contains `export stream MyKinesisStream doesn't exist` in an error message, then the test is successful. This error means that the service tried to export to the stream but the stream doesn't exist.

2. On the **MyKinesisStream** page, choose **Monitoring**. If the test is successful, you should see data in the **Put Records** charts. Depending on your connection, it might take a minute before the data is displayed.

⚠ Important

When you're finished testing, delete the Kinesis data stream to avoid incurring more charges.

Or, run the following commands to stop the Greengrass daemon. This prevents the core from sending messages until you're ready to continue testing.

```
cd /greengrass/ggc/core/  
sudo ./greengrassd stop
```

3. Remove the **TransferStream** Lambda function from the core.
 - a. Follow [the section called "Create a group version"](#) to create a new group version. but remove the `--function-definition-version-arn` option in the `create-group-version` command. Or, create a function definition version that doesn't include the **TransferStream** Lambda function.

📌 Note

By omitting the system `GGStreamManager` Lambda function from the deployed group version, you disable stream management on the core.

- b. Follow [the section called "Create a deployment"](#) to deploy the new group version.

To view logging information or troubleshoot issues with streams, check the logs for the `TransferStream` and `GGStreamManager` functions. You must have root permissions to read AWS IoT Greengrass logs on the file system.

- `TransferStream` writes log entries to `greengrass-root/ggc/var/log/user/region/account-id/TransferStream.log`.
- `GGStreamManager` writes log entries to `greengrass-root/ggc/var/log/system/GGStreamManager.log`.

If you need more troubleshooting information, you can set the Lambda logging level to `DEBUG` and then create and deploy a new group version.

See also

- [Manage data streams](#)
- [the section called “Use StreamManagerClient to work with streams”](#)
- [the section called “Export configurations for supported AWS Cloud destinations”](#)
- [the section called “Configure stream manager”](#)
- [the section called “Export data streams \(console\)”](#)
- [AWS Identity and Access Management \(IAM\) commands](#) in the *AWS CLI Command Reference*
- [AWS Lambda commands](#) in the *AWS CLI Command Reference*
- [AWS IoT Greengrass commands](#) in the *AWS CLI Command Reference*

Deploy secrets to the AWS IoT Greengrass core

This feature is available for AWS IoT Greengrass Core v1.7 and later.


AWS IoT Greengrass lets you authenticate with services and applications from Greengrass devices without hard-coding passwords, tokens, or other secrets.

AWS Secrets Manager is a service that you can use to securely store and manage your secrets in the cloud. AWS IoT Greengrass extends Secrets Manager to Greengrass core devices, so your [connectors](#) and Lambda functions can use local secrets to interact with services and applications. For example, the Twilio Notifications connector uses a locally stored authentication token.

To integrate a secret into a Greengrass group, you create a group resource that references the Secrets Manager secret. This *secret resource* references the cloud secret by ARN. To learn how to create, manage, and use secret resources, see [the section called “Work with secret resources”](#).

AWS IoT Greengrass encrypts your secrets while in transit and at rest. During group deployment, AWS IoT Greengrass fetches the secret from Secrets Manager and creates a local, encrypted copy on the Greengrass core. After you rotate your cloud secrets in Secrets Manager, redeploy the group to propagate the updated values to the core.

The following diagram shows the high-level process of deploying a secret to the core. Secrets are encrypted in transit and at rest.



Using AWS IoT Greengrass to store your secrets locally offers these advantages:

- **Decoupled from code (not hard-coded).** This supports centrally managed credentials and helps protect sensitive data from the risk of compromise.
- **Available for offline scenarios.** Connectors and functions can securely access local services and software when disconnected from the internet.
- **Controlled access to secrets.** Only authorized connectors and functions in the group can access your secrets. AWS IoT Greengrass uses private key encryption to secure your secrets. Secrets are encrypted in transit and at rest. For more information, see [the section called “Secrets encryption”](#).
- **Controlled rotation.** After you rotate your secrets in Secrets Manager, redeploy the Greengrass group to update the local copies of your secrets. For more information, see [the section called “Creating and managing secrets”](#).

Important

AWS IoT Greengrass doesn't automatically update the values of local secrets after cloud versions are rotated. To update local values, you must redeploy the group.

Secrets encryption

AWS IoT Greengrass encrypts secrets in transit and at rest.

Important

Make sure that your user-defined Lambda functions handle secrets securely and don't log any any sensitive data that's stored in the secret. For more information, see [Mitigate the Risks of Logging and Debugging Your Lambda Function](#) in the *AWS Secrets Manager User Guide*. Although this documentation specifically refers to rotation functions, the recommendation also applies to Greengrass Lambda functions.

Encryption in transit

AWS IoT Greengrass uses Transport Layer Security (TLS) to encrypt all communication over the internet and local network. This protects secrets while in transit, which occurs when secrets are retrieved from Secrets Manager and deployed to the core. For supported TLS cipher suites, see [the section called “TLS cipher suites support”](#).

Encryption at rest

AWS IoT Greengrass uses the private key specified in [config.json](#) for encryption of the secrets that are stored on the core. For this reason, secure storage of the private key is critical for protecting local secrets. In the AWS [shared responsibility model](#), it's the responsibility of the customer to guarantee secure storage of the private key on the core device.

AWS IoT Greengrass supports two modes of private key storage:

- Using hardware security modules. For more information, see [the section called “Hardware security integration”](#).

Note

Currently, AWS IoT Greengrass supports only the [PKCS#1 v1.5](#) padding mechanism for encryption and decryption of local secrets when using hardware-based private keys. If you're following vendor-provided instructions to manually generate hardware-based private keys, make sure to choose PKCS#1 v1.5. AWS IoT Greengrass doesn't support Optimal Asymmetric Encryption Padding (OAEP).

- Using file system permissions (default).

The private key is used to secure the data key, which is used to encrypt local secrets. The data key is rotated with each group deployment.

The AWS IoT Greengrass core is the only entity that has access to the private key. Greengrass connectors or Lambda functions that are affiliated with a secret resource get the value of the secret from the core.

Requirements

These are the requirements for local secret support:

- You must be using AWS IoT Greengrass Core v1.7 or later.
- To get the values of local secrets, your user-defined Lambda functions must use AWS IoT Greengrass Core SDK v1.3.0 or later.
- The private key used for local secrets encryption must be specified in the Greengrass configuration file. By default, AWS IoT Greengrass uses the core private key stored in the file

system. To provide your own private key, see [the section called “Specify the private key for secret encryption”](#). Only the RSA key type is supported.

Note

Currently, AWS IoT Greengrass supports only the [PKCS#1 v1.5](#) padding mechanism for encryption and decryption of local secrets when using hardware-based private keys. If you're following vendor-provided instructions to manually generate hardware-based private keys, make sure to choose PKCS#1 v1.5. AWS IoT Greengrass doesn't support Optimal Asymmetric Encryption Padding (OAEP).

- AWS IoT Greengrass must be granted permission to get your secret values. This allows AWS IoT Greengrass to fetch the values during group deployment. If you're using the default Greengrass service role, then AWS IoT Greengrass already has access to secrets with names that start with *greengrass-*. To customize access, see [the section called “Allow AWS IoT Greengrass to get secret values”](#).

Note

We recommend that you use this naming convention to identify the secrets that AWS IoT Greengrass is allowed to access, even if you customize permissions. The console uses different permissions to read your secrets, so it's possible that you can select secrets in the console that AWS IoT Greengrass doesn't have permission to fetch. Using a naming convention can help avoid a permission conflict, which results in a deployment error.

Specify the private key for secret encryption

In this procedure, you provide the path to a private key that's used for local secret encryption. This must be an RSA key with a minimum length of 2048 bits. For more information about private keys used on the AWS IoT Greengrass core, see [the section called “Security principals”](#).

AWS IoT Greengrass supports two modes of private key storage: hardware-based or file system-based (default). For more information, see [the section called “Secrets encryption”](#).

Follow this procedure only if you want to change the default configuration, which uses the core private key in the file system. These steps are written with the assumption that you created your group and core as described in [Module 2](#) of the Getting Started tutorial.

1. Open the [config.json](#) file that's located in the `/greengrass-root/config` directory.

Note

`greengrass-root` represents the path where the AWS IoT Greengrass Core software is installed on your device. Typically, this is the `/greengrass` directory.

2. In the `crypto.principals.SecretsManager` object, for the `privateKeyPath` property, enter the path of the private key:

- If your private key is stored in the file system, specify the absolute path to the key. For example:

```
"SecretsManager" : {  
  "privateKeyPath" : "file:///somepath/hash.private.key"  
}
```

- If your private key is stored in a hardware security module (HSM), specify the path using the [RFC 7512 PKCS#11](#) URI scheme. For example:

```
"SecretsManager" : {  
  "privateKeyPath" : "pkcs11:object=private-key-label;type=private"  
}
```

For more information, see [the section called “Hardware security configuration”](#).

Note

Currently, AWS IoT Greengrass supports only the [PKCS#1 v1.5](#) padding mechanism for encryption and decryption of local secrets when using hardware-based private keys. If you're following vendor-provided instructions to manually generate hardware-based private keys, make sure to choose PKCS#1 v1.5. AWS IoT Greengrass doesn't support Optimal Asymmetric Encryption Padding (OAEP).

Allow AWS IoT Greengrass to get secret values

In this procedure, you add an inline policy to the Greengrass service role that allows AWS IoT Greengrass to get the values of your secrets.

Follow this procedure only if you want to grant AWS IoT Greengrass custom permissions to your secrets or if your Greengrass service role doesn't include the `AWSGreengrassResourceAccessRolePolicy` managed policy. `AWSGreengrassResourceAccessRolePolicy` grants access to secrets with names that start with *greengrass-*.

1. Run the following CLI command to get the ARN of the Greengrass service role:

```
aws greengrass get-service-role-for-account --region region
```

The returned ARN contains the role name.

```
{
  "AssociatedAt": "time-stamp",
  "RoleArn": "arn:aws:iam::account-id:role/service-role/role-name"
}
```

You use the ARN or name in the following step.

2. Add an inline policy that allows the `secretsmanager:GetSecretValue` action. For instructions, see [Adding and removing IAM policies](#) in the *IAM User Guide*.

You can grant granular access by explicitly listing secrets or using a wildcard `*` naming scheme, or you can grant conditional access to versioned or tagged secrets. For example, the following policy allows AWS IoT Greengrass to read only the specified secrets.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "secretsmanager:GetSecretValue"
      ],
      "Resource": [
        "arn:aws:secretsmanager:region:account-id:secret:greengrass-SecretA-abc",
        "arn:aws:secretsmanager:region:account-id:secret:greengrass-SecretB-xyz"
      ]
    }
  ]
}
```

```
]
}
```

Note

If you use a customer-managed AWS KMS key to encrypt secrets, your Greengrass service role must also allow the `kms:Decrypt` action.

For more information about IAM policies for Secrets Manager, see [Authentication and access control for AWS Secrets Manager](#) and [Actions, resources, and context keys you can use in an IAM policy or secret policy for AWS Secrets Manager](#) in the *AWS Secrets Manager User Guide*.

See also

- [What is AWS Secrets Manager?](#) in the *AWS Secrets Manager User Guide*
- [PKCS #1: RSA Encryption Version 1.5](#)

Working with secret resources

AWS IoT Greengrass uses *secret resources* to integrate secrets from AWS Secrets Manager into a Greengrass group. A secret resource is a reference to a Secrets Manager secret. For more information, see [Deploy secrets to the core](#).

On the AWS IoT Greengrass core device, connectors and Lambda functions can use the secret resource to authenticate with services and applications, without hard-coding passwords, tokens, or other credentials.

Creating and managing secrets

In a Greengrass group, a secret resource references the ARN of a Secrets Manager secret. When the secret resource is deployed to the core, the value of the secret is encrypted and made available to affiliated connectors and Lambda functions. For more information, see [the section called “Secrets encryption”](#).

You use Secrets Manager to create and manage the cloud versions of your secrets. You use AWS IoT Greengrass to create, manage, and deploy your secret resources.

⚠ Important

We recommend that you follow the best practice of rotating your secrets in Secrets Manager. Then, deploy the Greengrass group to update the local copies of your secrets. For more information, see [Rotating your AWS Secrets Manager secrets](#) in the *AWS Secrets Manager User Guide*.

To make a secret available on the Greengrass core

1. Create a secret in Secrets Manager. This is the cloud version of your secret, which is centrally stored and managed in Secrets Manager. Management tasks include rotating secret values and applying resource policies.
2. Create a secret resource in AWS IoT Greengrass. This is a type of group resource that references the cloud secret by ARN. You can reference a secret only once per group.
3. Configure your connector or Lambda function. You must affiliate the resource with a connector or function by specifying corresponding parameters or properties. This allows them to get the value of the locally deployed secret resource. For more information, see [the section called “Using local secrets”](#).
4. Deploy the Greengrass group. During deployment, AWS IoT Greengrass fetches the value of the cloud secret and creates (or updates) the local secret on the core.

Secrets Manager logs an event in AWS CloudTrail each time that AWS IoT Greengrass retrieves a secret value. AWS IoT Greengrass doesn't log any events related to the deployment or usage of local secrets. For more information about Secrets Manager logging, see [Monitor the use of your AWS Secrets Manager secrets](#) in the *AWS Secrets Manager User Guide*.

Including staging labels in secret resources

Secrets Manager uses staging labels to identify specific versions of a secret value. Staging labels can be system-defined or user-defined. Secrets Manager assigns the `AWSCURRENT` label to the most recent version of the secret value. Staging labels are commonly used to manage secrets rotation. For more information about Secrets Manager versioning, see [Key terms and concepts for AWS Secrets Manager](#) in the *AWS Secrets Manager User Guide*.

Secret resources always include the `AWSCURRENT` staging label, and they can optionally include other staging labels if they're required by a Lambda function or connector. During group

deployment, AWS IoT Greengrass retrieves the values of the staging labels that are referenced in the group, and then creates or updates the corresponding values on the core.

Create and manage secret resources (console)

Creating secret resources (console)

In the AWS IoT Greengrass console, you create and manage secret resources from the **Secrets** tab on the group's **Resources** page. For tutorials that create a secret resource and add it to a group, see [the section called “How to create a secret resource \(console\)”](#) and [the section called “Get started with connectors \(console\)”](#).

Deployments	Resources			
Subscriptions	Local	Machine Learning	Secret	
Cores				
Devices				
Lambdas				
Resources				
Connectors				
Tags				
Settings				

Resource Name	Secret Name	Status	Labels
MyTwilioAuthToken	greengrass-TwilioAuthTo...	Unaffiliated	AWSCURRENT

Note

Alternatively, the console allows you to create a secret and secret resource when you configure a connector or Lambda function. You can do this from the connector's **Configure parameters** page or the Lambda function's **Resources** page.

Managing secret resources (console)

Management tasks for the secret resources in your Greengrass group include adding secret resources to the group, removing secret resources from the group, and changing the set of [staging labels](#) that are included in a secret resource.

If you point to a different secret from Secrets Manager, you must also edit any connectors that use the secret:

1. On the group configuration page, choose **Connectors**.
2. From the connector's contextual menu, choose **Edit**.
3. The **Edit parameters** page displays a message to inform you that the secret ARN changed. To confirm the change, choose **Save**.

If you delete a secret in Secrets Manager, remove the corresponding secret resource from the group and from connectors and Lambda functions that reference it. Otherwise, during group deployment, AWS IoT Greengrass returns an error that the secret can't be found. Also update your Lambda function code as needed.

Create and manage secret resources (CLI)

Creating secret resources (CLI)

In the AWS IoT Greengrass API, a secret is a type of group resource. The following example creates a resource definition with an initial version that includes a secret resource named `MySecretResource`. For a tutorial that creates a secret resource and adds it to a group version, see [the section called “Get started with connectors \(CLI\)”](#).

The secret resource references the ARN of the corresponding Secrets Manager secret and includes two staging labels in addition to `AWSCURRENT`, which is always included.

```
aws greengrass create-resource-definition --name MyGreengrassResources --initial-  
version '{  
  "Resources": [  
    {  
      "Id": "my-resource-id",  
      "Name": "MySecretResource",  
      "ResourceDataContainer": {  
        "SecretsManagerSecretResourceData": {  
          "ARN": "arn:aws:secretsmanager:us-  
west-2:123456789012:secret:greengrass-SomeSecret-KUj89s",  
          "AdditionalStagingLabelsToDownload": [  
            "Label1",  
            "Label2"  
          ]  
        }  
      }  
    ]  
  }  
}
```

```
}'
```

Managing secret resources (CLI)

Management tasks for the secret resources in your Greengrass group include adding secret resources to the group, removing secret resources from the group, and changing the set of [staging labels](#) that are included in a secret resource.

In the AWS IoT Greengrass API, these changes are implemented by using versions.

The AWS IoT Greengrass API uses versions to manage groups. Versions are immutable, so to add or change group components—for example, the group's client devices, functions, and resources—you must create versions of new or updated components. Then, you create and deploy a group version that contains the target version of each component. To learn more about groups, see [the section called “AWS IoT Greengrass groups”](#).

For example, to change the set of staging labels for a secret resource:

1. Create a resource definition version that contains the updated secret resource. The following example adds a third staging label to the secret resource from the previous section.

Note

To add more resources to the version, include them in the Resources array.

```
aws greengrass create-resource-definition --name MyGreengrassResources --initial-  
version '{  
  "Resources": [  
    {  
      "Id": "my-resource-id",  
      "Name": "MySecretResource",  
      "ResourceDataContainer": {  
        "SecretsManagerSecretResourceData": {  
          "ARN": "arn:aws:secretsmanager:us-  
west-2:123456789012:secret:greengrass-SomeSecret-KUj89s",  
          "AdditionalStagingLabelsToDownload": [  
            "Label1",  
            "Label2",  
            "Label3"  
          ]  
        }  
      }  
    ]  
  }  
}
```

```
}  
  }  
}   
]
```

2. If the ID of the secret resource is changed, update connectors and functions that use the secret resource. In the new versions, update the parameter or property that corresponds to the resource ID. If the ARN of the secret is changed, you must also update the corresponding parameter for any connectors that use the secret.

 Note

The resource ID is an arbitrary identifier that's provided by the customer.

3. Create a group version that contains the target version of each component that you want to send to the core.
4. Deploy the group version.

For a tutorial that shows how to create and deploy secret resources, connectors, and functions, see [the section called “Get started with connectors \(CLI\)”](#).

If you delete a secret in Secrets Manager, remove the corresponding secret resource from the group and from connectors and Lambda functions that reference it. Otherwise, during group deployment, AWS IoT Greengrass returns an error that the secret can't be found. Also update your Lambda function code as needed. You can remove a local secret by deploying a resource definition version that doesn't contain the corresponding secret resource.

Using local secrets in connectors and Lambda functions

Greengrass connectors and Lambda functions use local secrets to interact with services and applications. The `AWSCURRENT` value is used by default, but values for other [staging labels](#) included in the secret resource are also available.

Connectors and functions must be configured before they can access local secrets. This affiliates the secret resource with connector or function.

Connectors

If a connector requires access to a local secret, it provides parameters that you configure with the information it needs to access the secret.

- To learn how to do this in the AWS IoT Greengrass console, see [the section called “Get started with connectors \(console\)”](#).
- To learn how to do this with the AWS IoT Greengrass CLI, see [the section called “Get started with connectors \(CLI\)”](#).

For information about requirements for individual connectors, see [the section called “AWS-provided Greengrass connectors”](#).

The logic for accessing and using the secret is built into the connector.

Lambda functions

To allow a Greengrass Lambda function to access a local secret, you configure the function's properties.

- To learn how to do this in the AWS IoT Greengrass console, see [the section called “How to create a secret resource \(console\)”](#).
- To do this in the AWS IoT Greengrass API, you provide the following information in the `ResourceAccessPolicies` property.
 - `ResourceId`: The ID of the secret resource in the Greengrass group. This is the resource that references the ARN of the corresponding Secrets Manager secret.
 - `Permission`: The type of access that the function has to the resource. Only `ro` (read-only) permission is supported for secret resources.

The following example creates a Lambda function that can access the `MyApiKey` secret resource.

```
aws greengrass create-function-definition --name MyGreengrassFunctions --initial-  
version '{  
  "Functions": [  
    {  
      "Id": "MyLambdaFunction",  
      "FunctionArn": "arn:aws:lambda:us-  
west-2:123456789012:function:myFunction:1",  
      "FunctionConfiguration": {  
        "Pinned": false,  
        "MemorySize": 16384,
```

```

        "Timeout": 10,
        "Environment": {
            "ResourceAccessPolicies": [
                {
                    "ResourceId": "MyApiKey",
                    "Permission": "ro"
                }
            ],
            "AccessSysfs": true
        }
    }
}
]
}'

```

To access local secrets at runtime, Greengrass Lambda functions call the `get_secret_value` function from the `secretsmanager` client in the AWS IoT Greengrass Core SDK (v1.3.0 or later).

The following example shows how to use the AWS IoT Greengrass Core SDK for Python to get a secret. It passes the name of the secret to the `get_secret_value` function. `SecretId` can be the name or ARN of the Secrets Manager secret (not the secret resource).

```

import greengrasssdk

secrets_client = greengrasssdk.client("secretsmanager")
secret_name = "greengrass-MySecret-abc"

def function_handler(event, context):
    response = secrets_client.get_secret_value(SecretId=secret_name)
    secret = response.get("SecretString")

```

For text type secrets, the `get_secret_value` function returns a string. For binary type secrets, it returns a base64-encoded string.

⚠ Important

Make sure that your user-defined Lambda functions handle secrets securely and don't log any sensitive data that's stored in the secret. For more information, see [Mitigate the Risks of Logging and Debugging Your Lambda Function](#) in the *AWS Secrets Manager User Guide*. Although this documentation specifically refers to rotation functions, the recommendation also applies to Greengrass Lambda functions.

The current value of the secret is returned by default. This is the version that the `AWSCURRENT` staging label is attached to. To access a different version, pass the name of the corresponding staging label for the optional `VersionStage` argument. For example:

```
import greengrasssdk

secrets_client = greengrasssdk.client("secretsmanager")
secret_name = "greengrass-TestSecret"
secret_version = "MyTargetLabel"

# Get the value of a specific secret version
def function_handler(event, context):
    response = secrets_client.get_secret_value(
        SecretId=secret_name, VersionStage=secret_version
    )
    secret = response.get("SecretString")
```

For another example function that calls `get_secret_value`, see [Create a Lambda function deployment package](#).

How to create a secret resource (console)

This feature is available for AWS IoT Greengrass Core v1.7 and later.

This tutorial shows how to use the AWS Management Console to add a *secret resource* to a Greengrass group. A secret resource is a reference to a secret from AWS Secrets Manager. For more information, see [Deploy secrets to the core](#).

On the AWS IoT Greengrass core device, connectors and Lambda functions can use the secret resource to authenticate with services and applications, without hard-coding passwords, tokens, or other credentials.

In this tutorial, you start by creating a secret in the AWS Secrets Manager console. Then, in the AWS IoT Greengrass console, you add a secret resource to a Greengrass group from the group's **Resources** page. This secret resource references the Secrets Manager secret. Later, you attach the secret resource to a Lambda function, which allows the function to get the value of the local secret.

 Note

Alternatively, the console allows you to create a secret and secret resource when you configure a connector or Lambda function. You can do this from the connector's **Configure parameters** page or the Lambda function's **Resources** page.

Only connectors that contain parameters for secrets can access secrets. For a tutorial that shows how the Twilio Notifications connector uses a locally stored authentication token, see [the section called "Get started with connectors \(console\)"](#).

The tutorial contains the following high-level steps:

1. [Create a Secrets Manager secret](#)
2. [Add a secret resource to a group](#)
3. [Create a Lambda function deployment package](#)
4. [Create a Lambda function](#)
5. [Add the function to the group](#)
6. [Attach the secret resource to the function](#)
7. [Add subscriptions to the group](#)
8. [Deploy the group](#)
9. [the section called "Test the Lambda function"](#)

The tutorial should take about 20 minutes to complete.

Prerequisites

To complete this tutorial, you need:

- A Greengrass group and a Greengrass core (v1.7 or later). To learn how to create a Greengrass group and core, see [Getting started with AWS IoT Greengrass](#). The Getting Started tutorial also includes steps for installing the AWS IoT Greengrass Core software.
- AWS IoT Greengrass must be configured to support local secrets. For more information, see [Secrets Requirements](#).

Note

This requirement includes allowing access to your Secrets Manager secrets. If you're using the default Greengrass service role, Greengrass has permission to get the values of secrets with names that start with *greengrass-*.

- To get the values of local secrets, your user-defined Lambda functions must use AWS IoT Greengrass Core SDK v1.3.0 or later.

Step 1: Create a Secrets Manager secret

In this step, you use the AWS Secrets Manager console to create a secret.

1. Sign in to the [AWS Secrets Manager console](#).

Note

For more information about this process, see [Step 1: Create and store your secret in AWS Secrets Manager](#) in the *AWS Secrets Manager User Guide*.

2. Choose **Store a new secret**.
3. Under **Choose secret type**, choose **Other type of secret**.
4. Under **Specify the key-value pairs to be stored for this secret**:
 - For **Key**, enter **test**.
 - For **Value**, enter **abcdefghi**.
5. Keep **aws/secretsmanager** selected for the encryption key, and then choose **Next**.

Note

You aren't charged by AWS KMS if you use the default AWS managed key that Secrets Manager creates in your account.

6. For **Secret name**, enter **greengrass-TestSecret**, and then choose **Next**.

Note

By default, the Greengrass service role allows AWS IoT Greengrass to get the value of secrets with names that start with *greengrass-*. For more information, see [secrets requirements](#).

7. This tutorial doesn't require rotation, so choose **disable automatic rotation**, and then choose **Next**.
8. On the **Review** page, review your settings, and then choose **Store**.

Next, you create a secret resource in your Greengrass group that references the secret.

Step 2: Add a secret resource to a Greengrass group

In this step, you configure a group resource that references the Secrets Manager secret.

1. In the AWS IoT console navigation pane, under **Manage**, expand **Greengrass devices**, and then choose **Groups (V1)**.
2. Choose the group that you want to add the secret resource to.
3. On the group configuration page, choose the **Resources** tab, and then scroll down to the **Secrets** section. The **Secrets** section displays the secret resources that belong to the group. You can add, edit, and remove secret resources from this section.

Note

Alternatively, the console allows you to create a secret and secret resource when you configure a connector or Lambda function. You can do this from the connector's **Configure parameters** page or the Lambda function's **Resources** page.

4. Choose **Add** under the **Secrets** section.

5. On the **Add a secret resource** page, enter **MyTestSecret** in the **Resource name**.
6. Under **Secret**, choose **greengrass-TestSecret**.
7. In the **Select labels (Optional)** section, the **AWSCURRENT** staging label represents the latest version of the secret. This label is always included in a secret resource.

Note

This tutorial requires the **AWSCURRENT** label only. You can optionally include labels that are required by your Lambda function or connector.

8. Choose **Add resource**.

Step 3: Create a Lambda function deployment package

To create a Lambda function, you must first create a Lambda function *deployment package* that contains the function code and dependencies. Greengrass Lambda functions require the [AWS IoT Greengrass Core SDK](#) for tasks such as communicating with MQTT messages in the core environment and accessing local secrets. This tutorial creates a Python function, so you use the Python version of the SDK in the deployment package.

Note

To get the values of local secrets, your user-defined Lambda functions must use AWS IoT Greengrass Core SDK v1.3.0 or later.

1. From the [AWS IoT Greengrass Core SDK](#) downloads page, download the AWS IoT Greengrass Core SDK for Python to your computer.
2. Unzip the downloaded package to get the SDK. The SDK is the `greengrasssdk` folder.
3. Save the following Python code function in a local file named `secret_test.py`.

```
import greengrasssdk

secrets_client = greengrasssdk.client("secretsmanager")
iot_client = greengrasssdk.client("iot-data")
secret_name = "greengrass-TestSecret"
send_topic = "secrets/output"
```

```
def function_handler(event, context):
    """
    Gets a secret and publishes a message to indicate whether the secret was
    successfully retrieved.
    """
    response = secrets_client.get_secret_value(SecretId=secret_name)
    secret_value = response.get("SecretString")
    message = (
        f"Failed to retrieve secret {secret_name}."
        if secret_value is None
        else f"Successfully retrieved secret {secret_name}."
    )
    iot_client.publish(topic=send_topic, payload=message)
    print("Published: " + message)
```

The `get_secret_value` function supports the name or ARN of the Secrets Manager secret for the `SecretId` value. This example uses the secret name. For this example secret, AWS IoT Greengrass returns the key-value pair: `{"test": "abcdefghi"}`.

Important

Make sure that your user-defined Lambda functions handle secrets securely and don't log any any sensitive data that's stored in the secret. For more information, see [Mitigate the Risks of Logging and Debugging Your Lambda Function](#) in the *AWS Secrets Manager User Guide*. Although this documentation specifically refers to rotation functions, the recommendation also applies to Greengrass Lambda functions.


4. Zip the following items into a file named `secret_test_python.zip`. When you create the ZIP file, include only the code and dependencies, not the containing folder.
 - **secret_test.py**. App logic.
 - **greengrasssdk**. Required library for all Python Greengrass Lambda functions.

This is your Lambda function deployment package.

Step 4: Create a Lambda function

In this step, you use the AWS Lambda console to create a Lambda function and configure it to use your deployment package. Then, you publish a function version and create an alias.

1. First, create the Lambda function.
 - a. In the AWS Management Console, choose **Services**, and open the AWS Lambda console.
 - b. Choose **Create function** and then choose **Author from scratch**.
 - c. In the **Basic information** section, use the following values:
 - For **Function name**, enter **SecretTest**.
 - For **Runtime**, choose **Python 3.7**.
 - For **Permissions**, keep the default setting. This creates an execution role that grants basic Lambda permissions. This role isn't used by AWS IoT Greengrass.
 - d. At the bottom of the page, choose **Create function**.
2. Next, register the handler and upload your Lambda function deployment package.
 - a. On the **Code** tab, under **Code source**, choose **Upload from**. From the dropdown, choose **.zip file**.



- b. Choose **Upload**, then choose your `secret_test_python.zip` deployment package. Then, choose **Save**.
- c. On the **Code** tab for the function, under **Runtime settings**, choose **Edit**, and then enter the following values.
 - For **Runtime**, choose **Python 3.7**.
 - For **Handler**, enter `secret_test.function_handler`
- d. Choose **Save**.

Note

The **Test** button on the AWS Lambda console doesn't work with this function. The AWS IoT Greengrass Core SDK doesn't contain modules that are required to run your Greengrass Lambda functions independently in the AWS Lambda console. These modules (for example, `greengrass_common`) are supplied to the functions after they are deployed to your Greengrass core.

3. Now, publish the first version of your Lambda function and create an [alias for the version](#).

Note

Greengrass groups can reference a Lambda function by alias (recommended) or by version. Using an alias makes it easier to manage code updates because you don't have to change your subscription table or group definition when the function code is updated. Instead, you just point the alias to the new function version.

- a. From the **Actions** menu, choose **Publish new version**.
- b. For **Version description**, enter **First version**, and then choose **Publish**.
- c. On the **SecretTest: 1** configuration page, from the **Actions** menu, choose **Create alias**.
- d. On the **Create a new alias** page, use the following values:
 - For **Name**, enter **GG_SecretTest**.
 - For **Version**, choose **1**.

Note

AWS IoT Greengrass doesn't support Lambda aliases for **\$LATEST** versions.

- e. Choose **Create**.

Now you're ready to add the Lambda function to your Greengrass group and attach the secret resource.

Step 5: Add the Lambda function to the Greengrass group

In this step, you add the Lambda function to the Greengrass group in the AWS IoT console.

1. On the group configuration page, choose the **Lambda functions** tab.
2. Under the **My Lambda functions** section, choose **Add**.
3. For the **Lambda function**, choose **SecretTest**.
4. For the **Lambda function version**, choose the alias to the version that you published.

Next, configure the lifecycle of the Lambda function.

1. In the **Lambda function configuration** section, make the following updates.

Note

We recommend that you run your Lambda function without containerization unless your business case requires it. This helps enable access to your device GPU and camera without configuring device resources. If you run without containerization, you must also grant root access to your AWS IoT Greengrass Lambda functions.

a. To run without containerization:

- For **System user and group**, choose **Another user ID/group ID**. For **System user ID**, enter **0**. For **System group ID**, enter **0**.

This allows your Lambda function to run as root. For more information about running as root, see [the section called "Setting the default access identity for Lambda functions in a group"](#).

Tip

You also must update your `config.json` file to grant root access to your Lambda function. For the procedure, see [the section called "Running a Lambda function as root"](#).

- For **Lambda function containerization**, choose **No container**.

For more information about running without containerization, see [the section called “Considerations when choosing Lambda function containerization”](#).

- For **Timeout**, enter **10 seconds**.
- For **Pinned**, choose **True**.

For more information, see [the section called “Lifecycle configuration”](#).

- Under **Additional Parameter**, for **Read access to /sys directory**, choose **Enabled**.

b. **To run in containerized mode instead:**

 Note

We do not recommend running in containerized mode unless your business case requires it.

- For **System user and group**, choose **Use group default**.
- For **Lambda function containerization**, choose **Use group default**.
- For **Memory limit**, enter **1024 MB**.
- For **Timeout**, enter **10 seconds**.
- For **Pinned**, choose **True**.

For more information, see [the section called “Lifecycle configuration”](#).

- Under **Additional Parameters**, for **Read access to /sys directory**, choose **Enabled**.

2. Choose **Add Lambda function**.

Next, associate the secret resource with the function.

Step 6: Attach the secret resource to the Lambda function

In this step, you associate the secret resource to the Lambda function in your Greengrass group. This associates the resource with the function, which allows the function to get the value of the local secret.

1. On the group configuration page, choose the **Lambda functions** tab.
2. Choose the **SecretTest** function.

3. On the function's details page, choose **Resources**.
4. Scroll to the **Secrets** section and choose **Associate**.
5. Choose **MyTestSecret**, and then choose **Associate**.

Step 7: Add subscriptions to the Greengrass group

In this step, you add subscriptions that allow AWS IoT and the Lambda function to exchange messages. One subscription allows AWS IoT to invoke the function, and one allows the function to send output data to AWS IoT.

1. On the group configuration page, choose the **Subscriptions** tab, and then choose **Add Subscription**.
2. Create a subscription that allows AWS IoT to publish messages to the function.

On the group configuration page, choose the **Subscriptions** tab, and then choose **Add subscription**.

3. On the **Create a subscription** page, configure the source and target, as follows:
 - a. In **Source type**, choose **Lambda function**, and then choose **IoT Cloud**.
 - b. In **Target type**, choose **Service**, and then choose **SecretTest**.
 - c. In the **Topic filter**, enter **secrets/input**, and then choose **Create subscription**.
4. Add a second subscription. Choose the **Subscriptions** tab, choose **Add subscription**, and configure the source and target, as follows:
 - a. In **Source type**, choose **Services**, and then choose **SecretTest**.
 - b. In **Target type**, choose **Lambda function**, and then choose **IoT Cloud**.
 - c. In the **Topic filter**, enter **secrets/output**, and then choose **Create subscription**.

Step 8: Deploy the Greengrass group

Deploy the group to the core device. During deployment, AWS IoT Greengrass fetches the value of the secret from Secrets Manager and creates a local, encrypted copy on the core.

1. Make sure that the AWS IoT Greengrass core is running. Run the following commands in your Raspberry Pi terminal, as needed.

- a. To check whether the daemon is running:

```
ps aux | grep -E 'greengrass.*daemon'
```

If the output contains a root entry for `/greengrass/ggc/packages/ggc-version/bin/daemon`, then the daemon is running.

 Note

The version in the path depends on the AWS IoT Greengrass Core software version that's installed on your core device.

- b. To start the daemon:

```
cd /greengrass/ggc/core/  
sudo ./greengrassd start
```

2. On the group configuration page, choose **Deploy**.
3.
 - a. In the **Lambda functions** tab, under the **System Lambda functions** section, select **IP detector** and choose **Edit**.
 - b. In the **Edit IP detector settings** dialog box, select **Automatically detect and override MQTT broker endpoints**.
 - c. Choose **Save**.

This enables devices to automatically acquire connectivity information for the core, such as IP address, DNS, and port number. Automatic detection is recommended, but AWS IoT Greengrass also supports manually specified endpoints. You're only prompted for the discovery method the first time that the group is deployed.

 Note

If prompted, grant permission to create the [Greengrass service role](#) and associate it with your AWS account in the current AWS Region. This role allows AWS IoT Greengrass to access your resources in AWS services.

The **Deployments** page shows the deployment timestamp, version ID, and status. When completed, the status displayed for the deployment should be **Completed**.

For troubleshooting help, see [Troubleshooting](#).

Test the Lambda function

1. On the AWS IoT console home page, choose **Test**.
2. For **Subscribe to topic**, use the following values, and then choose **Subscribe**.

Property	Value
Subscription topic	secrets/output
MQTT payload display	Display payloads as strings

3. For **Publish to topic**, use the following values, and then choose **Publish** to invoke the function.

Property	Value
Topic	secrets/input
Message	Keep the default message. Publishing a message invokes the Lambda function, but the function in this tutorial doesn't process the message body.

If successful, the function publishes a "Success" message.

See also


- [Deploy secrets to the core](#)

Integrate with services and protocols using Greengrass connectors


This feature is available for AWS IoT Greengrass Core v1.7 and later.

Connectors in AWS IoT Greengrass are prebuilt modules that make it more efficient to interact with local infrastructure, device protocols, AWS, and other cloud services. By using connectors, you can spend less time learning new protocols and APIs and more time focusing on the logic that matters to your business.

The following diagram shows where connectors can fit into the AWS IoT Greengrass landscape.



Many connectors use MQTT messages to communicate with client devices and Greengrass Lambda functions in the group, or with AWS IoT and the local shadow service. In the following example, the Twilio Notifications connector receives MQTT messages from a user-defined Lambda function, uses a local reference of a secret from AWS Secrets Manager, and calls the Twilio API.



For tutorials that create this solution, see [the section called “Get started with connectors \(console\)”](#) and [the section called “Get started with connectors \(CLI\)”](#).

Greengrass connectors can help you extend device capabilities or create single-purpose devices. By using connectors, you can:

- Implement reusable business logic.
- Interact with cloud and local services, including AWS and third-party services.
- Ingest and process device data.
- Enable device-to-device calls using MQTT topic subscriptions and user-defined Lambda functions.

AWS provides a set of Greengrass connectors that simplify interactions with common services and data sources. These prebuilt modules enable scenarios for logging and diagnostics, replenishment, industrial data processing, and alarm and messaging. For more information, see [the section called “AWS-provided Greengrass connectors”](#).

Requirements

To use connectors, keep these points in mind:

- Each connector that you use has requirements that you must meet. These requirements might include the minimum AWS IoT Greengrass Core software version, device prerequisites, required

permissions, and limits. For more information, see [the section called “AWS-provided Greengrass connectors”](#).


- A Greengrass group can contain only one configured instance of a given connector. However, you can use the instance in multiple subscriptions. For more information, see [the section called “Configuration parameters”](#).
- When the [default containerization](#) for the Greengrass group is set to **No container**, the connectors in the group must run without containerization. To find connectors that support **No container** mode, see [the section called “AWS-provided Greengrass connectors”](#).

Using Greengrass connectors

A connector is a type of group component. Like other group components, such as client devices and user-defined Lambda functions, you add connectors to groups, configure their settings, and deploy them to the AWS IoT Greengrass core. Connectors run in the core environment.

You can deploy some connectors as simple standalone applications. For example, the Device Defender connector reads system metrics from the core device and sends them to AWS IoT Device Defender for analysis.

You can add other connectors as building blocks in larger solutions. The following example solution uses the Modbus-RTU Protocol Adapter connector to process messages from sensors and the Twilio Notifications connector to initiate Twilio messages.



Solutions often include user-defined Lambda functions that sit next to connectors and process the data that the connector sends or receives. In this example, the TempMonitor function receives data from Modbus-RTU Protocol Adapter, runs some business logic, and then sends data to Twilio Notifications.

To create and deploy a solution, you follow this general process:

1. Map out the high-level data flow. Identify the data sources, data channels, services, protocols, and resources that you need to work with. In the example solution, this includes data over the Modbus RTU protocol, the physical Modbus serial port, and Twilio.
2. Identify the connectors to include in the solution, and add them to your group. The example solution uses Modbus-RTU Protocol Adapter and Twilio Notifications. To help you find connectors that apply to your scenario, and to learn about their individual requirements, see [the section called “AWS-provided Greengrass connectors”](#).
3. Identify whether user-defined Lambda functions, client devices, or resources are needed, and then create and add them to the group. This might include functions that contain business logic or process data into a format required by another entity in the solution. The example solution uses functions to send Modbus RTU requests and initiate Twilio notifications. It also includes a local device resource for the Modbus RTU serial port and a secret resource for the Twilio authentication token.

 Note

Secret resources reference passwords, tokens, and other secrets from AWS Secrets Manager. Secrets can be used by connectors and Lambda functions to authenticate with services and applications. By default, AWS IoT Greengrass can access secrets with names that start with *greengrass-*. For more information, see [Deploy secrets to the core](#).

4. Create subscriptions that allow the entities in the solution to exchange MQTT messages. If a connector is used in a subscription, the connector and the message source or target must use the predefined topic syntax supported by the connector. For more information, see [the section called “Inputs and outputs”](#).
5. Deploy the group to the Greengrass core.

For information about creating and deploying a connector, see the following tutorials:

- [the section called “Get started with connectors \(console\)”](#)

- [the section called “Get started with connectors \(CLI\)”](#)

Configuration parameters

Many connectors provide parameters that let you customize the behavior or output. These parameters are used during initialization, at runtime, or at other times in the connector lifecycle.

Parameter types and usage vary by connector. For example, the SNS connector has a parameter that configures the default SNS topic, and Device Defender has a parameter that configures the data sampling rate.

A group version can contain multiple connectors, but only one instance of a given connector at a time. This means that each connector in the group can have only one active configuration. However, the connector instance can be used in multiple subscriptions in the group. For example, you can create subscriptions that allow many devices to send data to the Kinesis Firehose connector.

Parameters used to access group resources

Greengrass connectors use group resources to access the file system, ports, peripherals, and other local resources on the core device. If a connector requires access to a group resource, then it provides related configuration parameters.

Group resources include:

- [Local resources](#). Directories, files, ports, pins, and peripherals that are present on the Greengrass core device.
- [Machine learning resources](#). Machine learning models that are trained in the cloud and deployed to the core for local inference.
- [Secret resources](#). Local, encrypted copies of passwords, keys, tokens, or arbitrary text from AWS Secrets Manager. Connectors can securely access these local secrets and use them to authenticate to services or local infrastructure.

For example, parameters for Device Defender enable access to system metrics in the host / `proc` directory, and parameters for Twilio Notifications enable access to a locally stored Twilio authentication token.

Updating connector parameters

Parameters are configured when the connector is added to a Greengrass group. You can change parameter values after the connector is added.

- In the console: From the group configuration page, open **Connectors**, and from the connector's contextual menu, choose **Edit**.

Note

If the connector uses a secret resource that's later changed to reference a different secret, you must edit the connector's parameters and confirm the change.

- In the API: Create another version of the connector that defines the new configuration.

The AWS IoT Greengrass API uses versions to manage groups. Versions are immutable, so to add or change group components—for example, the group's client devices, functions, and resources—you must create versions of new or updated components. Then, you create and deploy a group version that contains the target version of each component.

After you make changes to the connector configuration, you must deploy the group to propagate the changes to the core.

Inputs and outputs

Many Greengrass connectors can communicate with other entities by sending and receiving MQTT messages. MQTT communication is controlled by subscriptions that allow a connector to exchange data with Lambda functions, client devices, and other connectors in the Greengrass group, or with AWS IoT and the local shadow service. To allow this communication, you must create subscriptions in the group that the connector belongs to. For more information, see [the section called “Managed subscriptions in the MQTT messaging workflow”](#).

Connectors can be message publishers, message subscribers, or both. Each connector defines the MQTT topics that it publishes or subscribes to. These predefined topics must be used in the subscriptions where the connector is a message source or message target. For tutorials that include steps for configuring subscriptions for a connector, see [the section called “Get started with connectors \(console\)”](#) and [the section called “Get started with connectors \(CLI\)”](#).

Note

Many connectors also have built-in modes of communication to interact with cloud or local services. These vary by connector and might require that you configure parameters or add permissions to the [group role](#). For information about connector requirements, see [the section called "AWS-provided Greengrass connectors"](#).

Input topics

Most connectors receive input data on MQTT topics. Some connectors subscribe to multiple topics for input data. For example, the Serial Stream connector supports two topics:

- `serial/+/read/#`
- `serial/+/write/#`

For this connector, read and write requests are sent to the corresponding topic. When you create subscriptions, make sure to use the topic that aligns with your implementation.

The `+` and `#` characters in the previous examples are wildcards. These wildcards allow subscribers to receive messages on multiple topics and publishers to customize the topics that they publish to.

- The `+` wildcard can appear anywhere in the topic hierarchy. It can be replaced by one hierarchy item.

As an example, for topic `sensor/+/input`, messages can be published to topics `sensor/id-123/input` but not to `sensor/group-a/id-123/input`.

- The `#` wildcard can appear only at the end of the topic hierarchy. It can be replaced by zero or more hierarchy items.

As an example, for topic `sensor/#`, messages can be published to `sensor/`, `sensor/id-123`, and `sensor/group-a/id-123`, but not to `sensor`.

Wildcard characters are valid only when subscribing to topics. Messages can't be published to topics that contain wildcards. Check the documentation for the connector for more information

about its input or output topic requirements. For more information, see [the section called “AWS-provided Greengrass connectors”](#).

Containerization support

By default, most connectors run on the Greengrass core in an isolated runtime environment that's managed by AWS IoT Greengrass. These runtime environments, called *containers*, provide isolation between connectors and the host system, which offers more security for the host and the connector.

However, this Greengrass containerization isn't supported in some environments, such as when you run AWS IoT Greengrass in a Docker container or on older Linux kernels without cgroups. In these environments, the connectors must run in **No container** mode. To find connectors that support **No container** mode, see [the section called “AWS-provided Greengrass connectors”](#). Some connectors run in this mode natively, and some connectors allow you to set the isolation mode.

You can also set the isolation mode to **No container** in environments that support Greengrass containerization, but we recommend using **Greengrass container** mode when possible.

Note

The [default containerization](#) setting for the Greengrass group doesn't apply to connectors.

Upgrading connector versions

Connector providers might release new versions of a connector that add features, fix issues, or improve performance. For information about available versions and related changes, see the [documentation for each connector](#).

In the AWS IoT console, you can check for new versions for the connectors in your Greengrass group.

1. In the AWS IoT console navigation pane, under **Manage**, expand **Greengrass devices**, and then choose **Groups (V1)**.
2. Under **Greengrass groups**, choose your group.
3. Choose **Connectors** to display the connectors in the group.

If the connector has a new version, an **Available** button appears in the **Upgrade** column.

4. To upgrade the connector version:

- a. On the **Connectors** page, in the **Upgrade** column, choose **Available**. The **Upgrade connector** page opens and displays the current parameter settings, where applicable.

Choose the new connector version, define parameters as needed, and then choose **Upgrade**.

- b. On the **Subscriptions** page, add new subscriptions in the group to replace any that use the connector as a source or target. Then, remove the old subscriptions.

Subscriptions reference connectors by version, so they become invalid if you change the connector version in the group.

- c. From the **Actions** menu, choose **Deploy** to deploy your changes to the core.

To upgrade a connector from the AWS IoT Greengrass API, create and deploy a group version that includes the updated connector and subscriptions. Use the same process as when you add a connector to a group. For detailed steps that show you how to use the AWS CLI to configure and deploy an example Twilio Notifications connector, see [the section called “Get started with connectors \(CLI\)”](#).

Logging for connectors

Greengrass connectors contain Lambda functions that write events and errors to Greengrass logs. Depending on your group settings, logs are written to CloudWatch Logs, the local file system, or both. Logs from connectors include the ARN of the corresponding function. The following example ARN is from the Kinesis Firehose connector:

```
arn:aws:lambda:aws-region:account-id:function:KinesisFirehoseClient:1
```

The default logging configuration writes info-level logs to the file system using the following directory structure:

```
greengrass-root/ggc/var/log/user/region/aws/function-name.log
```

For more information about Greengrass logging, see [the section called “Monitoring with AWS IoT Greengrass logs”](#).

AWS-provided Greengrass connectors

AWS provides the following connectors that support common AWS IoT Greengrass scenarios. For more information about how connectors work, see the following documentation:

- [Integrate with services and protocols using connectors](#)
- [Get started with connectors \(console\)](#) or [Get started with connectors \(CLI\)](#)

Connector	Description	Supported Lambda runtimes	Supports No container mode
CloudWatch Metrics	Publishes custom metrics to Amazon CloudWatch.	• Python 3.8 * • Python 3.7 • Python 2.7	Yes
Device Defender	Sends system metrics to AWS IoT Device Defender.	• Python 3.8 * • Python 3.7 • Python 2.7	No
Docker Application Deployment	Runs a Docker Compose file to start a Docker application on the core device.	• Python 3.8 • Python 3.7	Yes
IoT Analytics	Sends data from devices and sensors to AWS IoT Analytics.	• Python 3.8 * • Python 3.7 • Python 2.7	Yes
IoT Ethernet IP Protocol Adapter	Collects data from Ethernet/IP devices.	• Java 8	Yes

Connector	Description	Supported Lambda runtimes	Supports No container mode
IoT SiteWise	Sends data from devices and sensors to asset properties in AWS IoT SiteWise.	• Java 8	Yes
Kinesis Firehose	Sends data to Amazon Data Firehose delivery streams.	• Python 3.8 * • Python 3.7 • Python 2.7	Yes
ML Feedback	Publishes machine learning model input to the cloud and output to an MQTT topic.	• Python 3.8 • Python 3.7	No
ML Image Classification	Runs a local image classification inference service. This connector provides versions for several platforms.	• Python 3.8 * • Python 3.7 • Python 2.7	No
ML Object Detection	Runs a local object detection inference service. This connector provides versions for several platforms.	• Python 3.8 • Python 3.7	No
Modbus-RTU Protocol Adapter	Sends requests to Modbus RTU devices.	• Python 3.8 * • Python 3.7 • Python 2.7	No
Modbus-TCP Protocol Adapter	Collects data from ModbusTCP devices.	• Java 8	Yes

Connector	Description	Supported Lambda runtimes	Supports No container mode
Raspberry Pi GPIO	Controls GPIO pins on a Raspberry Pi core device.	• Python 3.8 * • Python 3.7 • Python 2.7	No
Serial Stream	Reads and writes to a serial port on the core device.	• Python 3.8 * • Python 3.7 • Python 2.7	No
ServiceNow MetricBase Integration	Publishes time series metrics to ServiceNow MetricBase.	• Python 3.8 * • Python 3.7 • Python 2.7	Yes
SNS	Sends messages to an Amazon SNS topic.	• Python 3.8 * • Python 3.7 • Python 2.7	Yes
Splunk Integration	Publishes data to Splunk HEC.	• Python 3.8 * • Python 3.7 • Python 2.7	Yes
Twilio Notifications	Initiates a Twilio text or voice message.	• Python 3.8 * • Python 3.7 • Python 2.7	Yes

* To use the Python 3.8 runtimes, you must create a symbolic link from the default Python 3.7 installation folder to the installed Python 3.8 binaries. For more information, see the connector-specific requirements.

 Note

We recommend that you [upgrade connector versions](#) from Python 2.7 to Python 3.7. Continued support for Python 2.7 connectors depends on AWS Lambda runtime support. For more information, see [Runtime support policy](#) in the *AWS Lambda Developer Guide*.

CloudWatch Metrics connector

The CloudWatch Metrics [connector](#) publishes custom metrics from Greengrass devices to Amazon CloudWatch. The connector provides a centralized infrastructure for publishing CloudWatch metrics, which you can use to monitor and analyze the Greengrass core environment, and act on local events. For more information, see [Using Amazon CloudWatch metrics](#) in the *Amazon CloudWatch User Guide*.

This connector receives metric data as MQTT messages. The connector batches metrics that are in the same namespace and publishes them to CloudWatch at regular intervals.

This connector has the following versions.

Version	ARN
5	arn:aws:greengrass: <i>region</i> ::/connectors/CloudWatchMetrics/versions/5
4	arn:aws:greengrass: <i>region</i> ::/connectors/CloudWatchMetrics/versions/4
3	arn:aws:greengrass: <i>region</i> ::/connectors/CloudWatchMetrics/versions/3

Version	ARN
2	arn:aws:greengrass: <i>region</i> ::/connectors/CloudWatchMetrics/versions/2
1	arn:aws:greengrass: <i>region</i> ::/connectors/CloudWatchMetrics/versions/1

For information about version changes, see the [Changelog](#).

Requirements

This connector has the following requirements:

Version 3 - 5

- AWS IoT Greengrass Core software v1.9.3 or later.
- [Python](#) version 3.7 or 3.8 installed on the core device and added to the PATH environment variable.

Note

To use Python 3.8, run the following command to create a symbolic link from the the default Python 3.7 installation folder to the installed Python 3.8 binaries.

```
sudo ln -s path-to-python-3.8/python3.8 /usr/bin/python3.7
```

This configures your device to meet the Python requirement for AWS IoT Greengrass.

- The [Greengrass group role](#) configured to allow the `cloudwatch:PutMetricData` action, as shown in the following example AWS Identity and Access Management (IAM) policy.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "Stmt1528133056761",
```

```
        "Action": [
            "cloudwatch:PutMetricData"
        ],
        "Effect": "Allow",
        "Resource": "*"
    }
]
```

For the group role requirement, you must configure the role to grant the required permissions and make sure the role has been added to the group. For more information, see [the section called “Manage the group role \(console\)”](#) or [the section called “Manage the group role \(CLI\)”](#).

For more information about CloudWatch permissions, see [Amazon CloudWatch permissions reference](#) in the *IAM User Guide*.

Versions 1 - 2

- AWS IoT Greengrass Core software v1.7 or later.
- [Python](#) version 2.7 installed on the core device and added to the PATH environment variable.
- The [Greengrass group role](#) configured to allow the `cloudwatch:PutMetricData` action, as shown in the following example AWS Identity and Access Management (IAM) policy.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "Stmt1528133056761",
      "Action": [
        "cloudwatch:PutMetricData"
      ],
      "Effect": "Allow",
      "Resource": "*"
    }
  ]
}
```

For the group role requirement, you must configure the role to grant the required permissions and make sure the role has been added to the group. For more information, see [the section called “Manage the group role \(console\)”](#) or [the section called “Manage the group role \(CLI\)”](#).

For more information about CloudWatch permissions, see [Amazon CloudWatch permissions reference](#) in the *IAM User Guide*.

Connector Parameters

This connector provides the following parameters:

Versions 4 - 5

PublishInterval

The maximum number of seconds to wait before publishing batched metrics for a given namespace. The maximum value is 900. To configure the connector to publish metrics as they are received (without batching), specify 0.

The connector publishes to CloudWatch after it receives 20 metrics in the same namespace or after the specified interval.

Note

The connector doesn't guarantee the order of publish events.

Display name in the AWS IoT console: **Publish interval**

Required: true

Type: string

Valid values: 0 - 900

Valid pattern: [0-9] | [1-9]\d | [1-9]\d\d | 900

PublishRegion

The AWS Region to post CloudWatch metrics to. This value overrides the default Greengrass metrics Region. It is required only when posting cross-Region metrics.

Display name in the AWS IoT console: **Publish region**

Required: false

Type: string

Valid pattern: `^$|([a-z]{2}-[a-z]+-\d{1})`

MemorySize

The memory (in KB) to allocate to the connector.

Display name in the AWS IoT console: **Memory size**

Required: true

Type: string

Valid pattern: `^[0-9]+$`

MaxMetricsToRetain

The maximum number of metrics across all namespaces to save in memory before they are replaced with new metrics. The minimum value is 2000.

This limit applies when there's no connection to the internet and the connector starts to buffer the metrics to publish later. When the buffer is full, the oldest metrics are replaced by new metrics. Metrics in a given namespace are replaced only by metrics in the same namespace.

Note

Metrics are not saved if the host process for the connector is interrupted. For example, this interruption can happen during group deployment or when the device restarts.

Display name in the AWS IoT console: **Maximum metrics to retain**

Required: true

Type: string

Valid pattern: `^([2-9]\d{3}|[1-9]\d{4,})$`

IsolationMode

The [containerization](#) mode for this connector. The default is `GreengrassContainer`, which means that the connector runs in an isolated runtime environment inside the AWS IoT Greengrass container.

Note

The default containerization setting for the group does not apply to connectors.

Display name in the AWS IoT console: **Container isolation mode**

Required: `false`

Type: `string`

Valid values: `GreengrassContainer` or `NoContainer`

Valid pattern: `^NoContainer$|^GreengrassContainer$`

Versions 1 - 3

PublishInterval

The maximum number of seconds to wait before publishing batched metrics for a given namespace. The maximum value is 900. To configure the connector to publish metrics as they are received (without batching), specify 0.

The connector publishes to CloudWatch after it receives 20 metrics in the same namespace or after the specified interval.

Note

The connector doesn't guarantee the order of publish events.

Display name in the AWS IoT console: **Publish interval**

Required: `true`

Type: `string`

Valid values: 0 - 900

Valid pattern: [0-9] | [1-9]\d | [1-9]\d\d | 900

PublishRegion

The AWS Region to post CloudWatch metrics to. This value overrides the default Greengrass metrics Region. It is required only when posting cross-Region metrics.

Display name in the AWS IoT console: **Publish region**

Required: false

Type: string

Valid pattern: ^\$ | ([a-z]{2}-[a-z]+-\d{1})

MemorySize

The memory (in KB) to allocate to the connector.

Display name in the AWS IoT console: **Memory size**

Required: true

Type: string

Valid pattern: ^[0-9]+\$

MaxMetricsToRetain

The maximum number of metrics across all namespaces to save in memory before they are replaced with new metrics. The minimum value is 2000.

This limit applies when there's no connection to the internet and the connector starts to buffer the metrics to publish later. When the buffer is full, the oldest metrics are replaced by new metrics. Metrics in a given namespace are replaced only by metrics in the same namespace.

Note

Metrics are not saved if the host process for the connector is interrupted. For example, this interruption can happen during group deployment or when the device restarts.

Display name in the AWS IoT console: **Maximum metrics to retain**

Required: true

Type: string

Valid pattern: `^([2-9]\d{3}|[1-9]\d{4,})$`

Create Connector Example (AWS CLI)

The following CLI command creates a `ConnectorDefinition` with an initial version that contains the CloudWatch Metrics connector.

```
aws greengrass create-connector-definition --name MyGreengrassConnectors --initial-  
version '{  
  "Connectors": [  
    {  
      "Id": "MyCloudWatchMetricsConnector",  
      "ConnectorArn": "arn:aws:greengrass:region::/connectors/CloudWatchMetrics/  
versions/4",  
      "Parameters": {  
        "PublishInterval" : "600",  
        "PublishRegion" : "us-west-2",  
        "MemorySize" : "16",  
        "MaxMetricsToRetain" : "2500",  
        "IsolationMode" : "GreengrassContainer"  
      }  
    }  
  ]  
}'
```

In the AWS IoT Greengrass console, you can add a connector from the group's **Connectors** page. For more information, see [the section called "Get started with connectors \(console\)"](#).

Input data

This connector accepts metrics on an MQTT topic and publishes the metrics to CloudWatch. Input messages must be in JSON format.

Topic filter in subscription

cloudwatch/metric/put

Message properties

`request`

Information about the metric in this message.

The request object contains the metric data to publish to CloudWatch. The metric values must meet the specifications of the [PutMetricData](#) API. Only the `namespace`, `metricData.metricName`, and `metricData.value` properties are required.

Required: `true`

Type: object that includes the following properties:

`namespace`

The user-defined namespace for the metric data in this request. CloudWatch uses namespaces as containers for metric data points.

Note

You can't specify a namespace that begins with the reserved string `AWS/`.

Required: `true`

Type: `string`

Valid pattern: `^[^:]*`

`metricData`

The data for the metric.

Required: `true`

Type: object that includes the following properties:

`metricName`

The name of the metric.

Required: `true`

Type: `string`

dimensions

The dimensions that are associated with the metric. Dimensions provide more information about the metric and its data. A metric can define up to 10 dimensions.

This connector automatically includes a dimension named `coreName`, where the value is the name of the core.

Required: false

Type: array of dimension objects that include the following properties:

name

The dimension name.

Required: false

Type: string

value

The dimension value.

Required: false

Type: string

timestamp

The time that the metric data was received, expressed as the number of seconds since Jan 1, 1970 00:00:00 UTC. If this value is omitted, the connector uses the time that it received the message.

Required: false

Type: timestamp

Note

If you use between versions 1 and 4 of this connector, we recommend that you retrieve the timestamp separately for each metric when you send multiple metrics from a single source. Don't use a variable to store the timestamp.

value

The value for the metric.

 Note

CloudWatch rejects values that are too small or too large. Values must be in the range of 8.515920e-109 to 1.174271e+108 (Base 10) or 2e-360 to 2e360 (Base 2). Special values (for example, NaN, +Infinity, -Infinity) are not supported.

Required: true

Type: double

unit

The unit of the metric.

Required: false

Type: string

Valid values: Seconds, Microseconds, Milliseconds, Bytes, Kilobytes, Megabytes, Gigabytes, Terabytes, Bits, Kilobits, Megabits, Gigabits, Terabits, Percent, Count, Bytes/Second, Kilobytes/Second, Megabytes/Second, Gigabytes/Second, Terabytes/Second, Bits/Second, Kilobits/Second, Megabits/Second, Gigabits/Second, Terabits/Second, Count/Second, None

Limits

All limits that are imposed by the CloudWatch [PutMetricData](#) API apply to metrics when using this connector. The following limits are especially important:

- 40 KB limit on API payload
- 20 metrics per API request
- 150 transactions per second (TPS) for the PutMetricData API

For more information, see [CloudWatch limits](#) in the *Amazon CloudWatch User Guide*.

Example input

```
{
  "request": {
    "namespace": "Greengrass",
    "metricData": {
      {
        "metricName": "latency",
        "dimensions": [
          {
            "name": "hostname",
            "value": "test_hostname"
          }
        ],
        "timestamp": 1539027324,
        "value": 123.0,
        "unit": "Seconds"
      }
    }
  }
}
```

Output data

This connector publishes status information as output data on an MQTT topic.

Topic filter in subscription

```
cloudwatch/metric/put/status
```

Example output: Success

The response includes the namespace of the metric data and the `RequestId` field from the CloudWatch response.

```
{
  "response": {
    "cloudwatch_rid": "70573243-d723-11e8-b095-75ff2EXAMPLE",
    "namespace": "Greengrass",
    "status": "success"
  }
}
```

Example output: Failure

```
{
  "response" : {
    "namespace": "Greengrass",
    "error": "InvalidInputException",
    "error_message": "cw metric is invalid",
    "status": "fail"
  }
}
```

Note

If the connector detects a retryable error (for example, connection errors), it retries the publish in the next batch.

Usage Example

Use the following high-level steps to set up an example Python 3.7 Lambda function that you can use to try out the connector.

Note

- If you use other Python runtimes, you can create a symlink from Python3.x to Python 3.7.
- The [Get started with connectors \(console\)](#) and [Get started with connectors \(CLI\)](#) topics contain detailed steps that show you how to configure and deploy an example Twilio Notifications connector.

1. Make sure you meet the [requirements](#) for the connector.

For the group role requirement, you must configure the role to grant the required permissions and make sure the role has been added to the group. For more information, see [the section called “Manage the group role \(console\)”](#) or [the section called “Manage the group role \(CLI\)”](#).

2. Create and publish a Lambda function that sends input data to the connector.

Save the [example code](#) as a PY file. Download and unzip the [AWS IoT Greengrass Core SDK for Python](#). Then, create a zip package that contains the PY file and the greengrasssdk folder at the root level. This zip package is the deployment package that you upload to AWS Lambda.

After you create the Python 3.7 Lambda function, publish a function version and create an alias.

3. Configure your Greengrass group.
 - a. Add the Lambda function by its alias (recommended). Configure the Lambda lifecycle as long-lived (or "Pinned": `true` in the CLI).
 - b. Add the connector and configure its [parameters](#).
 - c. Add subscriptions that allow the connector to receive [input data](#) and send [output data](#) on supported topic filters.
 - Set the Lambda function as the source, the connector as the target, and use a supported input topic filter.
 - Set the connector as the source, AWS IoT Core as the target, and use a supported output topic filter. You use this subscription to view status messages in the AWS IoT console.
4. Deploy the group.
5. In the AWS IoT console, on the **Test** page, subscribe to the output data topic to view status messages from the connector. The example Lambda function is long-lived and starts sending messages immediately after the group is deployed.

When you're finished testing, you can set the Lambda lifecycle to on-demand (or "Pinned": `false` in the CLI) and deploy the group. This stops the function from sending messages.

Example

The following example Lambda function sends an input message to the connector.

```
import greengrasssdk
import time
import json

iot_client = greengrasssdk.client('iot-data')
send_topic = 'cloudwatch/metric/put'
```

```
def create_request_with_all_fields():
    return {
        "request": {
            "namespace": "Greengrass_CW_Connector",
            "metricData": {
                "metricName": "Count1",
                "dimensions": [
                    {
                        "name": "test",
                        "value": "test"
                    }
                ],
                "value": 1,
                "unit": "Seconds",
                "timestamp": time.time()
            }
        }
    }

def publish_basic_message():
    messageToPublish = create_request_with_all_fields()
    print("Message To Publish: ", messageToPublish)
    iot_client.publish(topic=send_topic,
                      payload=json.dumps(messageToPublish))

publish_basic_message()

def lambda_handler(event, context):
    return
```

Licenses

The CloudWatch Metrics connector includes the following third-party software/licensing:

- [AWS SDK for Python \(Boto3\)](#)/Apache License 2.0
- [botocore](#)/Apache License 2.0
- [dateutil](#)/PSF License
- [docutils](#)/BSD License, GNU General Public License (GPL), Python Software Foundation License, Public Domain
- [jmespath](#)/MIT License
- [s3transfer](#)/Apache License 2.0

- [urllib3](#)/MIT License

This connector is released under the [Greengrass Core Software License Agreement](#).

Changelog

The following table describes the changes in each version of the connector.

Version	Changes
5	Fix to add support for duplicate timestamps in input data.
4	Added the <code>IsolationMode</code> parameter to configure the containerization mode for the connector.
3	Upgraded the Lambda runtime to Python 3.7, which changes the runtime requirement.
2	Fix to reduce excessive logging.
1	Initial release.

A Greengrass group can contain only one version of the connector at a time. For information about upgrading a connector version, see [the section called “Upgrading connector versions”](#).

See also

- [Integrate with services and protocols using connectors](#)
- [the section called “Get started with connectors \(console\)”](#)
- [the section called “Get started with connectors \(CLI\)”](#)
- [Using Amazon CloudWatch metrics](#) in the *Amazon CloudWatch User Guide*
- [PutMetricData](#) in the *Amazon CloudWatch API Reference*

Device Defender connector

The Device Defender [connector](#) notifies administrators of changes in the state of a Greengrass core device. This can help identify unusual behavior that might indicate a compromised device.

This connector reads system metrics from the `/proc` directory on the core device, and then publishes the metrics to AWS IoT Device Defender. For metrics reporting details, see [Device metrics document specification](#) in the *AWS IoT Developer Guide*.

This connector has the following versions.

Version	ARN
3	<code>arn:aws:greengrass: <i>region</i>::/connectors/DeviceDefender/versions/3</code>
2	<code>arn:aws:greengrass: <i>region</i>::/connectors/DeviceDefender/versions/2</code>
1	<code>arn:aws:greengrass: <i>region</i>::/connectors/DeviceDefender/versions/1</code>

For information about version changes, see the [Changelog](#).

Requirements

This connector has the following requirements:

Version 3

- AWS IoT Greengrass Core software v1.9.3 or later.
- [Python](#) version 3.7 or 3.8 installed on the core device and added to the PATH environment variable.

Note

To use Python 3.8, run the following command to create a symbolic link from the the default Python 3.7 installation folder to the installed Python 3.8 binaries.

```
sudo ln -s path-to-python-3.8/python3.8 /usr/bin/python3.7
```

This configures your device to meet the Python requirement for AWS IoT Greengrass.

- AWS IoT Device Defender configured to use the Detect feature to keep track of violations. For more information, see [Detect](#) in the *AWS IoT Developer Guide*.
- A [local volume resource](#) in the Greengrass group that points to the `/proc` directory. The resource must use the following properties:
 - Source path: `/proc`
 - Destination path: `/host_proc` (or a value that matches the [valid pattern](#))
 - `AutoAddGroupOwner`: `true`
- The [psutil](#) library installed on the Greengrass core. Version 5.7.0 is the latest version that is verified to work with the connector.
- The [cbor](#) library installed on the Greengrass core. Version 1.0.0 is the latest version that is verified to work with the connector.

Versions 1 - 2

- AWS IoT Greengrass Core software v1.7 or later.
- [Python](#) version 2.7 installed on the core device and added to the `PATH` environment variable.
- AWS IoT Device Defender configured to use the Detect feature to keep track of violations. For more information, see [Detect](#) in the *AWS IoT Developer Guide*.
- A [local volume resource](#) in the Greengrass group that points to the `/proc` directory. The resource must use the following properties:
 - Source path: `/proc`
 - Destination path: `/host_proc` (or a value that matches the [valid pattern](#))
 - `AutoAddGroupOwner`: `true`
- The [psutil](#) library installed on the Greengrass core.

- The [cbor](#) library installed on the Greengrass core.

Connector Parameters

This connector provides the following parameters:

SampleIntervalSeconds

The number of seconds between each cycle of gathering and reporting metrics. The minimum value is 300 seconds (5 minutes).

Display name in the AWS IoT console: **Metrics reporting interval**

Required: true

Type: string

Valid pattern: `^[0-9]*(?:3[0-9][0-9]|[4-9][0-9]{2}|[1-9][0-9]{3,})$`

ProcDestinationPath-ResourceId

The ID of the /proc volume resource.

Note

This connector is granted read-only access to the resource.

Display name in the AWS IoT console: **Resource for /proc directory**

Required: true

Type: string

Valid pattern: `[a-zA-Z0-9_-]+`

ProcDestinationPath

The destination path of the /proc volume resource.

Display name in the AWS IoT console: **Destination path of /proc resource**

Required: true

Type: string

Valid pattern: `\/[a-zA-Z0-9_-]+`

Create Connector Example (AWS CLI)

The following CLI command creates a `ConnectorDefinition` with an initial version that contains the Device Defender connector.

```
aws greengrass create-connector-definition --name MyGreengrassConnectors --initial-  
version '{  
  "Connectors": [  
    {  
      "Id": "MyDeviceDefenderConnector",  
      "ConnectorArn": "arn:aws:greengrass:region::/connectors/DeviceDefender/  
versions/3",  
      "Parameters": {  
        "SampleIntervalSeconds": "600",  
        "ProcDestinationPath": "/host_proc",  
        "ProcDestinationPath-ResourceId": "my-proc-resource"  
      }  
    }  
  ]  
}'
```

Note

The Lambda function in this connector has a [long-lived](#) lifecycle.

In the AWS IoT Greengrass console, you can add a connector from the group's **Connectors** page. For more information, see [the section called "Get started with connectors \(console\)"](#).

Input data

This connector doesn't accept MQTT messages as input data.

Output data

This connector publishes security metrics to AWS IoT Device Defender as output data.

Topic filter in subscription

`$aws/things/+/defender/metrics/json`

Note

This is the topic syntax that AWS IoT Device Defender expects. The connector replaces the + wildcard with the device name (for example, `$aws/things/thing-name/defender/metrics/json`).

Example output

For metrics reporting details, see [Device metrics document specification](#) in the *AWS IoT Developer Guide*.

```
{
  "header": {
    "report_id": 1529963534,
    "version": "1.0"
  },
  "metrics": {
    "listening_tcp_ports": {
      "ports": [
        {
          "interface": "eth0",
          "port": 24800
        },
        {
          "interface": "eth0",
          "port": 22
        },
        {
          "interface": "eth0",
          "port": 53
        }
      ],
      "total": 3
    },
    "listening_udp_ports": {
      "ports": [
        {
          "interface": "eth0",
```

```

        "port": 5353
      },
      {
        "interface": "eth0",
        "port": 67
      }
    ],
    "total": 2
  },
  "network_stats": {
    "bytes_in": 1157864729406,
    "bytes_out": 1170821865,
    "packets_in": 693092175031,
    "packets_out": 738917180
  },
  "tcp_connections": {
    "established_connections": {
      "connections": [
        {
          "local_interface": "eth0",
          "local_port": 80,
          "remote_addr": "192.168.0.1:8000"
        },
        {
          "local_interface": "eth0",
          "local_port": 80,
          "remote_addr": "192.168.0.1:8000"
        }
      ]
    },
    "total": 2
  }
}

```

Licenses

This connector is released under the [Greengrass Core Software License Agreement](#).

Changelog

The following table describes the changes in each version of the connector.

Version	Changes
3	Upgraded the Lambda runtime to Python 3.7, which changes the runtime requirement.
2	Fix to reduce excessive logging.
1	Initial release.

A Greengrass group can contain only one version of the connector at a time. For information about upgrading a connector version, see [the section called “Upgrading connector versions”](#).

See also

- [Integrate with services and protocols using connectors](#)
- [the section called “Get started with connectors \(console\)”](#)
- [the section called “Get started with connectors \(CLI\)”](#)
- [Device Defender](#) in the *AWS IoT Developer Guide*

Docker application deployment connector

The Greengrass Docker application deployment connector makes it easier to run your Docker images on an AWS IoT Greengrass core. The connector uses Docker Compose to start a multi-container Docker application from a `docker-compose.yml` file. Specifically, the connector runs `docker-compose` commands to manage Docker containers on a single core device. For more information, see [Overview of Docker Compose](#) in the Docker documentation. The connector can access Docker images stored in Docker container registries, such as Amazon Elastic Container Registry (Amazon ECR), Docker Hub, and private Docker trusted registries.

After you deploy the Greengrass group, the connector pulls the latest images and starts the Docker containers. It runs the `docker-compose pull` and `docker-compose up` command. Then, the connector publishes the status of the command to an [output MQTT topic](#). It also logs status information about running Docker containers. This makes it possible for you to monitor your application logs in Amazon CloudWatch. For more information, see [the section called “Monitoring with AWS IoT Greengrass logs”](#). The connector also starts Docker containers each

time the Greengrass daemon restarts. The number of Docker containers that can run on the core depends on your hardware.

The Docker containers run outside of the Greengrass domain on the core device, so they can't access the core's inter-process communication (IPC). However, you can configure some communication channels with Greengrass components, such as local Lambda functions. For more information, see [the section called "Communicating with Docker containers"](#).

You can use the connector for scenarios such as hosting a web server or MySQL server on your core device. Local services in your Docker applications can communicate with each other, other processes in the local environment, and cloud services. For example, you can run a web server on the core that sends requests from Lambda functions to a web service in the cloud.

This connector runs in [No container](#) isolation mode, so you can deploy it to a Greengrass group that runs without Greengrass containerization.

This connector has the following versions.

Version	ARN
7	arn:aws:greengrass: <i>region</i> ::/connectors/DockerApplicationDeployment/versions/7
6	arn:aws:greengrass: <i>region</i> ::/connectors/DockerApplicationDeployment/versions/6
5	arn:aws:greengrass: <i>region</i> ::/connectors/DockerApplicationDeployment/versions/5
4	arn:aws:greengrass: <i>region</i> ::/connectors/DockerApplicationDeployment/versions/4
3	arn:aws:greengrass: <i>region</i> ::/connectors/DockerApplicationDeployment/versions/3

Version	ARN
2	arn:aws:greengrass: <i>region</i> ::/connectors/DockerApplicationDeployment/versions/2
1	arn:aws:greengrass: <i>region</i> ::/connectors/DockerApplicationDeployment/versions/1

For information about version changes, see the [Changelog](#).

Requirements

This connector has the following requirements:

- AWS IoT Greengrass Core software v1.10 or later.

Note

This connector is not supported on OpenWrt distributions.

- [Python](#) version 3.7 or 3.8 installed on the core device and added to the PATH environment variable.

Note

To use Python 3.8, run the following command to create a symbolic link from the the default Python 3.7 installation folder to the installed Python 3.8 binaries.

```
sudo ln -s path-to-python-3.8/python3.8 /usr/bin/python3.7
```

This configures your device to meet the Python requirement for AWS IoT Greengrass.

- A minimum of 36 MB RAM on the Greengrass core for the connector to monitor running Docker containers. The total memory requirement depends on the number of Docker containers that run on the core.

- [Docker Engine](#) 1.9.1 or later installed on the Greengrass core. Version 19.0.3 is the latest version that is verified to work with the connector.

The docker executable must be in the `/usr/bin` or `/usr/local/bin` directory.

⚠ Important

We recommend that you install a credentials store to secure the local copies of your Docker credentials. For more information, see [the section called “Security notes”](#).

For information about installing Docker on Amazon Linux distributions, see [Docker basics for Amazon ECS](#) in the *Amazon Elastic Container Service Developer Guide*.

- [Docker Compose](#) installed on the Greengrass core. The `docker-compose` executable must be in the `/usr/bin` or `/usr/local/bin` directory.

The following Docker Compose versions are verified to work with the connector.

Connector version	Verified Docker Compose version
7	1.25.4
6	1.25.4
5	1.25.4
4	1.25.4
3	1.25.4
2	1.25.1
1	1.24.1

- A single Docker Compose file (for example, `docker-compose.yml`), stored in Amazon Simple Storage Service (Amazon S3). The format must be compatible with the version of Docker Compose installed on the core. You should test the file before you use it on your core. If you edit the file after you deploy the Greengrass group, you must redeploy the group to update your local copy on the core.

- A Linux user with permission to call the local Docker daemon and write to the directory that stores the local copy of your Compose file. For more information, see [Setting up the Docker user on the core](#).
- The [Greengrass group role](#) configured to allow the `s3:GetObject` action on the S3 bucket that contains your Compose file. This permission is shown in the following example IAM policy.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "AllowAccessToComposeFileS3Bucket",
      "Action": [
        "s3:GetObject",
        "s3:GetObjectVersion"
      ],
      "Effect": "Allow",
      "Resource": "arn:aws:s3:::bucket-name/*"
    }
  ]
}
```

Note

If your S3 bucket is versioning-enabled, then the role must be configured to allow the `s3:GetObjectVersion` action as well. For more information, see [Using versioning](#) in the *Amazon Simple Storage Service User Guide*.

For the group role requirement, you must configure the role to grant the required permissions and make sure the role has been added to the group. For more information, see [the section called “Manage the group role \(console\)”](#) or [the section called “Manage the group role \(CLI\)”](#).

- If your Docker Compose file references a Docker image stored in Amazon ECR, the [Greengrass group role](#) configured to allow the following:
 - `ecr:GetDownloadUrlForLayer` and `ecr:BatchGetImage` actions on your Amazon ECR repositories that contain the Docker images.
 - `ecr:GetAuthorizationToken` action on your resources.

Repositories must be in the same AWS account and AWS Region as the connector.

⚠ Important

Permissions in the group role can be assumed by all Lambda functions and connectors in the Greengrass group. For more information, see [the section called “Security notes”](#).

These permissions are shown in the following example policy.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "AllowGetEcrRepositories",
      "Effect": "Allow",
      "Action": [
        "ecr:GetDownloadUrlForLayer",
        "ecr:BatchGetImage"
      ],
      "Resource": [
        "arn:aws:ecr:region:account-id:repository/repository-name"
      ]
    },
    {
      "Sid": "AllowGetEcrAuthToken",
      "Effect": "Allow",
      "Action": "ecr:GetAuthorizationToken",
      "Resource": "*"
    }
  ]
}
```

For more information, see [Amazon ECR repository policy examples](#) in the *Amazon ECR User Guide*.

For the group role requirement, you must configure the role to grant the required permissions and make sure the role has been added to the group. For more information, see [the section called “Manage the group role \(console\)”](#) or [the section called “Manage the group role \(CLI\)”](#).

- If your Docker Compose file references a Docker image from [AWS Marketplace](#), the connector also has the following requirements:

- You must be subscribed to AWS Marketplace container products. For more information, see [Finding and subscribing to container products](#) in the *AWS Marketplace Subscribers Guide*.
- AWS IoT Greengrass must be configured to support local secrets, as described in [Secrets Requirements](#). The connector uses this feature only to retrieve your secrets from AWS Secrets Manager, not to store them.
- You must create a secret in Secrets Manager for each AWS Marketplace registry that stores a Docker image referenced in your Compose file. For more information, see [the section called "Accessing Docker images from private repositories"](#).
- If your Docker Compose file references a Docker image from private repositories in registries other than Amazon ECR, such as Docker Hub, the connector also has the following requirements:
 - AWS IoT Greengrass must be configured to support local secrets, as described in [Secrets Requirements](#). The connector uses this feature only to retrieve your secrets from AWS Secrets Manager, not to store them.
 - You must create a secret in Secrets Manager for each private repository that stores a Docker image referenced in your Compose file. For more information, see [the section called "Accessing Docker images from private repositories"](#).
- The Docker daemon must be running when you deploy a Greengrass group that contains this connector.

Accessing Docker images from private repositories

If you use credentials to access your Docker images, then you must allow the connector to access them. The way you do this depends on where the Docker image is located.

For Docker images stored Amazon ECR, you grant permission to get your authorization token in the Greengrass group role. For more information, see [the section called "Requirements"](#).

For Docker images stored in other private repositories or registries, you must create a secret in AWS Secrets Manager to store your login information. This includes Docker images that you subscribed to in AWS Marketplace. Create one secret for each repository. If you update your secrets in Secrets Manager, the changes propagate to the core the next time that you deploy the group.

Note

Secrets Manager is a service that you can use to securely store and manage your credentials, keys, and other secrets in the AWS Cloud. For more information, see [What is AWS Secrets Manager?](#) in the *AWS Secrets Manager User Guide*.

Each secret must contain the following keys:

Key	Value
username	The user name used to access the repository or registry.
password	The password used to access the repository or registry.
registryUrl	The endpoint of the registry. This must match the corresponding registry URL in the Compose file.

Note

To allow AWS IoT Greengrass to access a secret by default, the name of the secret must start with *greengrass-*. Otherwise, your Greengrass service role must grant access. For more information, see [the section called “Allow AWS IoT Greengrass to get secret values”](#).

To get login information for Docker images from AWS Marketplace

1. Get your password for Docker images from AWS Marketplace by using the `aws ecr get-login-password` command. For more information, see [get-login-password](#) in the *AWS CLI Command Reference*.

```
aws ecr get-login-password
```

2. Retrieve the registry URL for the Docker image. Open the AWS Marketplace website, and open the container product launch page. Under **Container Images**, choose **View container image details** to locate the user name and registry URL.

Use the retrieved user name, password, and registry URL to create a secret for each AWS Marketplace registry that stores Docker images referenced in your Compose file.

To create secrets (console)

In the AWS Secrets Manager console, choose **Other type of secrets**. Under **Specify the key-value pairs to be stored for this secret**, add rows for username, password, and registryUrl. For more information, see [Creating a basic secret](#) in the *AWS Secrets Manager User Guide*.

Specify the key/value pairs to be stored in this secret [Info](#)

Secret key/value	Plaintext	
username	Mary_Major	Remove
password	abc123xyz456	Remove
registryUrl	https://docker.io	Remove

[+ Add row](#)

To create secrets (CLI)

In the AWS CLI, use the Secrets Manager `create-secret` command, as shown in the following example. For more information, see [create-secret](#) in the *AWS CLI Command Reference*.

```
aws secretsmanager create-secret --name greengrass-MySecret --secret-string
[{"username": "Mary_Major"}, {"password": "abc123xyz456"}, {"registryUrl": "https://
docker.io"}]
```

Important

It is your responsibility to secure the `DockerComposeFileDestinationPath` directory that stores your Docker Compose file and the credentials for your Docker images from private repositories. For more information, see [the section called “Security notes”](#).

Parameters

This connector provides the following parameters:

Version 7

`DockerComposeFileS3Bucket`

The name of the S3 bucket that contains your Docker Compose file. When you create the bucket, make sure to follow the [rules for bucket names](#) described in the *Amazon Simple Storage Service User Guide*.

Display name in the AWS IoT console: **Docker Compose file in S3**

Note

In the console, the **Docker Compose file in S3** property combines the `DockerComposeFileS3Bucket`, `DockerComposeFileS3Key`, and `DockerComposeFileS3Version` parameters.

Required: true

Type: string

Valid pattern `[a-zA-Z0-9\\-\\.]{3,63}`

`DockerComposeFileS3Key`

The object key for your Docker Compose file in Amazon S3. For more information, including object key naming guidelines, see [Object key and metadata](#) in the *Amazon Simple Storage Service User Guide*.

Note

In the console, the **Docker Compose file in S3** property combines the `DockerComposeFileS3Bucket`, `DockerComposeFileS3Key`, and `DockerComposeFileS3Version` parameters.

Required: true

Type: string

Valid pattern .+

`DockerComposeFileS3Version`

The object version for your Docker Compose file in Amazon S3. For more information, including object key naming guidelines, see [Using versioning](#) in the *Amazon Simple Storage Service User Guide*.

Note

In the console, the **Docker Compose file in S3** property combines the `DockerComposeFileS3Bucket`, `DockerComposeFileS3Key`, and `DockerComposeFileS3Version` parameters.

Required: false

Type: string

Valid pattern .+

`DockerComposeFileDestinationPath`

The absolute path of the local directory used to store a copy of the Docker Compose file. This must be an existing directory. The user specified for `DockerUserId` must have permission to create a file in this directory. For more information, see [the section called "Setting up the Docker user on the core"](#).

⚠ Important

This directory stores your Docker Compose file and the credentials for your Docker images from private repositories. It is your responsibility to secure this directory. For more information, see [the section called “Security notes”](#).

Display name in the AWS IoT console: **Directory path for local Compose file**

Required: true

Type: string

Valid pattern `\/.*\/?`

Example: `/home/username/myCompose`

DockerUserId

The UID of the Linux user that the connector runs as. This user must belong to the `docker` Linux group on the core device and have write permissions to the `DockerComposeFileDestinationPath` directory. For more information, see [Setting up the Docker user on the core](#).

ℹ Note

We recommend that you avoid running as root unless absolutely necessary. If you do specify the root user, you must allow Lambda functions to run as root on the AWS IoT Greengrass core. For more information, see [the section called “Running a Lambda function as root”](#).

Display name in the AWS IoT console: **Docker user ID**

Required: false

Type: string

Valid pattern: `^[0-9]{1,5}$`

AWSecretsArnList

The Amazon Resource Names (ARNs) of the secrets in AWS Secrets Manager that contain the login information used to access your Docker images in private repositories. For more information, see [the section called "Accessing Docker images from private repositories"](#).

Display name in the AWS IoT console: **Credentials for private repositories**

Required: false. This parameter is required to access Docker images stored in private repositories.

Type: array of string

Valid pattern: `[( ? , ? ? "(arn:(aws(-[a-z]+)):secretsmanager:[a-z0-9-]+:[0-9]{12}:secret:([a-zA-Z0-9\\]+/)[a-zA-Z0-9/_+=, .@-]+-[a-zA-Z0-9]+)"]]`

DockerContainerStatusLogFrequency

The frequency (in seconds) at which the connector logs status information about the Docker containers running on the core. The default is 300 seconds (5 minutes).

Display name in the AWS IoT console: **Logging frequency**

Required: false

Type: string

Valid pattern: `^[1-9]{1}[0-9]{0,3}$`

ForceDeploy

Indicates whether to force the Docker deployment if it fails because of the improper cleanup of the last deployment. The default value is False.

Display name in the AWS IoT console: **Force deployment**

Required: false

Type: string

Valid pattern: `^(true|false)$`

DockerPullBeforeUp

Indicates whether the deployer should run `docker-compose pull` before running `docker-compose up` for a pull-down-up behavior. The default value is `True`.

Display name in the AWS IoT console: **Docker Pull Before Up**

Required: `false`

Type: `string`

Valid pattern: `^(true|false)$`

StopContainersOnNewDeployment

Indicates whether the connector should stop Docker Deployer managed docker containers when GGC is stopped (GGC stops when a new group is deployed, or the kernel is shut down). The default value is `True`.

Display name in the AWS IoT console: **Docker stop on new deployment**

Note

We recommend keeping this parameter set to its default `True` value. The parameter to `False` causes your Docker container to continue running even after terminating the AWS IoT Greengrass core or starting a new deployment. If you set this parameter to `False`, you must ensure that your Docker containers are maintained as necessary in the event of a `docker-compose` service name change or addition. For more information, see the `docker-compose` compose file documentation.

Required: `false`

Type: `string`

Valid pattern: `^(true|false)$`

DockerOfflineMode

Indicates whether to use the existing Docker Compose file when AWS IoT Greengrass starts offline. The default value is `False`.

Required: false

Type: string

Valid pattern: `^(true|false)$`

Version 6

DockerComposeFileS3Bucket

The name of the S3 bucket that contains your Docker Compose file. When you create the bucket, make sure to follow the [rules for bucket names](#) described in the *Amazon Simple Storage Service User Guide*.

Display name in the AWS IoT console: **Docker Compose file in S3**

Note

In the console, the **Docker Compose file in S3** property combines the `DockerComposeFileS3Bucket`, `DockerComposeFileS3Key`, and `DockerComposeFileS3Version` parameters.

Required: true

Type: string

Valid pattern `[a-zA-Z0-9\\-\\.]{3,63}`

DockerComposeFileS3Key

The object key for your Docker Compose file in Amazon S3. For more information, including object key naming guidelines, see [Object key and metadata](#) in the *Amazon Simple Storage Service User Guide*.

Note

In the console, the **Docker Compose file in S3** property combines the `DockerComposeFileS3Bucket`, `DockerComposeFileS3Key`, and `DockerComposeFileS3Version` parameters.

Required: true

Type: string

Valid pattern .+

DockerComposeFileS3Version

The object version for your Docker Compose file in Amazon S3. For more information, including object key naming guidelines, see [Using versioning](#) in the *Amazon Simple Storage Service User Guide*.

 Note

In the console, the **Docker Compose file in S3** property combines the DockerComposeFileS3Bucket, DockerComposeFileS3Key, and DockerComposeFileS3Version parameters.

Required: false

Type: string

Valid pattern .+

DockerComposeFileDestinationPath

The absolute path of the local directory used to store a copy of the Docker Compose file. This must be an existing directory. The user specified for DockerUserId must have permission to create a file in this directory. For more information, see [the section called "Setting up the Docker user on the core"](#).

 Important

This directory stores your Docker Compose file and the credentials for your Docker images from private repositories. It is your responsibility to secure this directory. For more information, see [the section called "Security notes"](#).

Display name in the AWS IoT console: **Directory path for local Compose file**

Required: true

Type: string

Valid pattern `\/.*\/?`

Example: `/home/username/myCompose`

DockerUserId

The UID of the Linux user that the connector runs as. This user must belong to the `docker` Linux group on the core device and have write permissions to the `DockerComposeFileDestinationPath` directory. For more information, see [Setting up the Docker user on the core](#).

Note

We recommend that you avoid running as root unless absolutely necessary. If you do specify the root user, you must allow Lambda functions to run as root on the AWS IoT Greengrass core. For more information, see [the section called “Running a Lambda function as root”](#).

Display name in the AWS IoT console: **Docker user ID**

Required: false

Type: string

Valid pattern: `^[0-9]{1,5}$`

AWSecretsArnList

The Amazon Resource Names (ARNs) of the secrets in AWS Secrets Manager that contain the login information used to access your Docker images in private repositories. For more information, see [the section called “Accessing Docker images from private repositories”](#).

Display name in the AWS IoT console: **Credentials for private repositories**

Required: false. This parameter is required to access Docker images stored in private repositories.

Type: array of string

Valid pattern: `[( ? , ? ? "(arn:(aws(-[a-z]+)):secretsmanager:[a-z0-9-]+:[0-9]{12}:secret:([a-zA-Z0-9\\]+)/[a-zA-Z0-9/_+=, .@-]+-[a-zA-Z0-9]+)")]`

DockerContainerStatusLogFrequency

The frequency (in seconds) at which the connector logs status information about the Docker containers running on the core. The default is 300 seconds (5 minutes).

Display name in the AWS IoT console: **Logging frequency**

Required: false

Type: string

Valid pattern: `^[1-9]{1}[0-9]{0,3}$`

ForceDeploy

Indicates whether to force the Docker deployment if it fails because of the improper cleanup of the last deployment. The default value is False.

Display name in the AWS IoT console: **Force deployment**

Required: false

Type: string

Valid pattern: `^(true|false)$`

DockerPullBeforeUp

Indicates whether the deployer should run `docker-compose pull` before running `docker-compose up` for a pull-down-up behavior. The default value is True.

Display name in the AWS IoT console: **Docker Pull Before Up**

Required: false

Type: string

Valid pattern: `^(true|false)$`

StopContainersOnNewDeployment

Indicates whether the connector should stop Docker Deployer managed docker containers when GGC is stopped (when a new group deployment is made, or the kernel is shutdown). The default value is `True`.

Display name in the AWS IoT console: **Docker stop on new deployment**

Note

We recommend keeping this parameter set to its default `True` value. The parameter to `False` causes your Docker container to continue running even after terminating the AWS IoT Greengrass core or starting a new deployment. If you set this parameter to `False`, you must ensure that your Docker containers are maintained as necessary in the event of a `docker-compose` service name change or addition. For more information, see the `docker-compose` compose file documentation.

Required: `false`

Type: `string`

Valid pattern: `^(true|false)$`

Version 5

DockerComposeFileS3Bucket

The name of the S3 bucket that contains your Docker Compose file. When you create the bucket, make sure to follow the [rules for bucket names](#) described in the *Amazon Simple Storage Service User Guide*.

Display name in the AWS IoT console: **Docker Compose file in S3**

Note

In the console, the **Docker Compose file in S3** property combines the `DockerComposeFileS3Bucket`, `DockerComposeFileS3Key`, and `DockerComposeFileS3Version` parameters.

Required: true

Type: string

Valid pattern [a-zA-Z0-9\\-\\.]{3,63}

DockerComposeFileS3Key

The object key for your Docker Compose file in Amazon S3. For more information, including object key naming guidelines, see [Object key and metadata](#) in the *Amazon Simple Storage Service User Guide*.

Note

In the console, the **Docker Compose file in S3** property combines the DockerComposeFileS3Bucket, DockerComposeFileS3Key, and DockerComposeFileS3Version parameters.

Required: true

Type: string

Valid pattern .+

DockerComposeFileS3Version

The object version for your Docker Compose file in Amazon S3. For more information, including object key naming guidelines, see [Using versioning](#) in the *Amazon Simple Storage Service User Guide*.

Note

In the console, the **Docker Compose file in S3** property combines the DockerComposeFileS3Bucket, DockerComposeFileS3Key, and DockerComposeFileS3Version parameters.

Required: false

Type: string

Valid pattern .+

DockerComposeFileDestinationPath

The absolute path of the local directory used to store a copy of the Docker Compose file. This must be an existing directory. The user specified for `DockerUserId` must have permission to create a file in this directory. For more information, see [the section called “Setting up the Docker user on the core”](#).

Important

This directory stores your Docker Compose file and the credentials for your Docker images from private repositories. It is your responsibility to secure this directory. For more information, see [the section called “Security notes”](#).

Display name in the AWS IoT console: **Directory path for local Compose file**

Required: true

Type: string

Valid pattern `\/.*\/?`

Example: `/home/username/myCompose`

DockerUserId

The UID of the Linux user that the connector runs as. This user must belong to the `docker` Linux group on the core device and have write permissions to the `DockerComposeFileDestinationPath` directory. For more information, see [Setting up the Docker user on the core](#).

Note

We recommend that you avoid running as root unless absolutely necessary. If you do specify the root user, you must allow Lambda functions to run as root on the AWS IoT Greengrass core. For more information, see [the section called “Running a Lambda function as root”](#).

Display name in the AWS IoT console: **Docker user ID**

Required: false

Type: string

Valid pattern: `^[0-9]{1,5}$`

`AWSecretsArnList`

The Amazon Resource Names (ARNs) of the secrets in AWS Secrets Manager that contain the login information used to access your Docker images in private repositories. For more information, see [the section called “Accessing Docker images from private repositories”](#).

Display name in the AWS IoT console: **Credentials for private repositories**

Required: false. This parameter is required to access Docker images stored in private repositories.

Type: array of string

Valid pattern: `[( ?, ? ?"(arn:(aws(-[a-z]+)):secretsmanager:[a-z0-9-]+:[0-9]{12}:secret:([a-zA-Z0-9\+\/][a-zA-Z0-9/_+=, .@-]+-[a-zA-Z0-9]+)"))]`

`DockerContainerStatusLogFrequency`

The frequency (in seconds) at which the connector logs status information about the Docker containers running on the core. The default is 300 seconds (5 minutes).

Display name in the AWS IoT console: **Logging frequency**

Required: false

Type: string

Valid pattern: `^[1-9]{1}[0-9]{0,3}$`

`ForceDeploy`

Indicates whether to force the Docker deployment if it fails because of the improper cleanup of the last deployment. The default value is `False`.

Display name in the AWS IoT console: **Force deployment**

Required: false

Type: string

Valid pattern: `^(true|false)$`

DockerPullBeforeUp

Indicates whether the deployer should run `docker-compose pull` before running `docker-compose up` for a pull-down-up behavior. The default value is `True`.

Display name in the AWS IoT console: **Docker Pull Before Up**

Required: false

Type: string

Valid pattern: `^(true|false)$`

Versions 2 - 4

DockerComposeFileS3Bucket

The name of the S3 bucket that contains your Docker Compose file. When you create the bucket, make sure to follow the [rules for bucket names](#) described in the *Amazon Simple Storage Service User Guide*.

Display name in the AWS IoT console: **Docker Compose file in S3**

Note

In the console, the **Docker Compose file in S3** property combines the `DockerComposeFileS3Bucket`, `DockerComposeFileS3Key`, and `DockerComposeFileS3Version` parameters.

Required: true

Type: string

Valid pattern `[a-zA-Z0-9\\-\\.]{3,63}`

DockerComposeFileS3Key

The object key for your Docker Compose file in Amazon S3. For more information, including object key naming guidelines, see [Object key and metadata](#) in the *Amazon Simple Storage Service User Guide*.

Note

In the console, the **Docker Compose file in S3** property combines the `DockerComposeFileS3Bucket`, `DockerComposeFileS3Key`, and `DockerComposeFileS3Version` parameters.

Required: true

Type: string

Valid pattern .+

DockerComposeFileS3Version

The object version for your Docker Compose file in Amazon S3. For more information, including object key naming guidelines, see [Using versioning](#) in the *Amazon Simple Storage Service User Guide*.

Note

In the console, the **Docker Compose file in S3** property combines the `DockerComposeFileS3Bucket`, `DockerComposeFileS3Key`, and `DockerComposeFileS3Version` parameters.

Required: false

Type: string

Valid pattern .+

DockerComposeFileDestinationPath

The absolute path of the local directory used to store a copy of the Docker Compose file. This must be an existing directory. The user specified for `DockerUserId` must have

permission to create a file in this directory. For more information, see [the section called “Setting up the Docker user on the core”](#).

 Important

This directory stores your Docker Compose file and the credentials for your Docker images from private repositories. It is your responsibility to secure this directory. For more information, see [the section called “Security notes”](#).

Display name in the AWS IoT console: **Directory path for local Compose file**

Required: true

Type: string

Valid pattern `\/.*\/?`

Example: `/home/username/myCompose`

DockerUserId

The UID of the Linux user that the connector runs as. This user must belong to the docker Linux group on the core device and have write permissions to the DockerComposeFileDestinationPath directory. For more information, see [Setting up the Docker user on the core](#).

 Note

We recommend that you avoid running as root unless absolutely necessary. If you do specify the root user, you must allow Lambda functions to run as root on the AWS IoT Greengrass core. For more information, see [the section called “Running a Lambda function as root”](#).

Display name in the AWS IoT console: **Docker user ID**

Required: false

Type: string

Valid pattern: `^[0-9]{1,5}$`

`AWSecretsArnList`

The Amazon Resource Names (ARNs) of the secrets in AWS Secrets Manager that contain the login information used to access your Docker images in private repositories. For more information, see [the section called "Accessing Docker images from private repositories"](#).

Display name in the AWS IoT console: **Credentials for private repositories**

Required: `false`. This parameter is required to access Docker images stored in private repositories.

Type: `array of string`

Valid pattern: `[( ? , ? ? "(arn:(aws(-[a-z]+)):secretsmanager:[a-z0-9-]+:[0-9]{12}:secret:([a-zA-Z0-9\\]+/)[a-zA-Z0-9/_+=, .@-]+-[a-zA-Z0-9]+)"]`

`DockerContainerStatusLogFrequency`

The frequency (in seconds) at which the connector logs status information about the Docker containers running on the core. The default is 300 seconds (5 minutes).

Display name in the AWS IoT console: **Logging frequency**

Required: `false`

Type: `string`

Valid pattern: `^[1-9]{1}[0-9]{0,3}$`

`ForceDeploy`

Indicates whether to force the Docker deployment if it fails because of the improper cleanup of the last deployment. The default value is `False`.

Display name in the AWS IoT console: **Force deployment**

Required: `false`

Type: `string`

Valid pattern: `^(true|false)$`

Version 1

DockerComposeFileS3Bucket

The name of the S3 bucket that contains your Docker Compose file. When you create the bucket, make sure to follow the [rules for bucket names](#) described in the *Amazon Simple Storage Service User Guide*.

Display name in the AWS IoT console: **Docker Compose file in S3**

Note

In the console, the **Docker Compose file in S3** property combines the `DockerComposeFileS3Bucket`, `DockerComposeFileS3Key`, and `DockerComposeFileS3Version` parameters.

Required: true

Type: string

Valid pattern `[a-zA-Z0-9\\-\\.]{3,63}`

DockerComposeFileS3Key

The object key for your Docker Compose file in Amazon S3. For more information, including object key naming guidelines, see [Object key and metadata](#) in the *Amazon Simple Storage Service User Guide*.

Note

In the console, the **Docker Compose file in S3** property combines the `DockerComposeFileS3Bucket`, `DockerComposeFileS3Key`, and `DockerComposeFileS3Version` parameters.

Required: true

Type: string

Valid pattern .+

DockerComposeFileS3Version

The object version for your Docker Compose file in Amazon S3. For more information, including object key naming guidelines, see [Using versioning](#) in the *Amazon Simple Storage Service User Guide*.

 Note

In the console, the **Docker Compose file in S3** property combines the DockerComposeFileS3Bucket, DockerComposeFileS3Key, and DockerComposeFileS3Version parameters.

Required: false

Type: string

Valid pattern .+

DockerComposeFileDestinationPath

The absolute path of the local directory used to store a copy of the Docker Compose file. This must be an existing directory. The user specified for DockerUserId must have permission to create a file in this directory. For more information, see [the section called "Setting up the Docker user on the core"](#).

 Important

This directory stores your Docker Compose file and the credentials for your Docker images from private repositories. It is your responsibility to secure this directory. For more information, see [the section called "Security notes"](#).

Display name in the AWS IoT console: **Directory path for local Compose file**

Required: true

Type: string

Valid pattern `\/.*\/?`

Example: `/home/username/myCompose`

DockerUserId

The UID of the Linux user that the connector runs as. This user must belong to the `docker` Linux group on the core device and have write permissions to the `DockerComposeFileDestinationPath` directory. For more information, see [Setting up the Docker user on the core](#).

 Note

We recommend that you avoid running as root unless absolutely necessary. If you do specify the root user, you must allow Lambda functions to run as root on the AWS IoT Greengrass core. For more information, see [the section called “Running a Lambda function as root”](#).

Display name in the AWS IoT console: **Docker user ID**

Required: false

Type: string

Valid pattern: `^[0-9]{1,5}$`

AWSecretsArnList

The Amazon Resource Names (ARNs) of the secrets in AWS Secrets Manager that contain the login information used to access your Docker images in private repositories. For more information, see [the section called “Accessing Docker images from private repositories”](#).

Display name in the AWS IoT console: **Credentials for private repositories**

Required: false. This parameter is required to access Docker images stored in private repositories.

Type: array of string

Valid pattern: `[( ?,? ?"(arn:(aws(-[a-z]+)):secretsmanager:[a-z0-9-]+:[0-9]{12}:secret:([a-zA-Z0-9\]+/)[a-zA-Z0-9/_+=, .@-]+-[a-zA-Z0-9]+)")]`

DockerContainerStatusLogFrequency

The frequency (in seconds) at which the connector logs status information about the Docker containers running on the core. The default is 300 seconds (5 minutes).

Display name in the AWS IoT console: **Logging frequency**

Required: false

Type: string

Valid pattern: `^[1-9]{1}[0-9]{0,3}$`

Create Connector Example (AWS CLI)

The following CLI command creates a `ConnectorDefinition` with an initial version that contains the Greengrass Docker application deployment connector.

```
aws greengrass create-connector-definition --name MyGreengrassConnectors --initial-
version '{
  "Connectors": [
    {
      "Id": "MyDockerAppplicationDeploymentConnector",
      "ConnectorArn": "arn:aws:greengrass:region::/connectors/
DockerApplicationDeployment/versions/5",
      "Parameters": {
        "DockerComposeFileS3Bucket": "amzn-s3-demo-bucket",
        "DockerComposeFileS3Key": "production-docker-compose.yml",
        "DockerComposeFileS3Version": "123",
        "DockerComposeFileDestinationPath": "/home/username/myCompose",
        "DockerUserId": "1000",
        "AWSecretsArnList": "[\"arn:aws:secretsmanager:region:account-
id:secret:greengrass-secret1-hash\", \"arn:aws:secretsmanager:region:account-
id:secret:greengrass-secret2-hash\"]",
        "DockerContainerStatusLogFrequency": "30",
        "ForceDeploy": "True",
        "DockerPullBeforeUp": "True"
      }
    }
  ]
}
```

```
]
}'
```

Note

The Lambda function in this connector has a [long-lived](#) lifecycle.

Input data

This connector doesn't require or accept input data.

Output data

This connector publishes the status of the `docker-compose up` command as output data.

Topic filter in subscription

```
dockerapplicationdeploymentconnector/message/status
```

Example output: Success

```
{
  "status": "success",
  "GreengrassDockerApplicationDeploymentStatus": "Successfully triggered docker-
compose up",
  "S3Bucket": "amzn-s3-demo-bucket",
  "ComposeFileName": "production-docker-compose.yml",
  "ComposeFileVersion": "123"
}
```

Example output: Failure

```
{
  "status": "fail",
  "error_message": "description of error",
  "error": "InvalidParameter"
}
```

The error type can be `InvalidParameter` or `InternalError`.

Setting up the Docker user on the AWS IoT Greengrass core

The Greengrass Docker application deployment connector runs as the user you specify for the `DockerUserId` parameter. If you don't specify a value, the connector runs as `ggc_user`, which is the default Greengrass access identity.

To allow the connector to interact with the Docker daemon, the Docker user must belong to the `docker` Linux group on the core. The Docker user must also have write permissions to the `DockerComposeFileDestinationPath` directory. This is where the connector stores your local `docker-compose.yml` file and Docker credentials.

Note

- We recommend that you create a Linux user instead of using the default `ggc_user`. Otherwise, any Lambda function in the Greengrass group can access the Compose file and Docker credentials.
- We recommend that you avoid running as root unless absolutely necessary. If you do specify the root user, you must allow Lambda functions to run as root on the AWS IoT Greengrass core. For more information, see [the section called “Running a Lambda function as root”](#).

1. Create the user. You can run the `useradd` command and include the optional `-u` option to assign a UID. For example:

```
sudo useradd -u 1234 user-name
```

2. Add the user to the `docker` group on the core. For example:

```
sudo usermod -aG docker user-name
```

For more information, including how to create the `docker` group, see [Manage Docker as a non-root user](#) in the Docker documentation.

3. Give the user permissions to write to the directory specified for the `DockerComposeFileDestinationPath` parameter. For example:
 - a. To set the user as the owner of the directory. This example uses the UID from step 1.

```
chown 1234 docker-compose-file-destination-path
```

- b. To give read and write permissions to the owner.

```
chmod 700 docker-compose-file-destination-path
```

For more information, see [How To Manage File And Folder Permissions In Linux](#) in the Linux Foundation documentation.

- c. If you didn't assign a UID when you created the user, or if you used an existing user, run the `id` command to look up the UID.

```
id -u user-name
```

You use the UID to configure the `DockerUserId` parameter for the connector.

Usage information

When you use the Greengrass Docker application deployment connector, you should be aware of the following implementation-specific usage information.

- **Fixed prefix for project names.** The connector prepends the `greengrassdockerapplicationdeployment` prefix to the names of the Docker containers that it starts. The connector uses this prefix as the project name in the `docker-compose` commands that it runs.
- **Logging behavior.** The connector writes status information and troubleshooting information to a log file. You can configure AWS IoT Greengrass to send logs to CloudWatch Logs and to write logs locally. For more information, see [the section called "Logging"](#). This is the path to the local log for the connector:

```
/greengrass-root/ggc/var/log/user/region/aws/DockerApplicationDeployment.log
```

You must have root permissions to access local logs.

- **Updating Docker images.** Docker caches images on the core device. If you update a Docker image and want to propagate the change to the core device, make sure to change the tag for the image in the Compose file. Changes take effect after the Greengrass group is deployed.

- **10-minute timeout for cleanup operations.** When the Greengrass daemon stops during a restart, the `docker-compose down` command is initiated. All Docker containers have a maximum of 10 minutes after `docker-compose down` is initiated to perform any cleanup operations. If the cleanup isn't completed in 10 minutes, you must clean up the remaining containers manually. For more information, see [docker rm](#) in the Docker CLI documentation.
- **Running Docker commands.** To troubleshoot issues, you can run Docker commands in a terminal window on the core device. For example, run the following command to see the Docker containers that were started by the connector:

```
docker ps --filter name="greengrassdockerapplicationdeployment"
```

- **Reserved resource ID.** The connector uses the `DOCKER_DEPLOYER_SECRET_RESOURCE_RESERVED_ID_`*index* ID for the Greengrass resources it creates in the Greengrass group. Resource IDs must be unique in the group, so don't assign a resource ID that might conflict with this reserved resource ID.
- **Offline mode.** When you set the `DockerOfflineMode` configuration parameter to `True`, then the Docker connector is able to operate in *offline mode*. This can happen when a Greengrass group deployment restarts while the core device is offline, and the connector cannot establish a connection to Amazon S3 or Amazon ECR to retrieve the Docker Compose file.

With offline mode enabled, the connector attempts to download your Compose file, and run `docker login` commands as it would for a normal restart. If these attempts fail, then the connector looks for a locally stored Compose file in the folder that was specified using the `DockerComposeFileDestinationPath` parameter. If a local Compose file exists, then the connector follows the normal sequence of `docker-compose` commands and pulls from local images. If the Compose file or the local images are not present, then the connector fails. The behavior of the `ForceDeploy` and `StopContainersOnNewDeployment` parameters remains the same in offline mode.

Communicating with Docker containers

AWS IoT Greengrass supports the following communication channels between Greengrass components and Docker containers:

- Greengrass Lambda functions can use REST APIs to communicate with processes in Docker containers. You can set up a server in a Docker container that opens a port. Lambda functions can communicate with the container on this port.

- Processes in Docker containers can exchange MQTT messages through the local Greengrass message broker. You can set up the Docker container as a client device in the Greengrass group and then create subscriptions to allow the container to communicate with Greengrass Lambda functions, client devices, and other connectors in the group, or with AWS IoT and the local shadow service. For more information, see [the section called “Configure MQTT communication with Docker containers”](#).
- Greengrass Lambda functions can update a shared file to pass information to Docker containers. You can use the Compose file to bind mount the shared file path for a Docker container.

Configure MQTT communication with Docker containers

You can configure a Docker container as a client device and add it to a Greengrass group. Then, you can create subscriptions that allow MQTT communication between the Docker container and Greengrass components or AWS IoT. In the following procedure, you create a subscription that allows the Docker container device to receive shadow update messages from the local shadow service. You can follow this pattern to create other subscriptions.

Note

This procedure assumes that you have already created a Greengrass group and a Greengrass core (v1.10 or later). For information about creating a Greengrass group and core, see [Getting started with AWS IoT Greengrass](#).

To configure a Docker container as a client device and add it to a Greengrass group

1. Create a folder on the core device to store the certificates and keys used to authenticate the Greengrass device.

The file path must be mounted on the Docker container you want to start. The following snippet shows how to mount a file path in your Compose file. In this example, *path-to-device-certs* represents the folder you created in this step.

```
version: '3.3'
services:
  myService:
    image: user-name/repo:image-tag
    volumes:
```



```
- /path-to-device-certs/:/path-accessible-in-container
```

2. In the AWS IoT console navigation pane, under **Manage**, expand **Greengrass devices**, and then choose **Groups (V1)**.
3. Choose the target group.
4. On the group configuration page, choose **Client devices**, and then choose **Associate**.
5. In the **Associate a client device with this group** modal, choose **Create new AWS IoT thing**.

The **Create things** page opens in a new tab.

6. On the **Create things** page, choose **Create single thing**, and then choose **Next**.
7. On the **Specify thing properties** page, enter a name for the device, and then choose **Next**.
8. On the **Configure device certificate** page, choose **Next**.
9. On the **Attach policies to certificate** page, do one of the following:
 - Select an existing policy that grants permissions that client devices require, and then choose **Create thing**.

A modal opens where you can download the certificates and keys that the device uses to connect to the AWS Cloud and the core.

- Create and attach a new policy that grants client device permissions. Do the following:
 - a. Choose **Create policy**.

The **Create policy** page opens in a new tab.

- b. On the **Create policy** page, do the following:
 - i. For **Policy name**, enter a name that describes the policy, such as **GreengrassV1ClientDevicePolicy**.
 - ii. On the **Policy statements** tab, under **Policy document**, choose **JSON**.
 - iii. Enter the following policy document. This policy allows the client device to discover Greengrass cores and communicate on all MQTT topics. For information about how to restrict this policy's access, see [Device authentication and authorization for AWS IoT Greengrass](#).

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
```

```
    "Effect": "Allow",
    "Action": [
        "iot:Publish",
        "iot:Subscribe",
        "iot:Connect",
        "iot:Receive"
    ],
    "Resource": [
        "*"
    ]
},
{
    "Effect": "Allow",
    "Action": [
        "greengrass:*"
    ],
    "Resource": [
        "*"
    ]
}
]
```

- iv. Choose **Create** to create the policy.
- c. Return to the browser tab with the **Attach policies to certificate** page open. Do the following:
 - i. In the **Policies** list, select the policy that you created, such as **GreengrassV1ClientDevicePolicy**.

If you don't see the policy, choose the refresh button.

- ii. Choose **Create thing**.

A modal opens where you can download the certificates and keys that the device uses to connect to the AWS Cloud and the core.

10. In the **Download certificates and keys** modal, download the device's certificates.

Important

Before you choose **Done**, download the security resources.

Do the following:

- a. For **Device certificate**, choose **Download** to download the device certificate.
- b. For **Public key file**, choose **Download** to download the public key for the certificate.
- c. For **Private key file**, choose **Download** to download the private key file for the certificate.
- d. Review [Server Authentication](#) in the *AWS IoT Developer Guide* and choose the appropriate root CA certificate. We recommend that you use Amazon Trust Services (ATS) endpoints and ATS root CA certificates. Under **Root CA certificates**, choose **Download** for a root CA certificate.
- e. Choose **Done**.

Make a note of the certificate ID that's common in the file names for the device certificate and keys. You need it later.

11. Copy the certificates and keys into the folder that you created in step 1.

Next, create a subscription in the group. For this example, you create a subscription allows the Docker container device to receive MQTT messages from the local shadow service.

 Note

The maximum size of a shadow document is 8 KB. For more information, see [AWS IoT quotas](#) in the *AWS IoT Developer Guide*.

To create a subscription that allows the Docker container device to receive MQTT messages from the local shadow service

1. On the group configuration page, choose the **Subscriptions** tab, and then choose **Add Subscription**.
2. On the **Select your source and target** page, configure the source and target, as follows:
 - a. For **Select a source**, choose **Services**, and then choose **Local Shadow Service**.
 - b. For **Select a target**, choose **Devices**, and then choose your device.
 - c. Choose **Next**.

- d. On the **Filter your data with a topic** page, for **Topic filter**, choose **\$aws/things/MyDockerDevice/shadow/update/accepted**, and then choose **Next**. Replace **MyDockerDevice** with the name of the device that you created earlier.
- e. Choose **Finish**.

Include the following code snippet in the Docker image that you reference in your Compose file. This is the Greengrass device code. Also, add code in your Docker container that starts the Greengrass device inside the container. It can run as a separate process in the image or in a separate thread.

```
import os
import sys
import time
import uuid

from AWSIoTPythonSDK.core.greengrass.discovery.providers import DiscoveryInfoProvider
from AWSIoTPythonSDK.exception.AWSIoTExceptions import DiscoveryInvalidRequestException
from AWSIoTPythonSDK.MQTTLib import AWSIoTMQTTClient

# Replace thingName with the name you registered for the Docker device.
thingName = "MyDockerDevice"
clientId = thingName

# Replace host with the IoT endpoint for your &AWS-account;.
host = "myPrefix.iot.region.amazonaws.com"

# Replace topic with the topic where the Docker container subscribes.
topic = "$aws/things/MyDockerDevice/shadow/update/accepted"

# Replace these paths based on the download location of the certificates for the Docker
  container.
rootCAPath = "/path-accessible-in-container/AmazonRootCA1.pem"
certificatePath = "/path-accessible-in-container/certId-certificate.pem.crt"
privateKeyPath = "/path-accessible-in-container/certId-private.pem.key"

# Discover Greengrass cores.
discoveryInfoProvider = DiscoveryInfoProvider()
discoveryInfoProvider.configureEndpoint(host)
discoveryInfoProvider.configureCredentials(rootCAPath, certificatePath, privateKeyPath)
discoveryInfoProvider.configureTimeout(10) # 10 seconds.
```

```
GROUP_CA_PATH = "./groupCA/"
MQTT_QOS = 1

discovered = False
groupCA = None
coreInfo = None

try:
    # Get discovery info from AWS IoT.
    discoveryInfo = discoveryInfoProvider.discover(thingName)
    caList = discoveryInfo.getAllCas()
    coreList = discoveryInfo.getAllCores()

    # Use first discovery result.
    groupId, ca = caList[0]
    coreInfo = coreList[0]

    # Save the group CA to a local file.
    groupCA = GROUP_CA_PATH + groupId + "_CA_" + str(uuid.uuid4()) + ".crt"
    if not os.path.exists(GROUP_CA_PATH):
        os.makedirs(GROUP_CA_PATH)
    groupCAFile = open(groupCA, "w")
    groupCAFile.write(ca)
    groupCAFile.close()
    discovered = True
except DiscoveryInvalidRequestException as e:
    print("Invalid discovery request detected!")
    print("Type: %s" % str(type(e)))
    print("Error message: %s" % str(e))
    print("Stopping...")
except BaseException as e:
    print("Error in discovery!")
    print("Type: %s" % str(type(e)))
    print("Error message: %s" % str(e))
    print("Stopping...")

myAWSIoTMQTTClient = AWSIoTMQTTClient(clientId)
myAWSIoTMQTTClient.configureCredentials(groupCA, privateKeyPath, certificatePath)

# Try to connect to the Greengrass core.
connected = False
for connectivityInfo in coreInfo.connectivityInfoList:
```

```
currentHost = connectivityInfo.host
currentPort = connectivityInfo.port
myAWSIoTMQTTClient.configureEndpoint(currentHost, currentPort)
try:
    myAWSIoTMQTTClient.connect()
    connected = True
except BaseException as e:
    print("Error in connect!")
    print("Type: %s" % str(type(e)))
    print("Error message: %s" % str(e))
if connected:
    break

if not connected:
    print("Cannot connect to core %s. Exiting..." % coreInfo.coreThingArn)
    sys.exit(-2)

# Handle the MQTT message received from GGShadowService.
def customCallback(client, userdata, message):
    print("Received an MQTT message")
    print(message)

# Subscribe to the MQTT topic.
myAWSIoTMQTTClient.subscribe(topic, MQTT_QOS, customCallback)

# Keep the process alive to listen for messages.
while True:
    time.sleep(1)
```

Security notes

When you use the Greengrass Docker application deployment connector, be aware of the following security considerations.

Local storage of the Docker Compose file

The connector stores a copy of your Compose file in the directory specified for the `DockerComposeFileDestinationPath` parameter.

It's your responsibility to secure this directory. You should use file system permissions to restrict access to the directory.

Local storage of the Docker credentials

If your Docker images are stored in private repositories, the connector stores your Docker credentials in the directory specified for the `DockerComposeFileDestinationPath` parameter.

It's your responsibility to secure these credentials. For example, you should use [credential-helper](#) on the core device when you install Docker Engine.

Install Docker Engine from a trusted source

It's your responsibility to install Docker Engine from a trusted source. This connector uses the Docker daemon on the core device to access your Docker assets and manage Docker containers.

Scope of Greengrass group role permissions

Permissions that you add in the Greengrass group role can be assumed by all Lambda functions and connectors in the Greengrass group. This connector requires access to your Docker Compose file stored in an S3 bucket. It also requires access to your Amazon ECR authorization token if your Docker images are stored in a private repository in Amazon ECR.

Licenses

The Greengrass Docker application deployment connector includes the following third-party software/licensing:

- [AWS SDK for Python \(Boto3\)](#)/Apache License 2.0
- [botocore](#)/Apache License 2.0
- [dateutil](#)/PSF License
- [docutils](#)/BSD License, GNU General Public License (GPL), Python Software Foundation License, Public Domain
- [jmespath](#)/MIT License
- [s3transfer](#)/Apache License 2.0
- [urllib3](#)/MIT License

This connector is released under the [Greengrass Core Software License Agreement](#).

Changelog

The following table describes the changes in each version of the connector.

Version	Changes
7	Added <code>DockerOfflineMode</code> to use an existing Docker Compose file when AWS IoT Greengrass starts offline. Implemented retries for the <code>docker login</code> command. Support for 32-bit UIDs.
6	Added <code>StopContainersOnNewDeployment</code> to override container clean up when a new deployment is made or GGC stops. Safer shutdown and start up mechanisms. YAML validation bug fix.
5	Images are pulled before running <code>docker-compose down</code> .
4	Added pull-before-up behavior to update Docker images.
3	Fixed an issue with finding environment variables.
2	Added the <code>ForceDeploy</code> parameter.
1	Initial release.

A Greengrass group can contain only one version of the connector at a time. For information about upgrading a connector version, see [the section called “Upgrading connector versions”](#).

See also

- [Integrate with services and protocols using connectors](#)
- [the section called “Get started with connectors \(console\)”](#)
- [the section called “Get started with connectors \(CLI\)”](#)

IoT Analytics connector

Warning

This connector has moved into the *extended life phase*, and AWS IoT Greengrass won't release updates that provide features, enhancements to existing features, security patches, or bug fixes. For more information, see [AWS IoT Greengrass Version 1 maintenance policy](#).

The IoT Analytics connector sends local device data to AWS IoT Analytics. You can use this connector as a central hub to collect data from sensors on the Greengrass core device and from [connected client devices](#). The connector sends the data to AWS IoT Analytics channels in the current AWS account and Region. It can send data to a default destination channel and to dynamically specified channels.

Note

AWS IoT Analytics is a fully managed service that allows you to collect, store, process, and query IoT data. In AWS IoT Analytics, the data can be further analyzed and processed. For example, it can be used to train ML models for monitoring machine health or to test new modeling strategies. For more information, see [What is AWS IoT Analytics?](#) in the *AWS IoT Analytics User Guide*.

The connector accepts formatted and unformatted data on [input MQTT topics](#). It supports two predefined topics where the destination channel is specified inline. It can also receive messages on customer-defined topics that are [configured in subscriptions](#). This can be used to route messages from client devices that publish to fixed topics or handle unstructured or stack-dependent data from resource-constrained devices.

This connector uses the [BatchPutMessage](#) API to send data (as a JSON or base64-encoded string) to the destination channel. The connector can process raw data into a format that conforms to API requirements. The connector buffers input messages in per-channel queues and asynchronously processes the batches. It provides parameters that allow you to control queueing and batching behavior and to restrict memory consumption. For example, you can configure the maximum queue size, batch interval, memory size, and number of active channels.

This connector has the following versions.

Version	ARN
4	arn:aws:greengrass: <i>region</i> ::/connectors/IoTAnalytics/versions/4
3	arn:aws:greengrass: <i>region</i> ::/connectors/IoTAnalytics/versions/3
2	arn:aws:greengrass: <i>region</i> ::/connectors/IoTAnalytics/versions/2
1	arn:aws:greengrass: <i>region</i> ::/connectors/IoTAnalytics/versions/1

For information about version changes, see the [Changelog](#).

Requirements

This connector has the following requirements:

Version 3 - 4

- AWS IoT Greengrass Core software v1.9.3 or later.
- [Python](#) version 3.7 or 3.8 installed on the core device and added to the PATH environment variable.

Note

To use Python 3.8, run the following command to create a symbolic link from the the default Python 3.7 installation folder to the installed Python 3.8 binaries.

```
sudo ln -s path-to-python-3.8/python3.8 /usr/bin/python3.7
```

This configures your device to meet the Python requirement for AWS IoT Greengrass.

- This connector can be used only in Amazon Web Services Regions where both [AWS IoT Greengrass](#) and [AWS IoT Analytics](#) are supported.
- All related AWS IoT Analytics entities and workflows are created and configured. The entities include channels, pipeline, datastores, and datasets. For more information, see the [AWS CLI](#) or [console](#) procedures in the *AWS IoT Analytics User Guide*.

Note

Destination AWS IoT Analytics channels must use the same account and be in the same AWS Region as this connector.

- The [Greengrass group role](#) configured to allow the `iotanalytics:BatchPutMessage` action on destination channels, as shown in the following example IAM policy. The channels must be in the current AWS account and Region.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "Stmt1528133056761",
      "Action": [
        "iotanalytics:BatchPutMessage"
      ],
      "Effect": "Allow",
      "Resource": [
        "arn:aws:iotanalytics:region:account-id:channel/channel_1_name",
        "arn:aws:iotanalytics:region:account-id:channel/channel_2_name"
      ]
    }
  ]
}
```

For the group role requirement, you must configure the role to grant the required permissions and make sure the role has been added to the group. For more information, see [the section called “Manage the group role \(console\)”](#) or [the section called “Manage the group role \(CLI\)”](#).

Versions 1 - 2

- AWS IoT Greengrass Core software v1.7 or later.
- [Python](#) version 2.7 installed on the core device and added to the PATH environment variable.
- This connector can be used only in Amazon Web Services Regions where both [AWS IoT Greengrass](#) and [AWS IoT Analytics](#) are supported.
- All related AWS IoT Analytics entities and workflows are created and configured. The entities include channels, pipeline, datastores, and datasets. For more information, see the [AWS CLI](#) or [console](#) procedures in the *AWS IoT Analytics User Guide*.

Note

Destination AWS IoT Analytics channels must use the same account and be in the same AWS Region as this connector.

- The [Greengrass group role](#) configured to allow the `iotanalytics:BatchPutMessage` action on destination channels, as shown in the following example IAM policy. The channels must be in the current AWS account and Region.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "Stmt1528133056761",
      "Action": [
        "iotanalytics:BatchPutMessage"
      ],
      "Effect": "Allow",
      "Resource": [
        "arn:aws:iotanalytics:region:account-id:channel/channel_1_name",
        "arn:aws:iotanalytics:region:account-id:channel/channel_2_name"
      ]
    }
  ]
}
```

For the group role requirement, you must configure the role to grant the required permissions and make sure the role has been added to the group. For more information, see [the section called “Manage the group role \(console\)”](#) or [the section called “Manage the group role \(CLI\)”](#).

Parameters

MemorySize

The amount of memory (in KB) to allocate to this connector.

Display name in the AWS IoT console: **Memory size**

Required: true

Type: string

Valid pattern: `^[0-9]+$`

PublishRegion

The AWS Region that your AWS IoT Analytics channels are created in. Use the same Region as the connector.

Note

This must also match the Region for the channels that are specified in the [group role](#).

Display name in the AWS IoT console: **Publish region**

Required: false

Type: string

Valid pattern: `^$|([a-z]{2}-[a-z]+\-\d{1})`

PublishInterval

The interval (in seconds) for publishing a batch of received data to AWS IoT Analytics.

Display name in the AWS IoT console: **Publish interval**

Required: false

Type: string

Default value: 1

Valid pattern: `$|^[0-9]+$`

IotAnalyticsMaxActiveChannels

The maximum number of AWS IoT Analytics channels that the connector actively watches for. This must be greater than 0, and at least equal to the number of channels that you expect the connector to publish to at a given time.

You can use this parameter to restrict memory consumption by limiting the total number of queues that the connector can manage at a given time. A queue is deleted when all queued messages are sent.

Display name in the AWS IoT console: **Maximum number of active channels**

Required: false

Type: string

Default value: 50

Valid pattern: `^$|^[1-9][0-9]*$`

IotAnalyticsQueueDropBehavior

The behavior for dropping messages from a channel queue when the queue is full.

Display name in the AWS IoT console: **Queue drop behavior**

Required: false

Type: string

Valid values: DROP_NEWEST or DROP_OLDEST

Default value: DROP_NEWEST

Valid pattern: `^DROP_NEWEST$|^DROP_OLDEST$`

IotAnalyticsQueueSizePerChannel

The maximum number of messages to retain in memory (per channel) before the messages are submitted or dropped. This must be greater than 0.

Display name in the AWS IoT console: **Maximum queue size per channel**

Required: false

Type: string

Default value: 2048

Valid pattern: `^$|^[1-9][0-9]*$`

`IotAnalyticsBatchSizePerChannel`

The maximum number of messages to send to an AWS IoT Analytics channel in one batch request. This must be greater than 0.

Display name in the AWS IoT console: **Maximum number of messages to batch per channel**

Required: false

Type: string

Default value: 5

Valid pattern: `^$|^[1-9][0-9]*$`

`IotAnalyticsDefaultChannelName`

The name of the AWS IoT Analytics channel that this connector uses for messages that are sent to a customer-defined input topic.

Display name in the AWS IoT console: **Default channel name**

Required: false

Type: string

Valid pattern: `^[a-zA-Z0-9_]$`

`IsolationMode`

The [containerization](#) mode for this connector. The default is `GreengrassContainer`, which means that the connector runs in an isolated runtime environment inside the AWS IoT Greengrass container.

Note

The default containerization setting for the group does not apply to connectors.

Display name in the AWS IoT console: **Container isolation mode**

Required: false

Type: string

Valid values: GreengrassContainer or NoContainer

Valid pattern: ^NoContainer\$|^GreengrassContainer\$

Create Connector Example (AWS CLI)

The following CLI command creates a ConnectorDefinition with an initial version that contains the IoT Analytics connector.

```
aws greengrass create-connector-definition --name MyGreengrassConnectors --initial-  
version '{  
  "Connectors": [  
    {  
      "Id": "MyIoTAnalyticsApplication",  
      "ConnectorArn": "arn:aws:greengrass:region::/connectors/IoTAnalytics/  
versions/3",  
      "Parameters": {  
        "MemorySize": "65535",  
        "PublishRegion": "us-west-1",  
        "PublishInterval": "2",  
        "IotAnalyticsMaxActiveChannels": "25",  
        "IotAnalyticsQueueDropBehavior": "DROP_OLDEST",  
        "IotAnalyticsQueueSizePerChannel": "1028",  
        "IotAnalyticsBatchSizePerChannel": "5",  
        "IotAnalyticsDefaultChannelName": "my_channel"  
      }  
    }  
  ]  
}'
```

Note

The Lambda function in this connector has a [long-lived](#) lifecycle.

In the AWS IoT Greengrass console, you can add a connector from the group's **Connectors** page. For more information, see [the section called "Get started with connectors \(console\)"](#).

Input data

This connector accepts data on predefined and customer-defined MQTT topics. Publishers can be client devices, Lambda functions, or other connectors.

Predefined topics

The connector supports the following two structured MQTT topics that allow publishers to specify the channel name inline.

- A [formatted message](#) on the `iotanalytics/channels/+/messages/put` topic. The IoT data in these input messages must be formatted as a JSON or base64-encoded string.
- An unformatted message on the `iotanalytics/channels/+/messages/binary/put` topic. Input messages received on this topic are treated as binary data and can contain any data type.

To publish to predefined topics, replace the `+` wildcard with the channel name. For example:

```
iotanalytics/channels/my_channel/messages/put
```

Customer-defined topics

The connector supports the `#` topic syntax, which allows it to accept input messages on any MQTT topic that you configure in a subscription. We recommend that you specify a topic path instead of using only the `#` wildcard in your subscriptions. These messages are sent to the default channel that you specify for the connector.

Input messages on customer-defined topics are treated as binary data. They can use any message format and can contain any data type. You can use customer-defined topics to route messages from devices that publish to fixed topics. You can also use them to accept input data from client devices that can't process the data into a formatted message to send to the connector.

For more information about subscriptions and MQTT topics, see [the section called "Inputs and outputs"](#).

The group role must allow the `iotanalytics:BatchPutMessage` action on all destination channels. For more information, see [the section called "Requirements"](#).

Topic filter: `iotanalytics/channels/+/messages/put`

Use this topic to send formatted messages to the connector and dynamically specify a destination channel. This topic also allows you to specify an ID that's returned in the response output. The connector verifies that IDs are unique for each message in the outbound `BatchPutMessage` request that it sends to AWS IoT Analytics. A message that has a duplicate ID is dropped.

Input data sent to this topic must use the following message format.

Message properties

`request`

The data to send to the specified channel.

Required: `true`

Type: object that includes the following properties:

`message`

The device or sensor data as a JSON or base64-encoded string.

Required: `true`

Type: `string`

`id`

An arbitrary ID for the request. This property is used to map an input request to an output response. When specified, the `id` property in the response object is set to this value. If you omit this property, the connector generates an ID.

Required: `false`

Type: `string`

Valid pattern: `.*`

Example input

```
{
```

```
"request": {
  "message" : "{\"temp\":23.33}"
},
"id" : "req123"
}
```

Topic filter: `iotanalytics/channels/+/messages/binary/put`

Use this topic to send unformatted messages to the connector and dynamically specify a destination channel.

The connector data doesn't parse the input messages received on this topic. It treats them as binary data. Before sending the messages to AWS IoT Analytics, the connector encodes and formats them to conform with BatchPutMessage API requirements:

- The connector base64-encodes the raw data and includes the encoded payload in an outbound BatchPutMessage request.
- The connector generates and assigns an ID to each input message.

 Note

The connector's response output doesn't include an ID correlation for these input messages.

Message properties

None.

Topic filter: `#`

Use this topic to send any message format to the default channel. This is especially useful when your client devices publish to fixed topics or when you want to send data to the default channel from client devices that can't process the data into the connector's [supported message format](#).

You define the topic syntax in the subscription that you create to connect this connector to the data source. We recommend that you specify a topic path instead of using only the `#` wildcard in your subscriptions.

The connector data doesn't parse the messages that are published to this input topic. All input messages are treated as binary data. Before sending the messages to AWS IoT Analytics, the connector encodes and formats them to conform with BatchPutMessage API requirements:

- The connector base64-encodes the raw data and includes the encoded payload in an outbound BatchPutMessage request.
- The connector generates and assigns an ID to each input message.

Note

The connector's response output doesn't include an ID correlation for these input messages.

Message properties

None.

Output data

This connector publishes status information as output data on an MQTT topic. This information contains the response returned by AWS IoT Analytics for each input message that it receives and sends to AWS IoT Analytics.

Topic filter in subscription

iotanalytics/messages/put/status

Example output: Success

```
{
  "response" : {
    "status" : "success"
  },
  "id" : "req123"
}
```

Example output: Failure

```
{
  "response" : {
    "status" : "fail",
    "error" : "ResourceNotFoundException",
    "error_message" : "A resource with the specified name could not be found."
  },
}
```

```
"id" : "req123"
}
```

Note

If the connector detects a retryable error (for example, connection errors), it retries the publish in the next batch. Exponential backoff is handled by the AWS SDK. Requests with retryable errors are added back to the channel queue for further publishing according to the `IotAnalyticsQueueDropBehavior` parameter.

Usage Example

Use the following high-level steps to set up an example Python 3.7 Lambda function that you can use to try out the connector.

Note

- If you use other Python runtimes, you can create a symlink from Python3.x to Python 3.7.
- The [Get started with connectors \(console\)](#) and [Get started with connectors \(CLI\)](#) topics contain detailed steps that show you how to configure and deploy an example Twilio Notifications connector.

1. Make sure you meet the [requirements](#) for the connector.

For the group role requirement, you must configure the role to grant the required permissions and make sure the role has been added to the group. For more information, see [the section called “Manage the group role \(console\)”](#) or [the section called “Manage the group role \(CLI\)”](#).

2. Create and publish a Lambda function that sends input data to the connector.

Save the [example code](#) as a PY file. Download and unzip the [AWS IoT Greengrass Core SDK for Python](#). Then, create a zip package that contains the PY file and the `greengrasssdk` folder at the root level. This zip package is the deployment package that you upload to AWS Lambda.

After you create the Python 3.7 Lambda function, publish a function version and create an alias.

3. Configure your Greengrass group.

- a. Add the Lambda function by its alias (recommended). Configure the Lambda lifecycle as long-lived (or "Pinned": `true` in the CLI).
- b. Add the connector and configure its [parameters](#).
- c. Add subscriptions that allow the connector to receive [input data](#) and send [output data](#) on supported topic filters.
 - Set the Lambda function as the source, the connector as the target, and use a supported input topic filter.
 - Set the connector as the source, AWS IoT Core as the target, and use a supported output topic filter. You use this subscription to view status messages in the AWS IoT console.

4. Deploy the group.

5. In the AWS IoT console, on the **Test** page, subscribe to the output data topic to view status messages from the connector. The example Lambda function is long-lived and starts sending messages immediately after the group is deployed.

When you're finished testing, you can set the Lambda lifecycle to on-demand (or "Pinned": `false` in the CLI) and deploy the group. This stops the function from sending messages.

Example

The following example Lambda function sends an input message to the connector.

```
import greengrasssdk
import time
import json

iot_client = greengrasssdk.client('iot-data')
send_topic = 'iotanalytics/channels/my_channel/messages/put'

def create_request_with_all_fields():
    return {
        "request": {
            "message" : "{\"temp\":23.33}"
        },
        "id" : "req_123"
    }
```

```
def publish_basic_message():
    messageToPublish = create_request_with_all_fields()
    print("Message To Publish: ", messageToPublish)
    iot_client.publish(topic=send_topic,
                       payload=json.dumps(messageToPublish))

publish_basic_message()

def lambda_handler(event, context):
    return
```

Limits

This connector is subject to the following limits.

- All limits imposed by the AWS SDK for Python (Boto3) for the AWS IoT Analytics [batch_put_message](#) action.
- All quotas imposed by the AWS IoT Analytics [BatchPutMessage](#) API. For more information, see [Service Quotas](#) for AWS IoT Analytics in the *AWS General Reference*.
 - 100,000 messages per second per channel.
 - 100 messages per batch.
 - 128 KB per message.

This API uses channel names (not channel ARNs), so sending data to cross-region or cross-account channels is not supported.

- All quotas imposed by the AWS IoT Greengrass Core. For more information, see [Service Quotas](#) for the AWS IoT Greengrass core in the *AWS General Reference*.

The following quotas might be especially applicable:

- Maximum size of messages sent by a device is 128 KB.
- Maximum message queue size in the Greengrass core router is 2.5 MB.
- Maximum length of a topic string is 256 bytes of UTF-8 encoded characters.

Licenses

The IoT Analytics connector includes the following third-party software/licensing:

- [AWS SDK for Python \(Boto3\)](#)/Apache License 2.0

- [botocore](#)/Apache License 2.0
- [dateutil](#)/PSF License
- [docutils](#)/BSD License, GNU General Public License (GPL), Python Software Foundation License, Public Domain
- [jmespath](#)/MIT License
- [s3transfer](#)/Apache License 2.0
- [urllib3](#)/MIT License

This connector is released under the [Greengrass Core Software License Agreement](#).

Changelog

The following table describes the changes in each version of the connector.

Version	Changes
4	Adds the <code>IsolationMode</code> parameter to configure the containerization mode for the connector.
3	Upgraded the Lambda runtime to Python 3.7, which changes the runtime requirement.
2	Fix to reduce excessive logging.
1	Initial release.

A Greengrass group can contain only one version of the connector at a time. For information about upgrading a connector version, see [the section called “Upgrading connector versions”](#).

See also

- [Integrate with services and protocols using connectors](#)
- [the section called “Get started with connectors \(console\)”](#)
- [the section called “Get started with connectors \(CLI\)”](#)
- [What is AWS IoT Analytics?](#) in the *AWS IoT Analytics User Guide*

IoT Ethernet IP Protocol Adapter connector

The IoT Ethernet IP Protocol Adapter [connector](#) collects data from local devices using the Ethernet/IP protocol. You can use this connector to collect data from multiple devices and publish it to a StreamManager message stream.

You can also use this connector with the IoT SiteWise connector and your IoT SiteWise gateway. Your gateway must supply the configuration for the connector. For more information, see [Configure an Ethernet/IP \(EIP\) source](#) in the IoT SiteWise user guide.

Note

This connector runs in [No container](#) isolation mode, so you can deploy it to a AWS IoT Greengrass group running in a Docker container.

This connector has the following versions.

Version	ARN
2 (recommended)	arn:aws:greengrass: <i>region</i> ::/connectors/IoTEIPProtocolAdapter/versions/2
1	arn:aws:greengrass: <i>region</i> ::/connectors/IoTEIPProtocolAdapter/versions/1

For information about version changes, see the [Changelog](#).

Requirements

This connector has the following requirements:

Version 1 and 2

- AWS IoT Greengrass Core software v1.10.2 or later.
- Stream manager enabled on the AWS IoT Greengrass group.

- Java 8 installed on the core device and added to the PATH environment variable.
- A minimum of 256 MB additional RAM. This requirement is in addition to AWS IoT Greengrass Core memory requirements.

 Note

This connector is available only in the following Regions:

- cn-north-1
- ap-southeast-1
- ap-southeast-2
- eu-central-1
- eu-west-1
- us-east-1
- us-west-2

Connector Parameters

This connector supports the following parameters:

LocalStoragePath

The directory on the AWS IoT Greengrass host that the IoT SiteWise connector can write persistent data to. The default directory is `/var/sitewise`.

Display name in the AWS IoT console: **Local storage path**

Required: false

Type: string

Valid pattern: `^\s*$|\/`.

ProtocolAdapterConfiguration

The set of Ethernet/IP collector configurations that the connector collect data from or connect to. This can be an empty list.

Display name in the AWS IoT console: **Protocol Adapter Configuration**

Required: true

Type: A well-formed JSON string that defines the set of supported feedback configurations.

The following is an example of a ProtocolAdapterConfiguration:

```
{
  "sources": [
    {
      "type": "EIPSource",
      "name": "TestSource",
      "endpoint": {
        "ipAddress": "52.89.2.42",
        "port": 44818
      },
      "destination": {
        "type": "StreamManager",
        "streamName": "MyOutput_Stream",
        "streamBufferSize": 10
      },
      "destinationPathPrefix": "EIPSource_Prefix",
      "propertyGroups": [
        {
          "name": "DriveTemperatures",
          "scanMode": {
            "type": "POLL",
            "rate": 10000
          },
          "tagPathDefinitions": [
            {
              "type": "EIPTagPath",
              "path": "arrayREAL[0]",
              "dstDataType": "double"
            }
          ]
        }
      ]
    }
  ]
}
```

Create Connector Example (AWS CLI)

The following CLI command creates a `ConnectorDefinition` with an initial version that contains the IoT Ethernet IP Protocol Adapter connector.

```
aws greengrass create-connector-definition --name MyGreengrassConnectors --initial-
version
'{
  "Connectors": [
    {
      "Id": "MyIoTEIPProtocolConnector",
      "ConnectorArn": "arn:aws:greengrass:region::/connectors/
IoTEIPProtocolAdaptor/versions/2",
      "Parameters": {
        "ProtocolAdaptorConfiguration": "{ \"sources\": [{ \"type
\": \"EIPSource\", \"name\": \"Source1\", \"endpoint\": { \"ipAddress\":
\"54.245.77.218\", \"port\": 44818 }, \"destinationPathPrefix\": \"EIPConnector_Prefix
\", \"propertyGroups\": [{ \"name\": \"Values\", \"scanMode\": { \"type\": \"POLL\",
\"rate\": 2000 }, \"tagPathDefinitions\": [{ \"type\": \"EIPTagPath\", \"path\":
\"arrayREAL[0]\", \"dstDataType\": \"double\" }]]]]}",
        "LocalStoragePath": "/var/MyIoTEIPProtocolConnectorState"
      }
    }
  ]
}'
```

Note

The Lambda function in this connector has a [long-lived](#) lifecycle.

Input data

This connector doesn't accept MQTT messages as input data.

Output data

This connector publishes data to `StreamManager`. You must configure the destination message stream. The output messages are of the following structure:

```
{
```

```
"alias": "string",
"messages": [
  {
    "name": "string",
    "value": boolean|double|integer|string,
    "timestamp": number,
    "quality": "string"
  }
]
```

Licenses

The IoT Ethernet IP Protocol Adapter connector includes the following third-party software/licensing:

- [Ethernet/IP client](#)
- [MapDB](#)
- [Elsa](#)

This connector is released under the [Greengrass Core Software License Agreement](#).

Changelog

The following table describes the changes in each version of the connector.

Version	Changes	Date
2	This version contains bug fixes.	December 23, 2021
1	Initial release.	December 15, 2020

A Greengrass group can contain only one version of the connector at a time. For information about upgrading a connector version, see [the section called “Upgrading connector versions”](#).

See also

- [Integrate with services and protocols using connectors](#)

- [the section called “Get started with connectors \(console\)”](#)
- [the section called “Get started with connectors \(CLI\)”](#)

IoT SiteWise connector

The IoT SiteWise connector sends local device and equipment data to asset properties in AWS IoT SiteWise. You can use this connector to collect data from multiple OPC-UA servers and publish it to IoT SiteWise. The connector sends the data to asset properties in the current AWS account and Region.

Note

IoT SiteWise is a fully managed service that collects, processes, and visualizes data from industrial devices and equipment. You can configure asset properties that process raw data sent from this connector to your assets' measurement properties. For example, you can define a transform property that converts a device's Celsius temperature data points to Fahrenheit, or you can define a metric property that calculates the average hourly temperature. For more information, see [What is AWS IoT SiteWise?](#) in the *AWS IoT SiteWise User Guide*.

The connector sends data to IoT SiteWise with the OPC-UA data stream paths sent from the OPC-UA servers. For example, the data stream path `/company/windfarm/3/turbine/7/temperature` might represent the temperature sensor of turbine #7 at wind farm #3. If the AWS IoT Greengrass core loses connection to the internet, the connector caches data until it can successfully connect to the AWS Cloud. You can configure the maximum disk buffer size used for caching data. If the cache size exceeds the maximum disk buffer size, the connector discards the oldest data from the queue.

After you configure and deploy the IoT SiteWise connector, you can add a gateway and OPC-UA sources in the [IoT SiteWise console](#). When you configure a source in the console, you can filter or prefix the OPC-UA data stream paths sent by the IoT SiteWise connector. For instructions to finish setting up your gateway and sources, see [Adding the gateway](#) in the *AWS IoT SiteWise User Guide*.

IoT SiteWise receives data only from data streams that you have mapped to the measurement properties of IoT SiteWise assets. To map data streams to asset properties, you can set a property's alias to be equivalent to an OPC-UA data stream path. To learn about defining asset models and creating assets, see [Modeling industrial assets](#) in the *AWS IoT SiteWise User Guide*.

Notes

You can use stream manager to upload data to IoT SiteWise from sources other than OPC-UA servers. Stream manager also provides customizable support for persistence and bandwidth management. For more information, see [Manage data streams](#).

This connector runs in [No container](#) isolation mode, so you can deploy it to a Greengrass group running in a Docker container.

This connector has the following versions.

Version	ARN
12 (recommended)	<code>arn:aws:greengrass: <i>region</i>::/connectors/IoTSiteWise/versions/ 12</code>
11	<code>arn:aws:greengrass: <i>region</i>::/connectors/IoTSiteWise/versions/ 11</code>
10	<code>arn:aws:greengrass: <i>region</i>::/connectors/IoTSiteWise/versions/ 10</code>
9	<code>arn:aws:greengrass: <i>region</i>::/connectors/IoTSiteWise/versions/ 9</code>
8	<code>arn:aws:greengrass: <i>region</i>::/connectors/IoTSiteWise/versions/ 8</code>
7	<code>arn:aws:greengrass: <i>region</i>::/connectors/IoTSiteWise/versions/ 7</code>

Version	ARN
6	arn:aws:greengrass: <i>region</i> ::/connectors/IoTSiteWise/versions/ 6
5	arn:aws:greengrass: <i>region</i> ::/connectors/IoTSiteWise/versions/ 5
4	arn:aws:greengrass: <i>region</i> ::/connectors/IoTSiteWise/versions/ 4
3	arn:aws:greengrass: <i>region</i> ::/connectors/IoTSiteWise/versions/ 3
2	arn:aws:greengrass: <i>region</i> ::/connectors/IoTSiteWise/versions/ 2
1	arn:aws:greengrass: <i>region</i> ::/connectors/IoTSiteWise/versions/ 1

For information about version changes, see the [Changelog](#).

Requirements

This connector has the following requirements:

Version 9, 10, 11, and 12

Important

This version introduces new requirements: AWS IoT Greengrass Core software v1.10.2 and [stream manager](#).

- AWS IoT Greengrass Core software v1.10.2.
- [Stream manager](#) enabled on the Greengrass group.
- Java 8 installed on the core device and added to the PATH environment variable.
- This connector can be used only in Amazon Web Services Regions where both [AWS IoT Greengrass](#) and [IoT SiteWise](#) are supported.
- An IAM policy added to the Greengrass group role. This role allows the AWS IoT Greengrass group access to the `iotsitewise:BatchPutAssetPropertyValue` action on the target root asset and its children, as shown in the following example. You can remove the Condition from the policy to allow the connector to access all of your IoT SiteWise assets.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": "iotsitewise:BatchPutAssetPropertyValue",
      "Resource": "*",
      "Condition": {
        "StringLike": {
          "iotsitewise:assetHierarchyPath": [
            "/root node asset ID",
            "/root node asset ID/*"
          ]
        }
      }
    }
  ]
}
```

For more information, see [Adding and removing IAM policies](#) in the *IAM User Guide*.

Versions 6, 7, and 8

Important

This version introduces new requirements: AWS IoT Greengrass Core software v1.10.0 and [stream manager](#).

- AWS IoT Greengrass Core software v1.10.0.
- [Stream manager](#) enabled on the Greengrass group.
- Java 8 installed on the core device and added to the PATH environment variable.
- This connector can be used only in Amazon Web Services Regions where both [AWS IoT Greengrass](#) and [IoT SiteWise](#) are supported.
- An IAM policy added to the Greengrass group role. This role allows the AWS IoT Greengrass group access to the `iotsitewise:BatchPutAssetPropertyValue` action on the target root asset and its children, as shown in the following example. You can remove the Condition from the policy to allow the connector to access all of your IoT SiteWise assets.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": "iotsitewise:BatchPutAssetPropertyValue",
      "Resource": "*",
      "Condition": {
        "StringLike": {
          "iotsitewise:assetHierarchyPath": [
            "/root node asset ID",
            "/root node asset ID/*"
          ]
        }
      }
    }
  ]
}
```

For more information, see [Adding and removing IAM policies](#) in the *IAM User Guide*.

Version 5

- AWS IoT Greengrass Core software v1.9.4.
- Java 8 installed on the core device and added to the PATH environment variable.
- This connector can be used only in Amazon Web Services Regions where both [AWS IoT Greengrass](#) and [IoT SiteWise](#) are supported.

- An IAM policy added to the Greengrass group role. This role allows the AWS IoT Greengrass group access to the `iotsitewise:BatchPutAssetPropertyValue` action on the target root asset and its children, as shown in the following example. You can remove the `Condition` from the policy to allow the connector to access all of your IoT SiteWise assets.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": "iotsitewise:BatchPutAssetPropertyValue",
      "Resource": "*",
      "Condition": {
        "StringLike": {
          "iotsitewise:assetHierarchyPath": [
            "/root node asset ID",
            "/root node asset ID/*"
          ]
        }
      }
    }
  ]
}
```

For more information, see [Adding and removing IAM policies](#) in the *IAM User Guide*.

Version 4

- AWS IoT Greengrass Core software v1.10.0.
- Java 8 installed on the core device and added to the `PATH` environment variable.
- This connector can be used only in Amazon Web Services Regions where both [AWS IoT Greengrass](#) and [IoT SiteWise](#) are supported.
- An IAM policy added to the Greengrass group role. This role allows the AWS IoT Greengrass group access to the `iotsitewise:BatchPutAssetPropertyValue` action on the target root asset and its children, as shown in the following example. You can remove the `Condition` from the policy to allow the connector to access all of your IoT SiteWise assets.

```
{
  "Version": "2012-10-17",
```

```

    "Statement": [
      {
        "Effect": "Allow",
        "Action": "iotsitewise:BatchPutAssetPropertyValue",
        "Resource": "*",
        "Condition": {
          "StringLike": {
            "iotsitewise:assetHierarchyPath": [
              "/root node asset ID",
              "/root node asset ID/*"
            ]
          }
        }
      }
    ]
  }
}

```

For more information, see [Adding and removing IAM policies](#) in the *IAM User Guide*.

Version 3

- AWS IoT Greengrass Core software v1.9.4.
- Java 8 installed on the core device and added to the PATH environment variable.
- This connector can be used only in Amazon Web Services Regions where both [AWS IoT Greengrass](#) and [IoT SiteWise](#) are supported.
- An IAM policy added to the Greengrass group role. This role allows the AWS IoT Greengrass group access to the `iotsitewise:BatchPutAssetPropertyValue` action on the target root asset and its children, as shown in the following example. You can remove the Condition from the policy to allow the connector to access all of your IoT SiteWise assets.

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": "iotsitewise:BatchPutAssetPropertyValue",
      "Resource": "*",
      "Condition": {
        "StringLike": {
          "iotsitewise:assetHierarchyPath": [

```

```
}  
[  
    }  
    }  
}  
]  
["/root node asset ID",  
 "/root node asset ID/*"]  
]
```

For more information, see [Adding and removing IAM policies](#) in the *IAM User Guide*.

Versions 1 and 2

- AWS IoT Greengrass Core software v1.9.4.
- Java 8 installed on the core device and added to the PATH environment variable.
- This connector can be used only in Amazon Web Services Regions where both [AWS IoT Greengrass](#) and [IoT SiteWise](#) are supported.
- An IAM policy added to the Greengrass group role that allows access to AWS IoT Core and the `iotsitewise:BatchPutAssetPropertyValue` action on the target root asset and its children, as shown in the following example. You can remove the Condition from the policy to allow the connector to access all of your IoT SiteWise assets.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": "iotsitewise:BatchPutAssetPropertyValue",
      "Resource": "*",
      "Condition": {
        "StringLike": {
          "iotsitewise:assetHierarchyPath": [
            "/root node asset ID",
            "/root node asset ID/*"
          ]
        }
      }
    },
    {
      "Effect": "Allow",
```

```
        "Action": [
            "iot:Connect",
            "iot:DescribeEndpoint",
            "iot:Publish",
            "iot:Receive",
            "iot:Subscribe"
        ],
        "Resource": "*"
    }
]
```

For more information, see [Adding and removing IAM identity permissions](#) in the *IAM User Guide*.

Parameters

Versions 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, and 12

SiteWiseLocalStoragePath

The directory on the AWS IoT Greengrass host that the IoT SiteWise connector can write persistent data to. Defaults to `/var/sitewise`.

Display name in the AWS IoT console: **Local storage path**

Required: false

Type: string

Valid pattern: `^\s*$|\/`.

AWSecretsArnList

A list of secrets in AWS Secrets Manager that each contain a OPC-UA user name and password key-value pair. Each secret must be a key-value pair type secret.

Display name in the AWS IoT console: **List of ARNs for OPC-UA username/password secrets**

Required: false

Type: `JSONArrayOfStrings`

Valid pattern: `\[( ? , ? ?\"(arn:(aws(-[a-z]+)*):secretsmanager:[a-z0-9\\-]+:[0-9]{12}:secret:([a-zA-Z0-9\\\\\\\\]+\\\/)*[a-zA-Z0-9\\\/_+=, .@\\-]+-[a-zA-Z0-9]+)*\\")*\]`

MaximumBufferSize

The maximum size in GB for IoT SiteWise disk usage. Defaults to 10GB.

Display name in the AWS IoT console: **Maximum disk buffer size**

Required: false

Type: string

Valid pattern: `^\\s*$|[0-9]+`

Version 1

SiteWiseLocalStoragePath

The directory on the AWS IoT Greengrass host that the IoT SiteWise connector can write persistent data to. Defaults to `/var/sitewise`.

Display name in the AWS IoT console: **Local storage path**

Required: false

Type: string

Valid pattern: `^\\s*$|\\/.`

SiteWiseOpcuaUserIdentityTokenSecretArn

The secret in AWS Secrets Manager that contains the OPC-UA user name and password key-value pair. This secret must be a key-value pair type secret.

Display name in the AWS IoT console: **ARN of OPC-UA username/password secret**

Required: false

Type: string

Valid pattern: `^$|arn:(aws(-[a-z]+)*):secretsmanager:[a-z0-9\\-]+:[0-9]{12}:secret:([a-zA-Z0-9\\+\\/*]*[a-zA-Z0-9/_+=, .@\\-]+-[a-zA-Z0-9]+`

`SiteWiseOpcuaUserIdentityTokenSecretArn-ResourceId`

The secret resource in the AWS IoT Greengrass group that references an OPC-UA user name and password secret.

Display name in the AWS IoT console: **OPC-UA username/password secret resource**

Required: false

Type: string

Valid pattern: `^$|.+`

`MaximumBufferSize`

The maximum size in GB for IoT SiteWise disk usage. Defaults to 10GB.

Display name in the AWS IoT console: **Maximum disk buffer size**

Required: false

Type: string

Valid pattern: `^\\s*$|[0-9]+`

Create Connector Example (AWS CLI)

The following AWS CLI command creates a `ConnectorDefinition` with an initial version that contains the IoT SiteWise connector.

```
aws greengrass create-connector-definition --name MyGreengrassConnectors --initial-
version '{
  "Connectors": [
    {
      "Id": "MyIoTSiteWiseConnector",
      "ConnectorArn": "arn:aws:greengrass:region::/connectors/IoTSiteWise/
versions/11"
    }
  ]
}'
```


Note

The Lambda functions in this connector have a [long-lived](#) lifecycle.

In the AWS IoT Greengrass console, you can add a connector from the group's **Connectors** page. For more information, see [the section called "Get started with connectors \(console\)"](#).

Input data

This connector doesn't accept MQTT messages as input data.

Output data

This connector doesn't publish MQTT messages as output data.

Limits

This connector is subject to the following all limits imposed by IoT SiteWise, including the following. For more information, see [AWS IoT SiteWise endpoints and quotas](#) in the *AWS General Reference*.

- Maximum number of gateways per AWS account.
- Maximum number of OPC-UA sources per gateway.
- Maximum rate of timestamp-quality-value (TQV) data points stored per AWS account.
- Maximum rate of TQV data points stored per asset property.

Licenses

Version 9, 10, 11, and 12

The IoT SiteWise connector includes the following third-party software/licensing:

- [MapDB](#)
- [Elsa](#)
- [Eclipse Milo](#)

This connector is released under the [Greengrass Core Software License Agreement](#).

Versions 6, 7, and 8

The IoT SiteWise connector includes the following third-party software/licensing:

- [Milo](#) / EDL 1.0

This connector is released under the [Greengrass Core Software License Agreement](#).

Versions 1, 2, 3, 4, and 5

The IoT SiteWise connector includes the following third-party software/licensing:

- [Milo](#) / EDL 1.0
- [Chronicle-Queue](#) / Apache License 2.0

This connector is released under the [Greengrass Core Software License Agreement](#).

Changelog

The following table describes the changes in each version of the connector.

Version	Changes	Date
12	• This version contains bug fixes.	December 22, 2021
11	• Support for strings that contain hidden or unprintable characters. Hidden and unprintable characters are automatically removed before the strings are sent to the AWS Cloud. • Fixed an issue that caused the IoT SiteWise gateway to infinitely retry invalid requests.	March 24, 2021

Version	Changes	Date
	- Fixed an issue that caused a corrupted checkpoint when the IoT SiteWise gateway was connected to a high-frequency data source. - Improved error messages to help troubleshoot the gateway configuration.	
10	Configured StreamManager to improve handling when the source connection is lost and re-established. This version also accepts OPC-UA values with a ServerTimestamp when no SourceTimestamp is available.	January 22, 2021
9	Support launched for custom Greengrass StreamManager stream destinations, OPC-UA deadbanding, custom scan mode and custom scan rate. Also includes improved performance during configuration updates made from the IoT SiteWise gateway.	December 15, 2020
8	Improved stability when the connector experiences intermittent network connectivity.	November 19, 2020

Version	Changes	Date
7	Fixed an issue with gateway metrics.	August 14, 2020
6	Added support for CloudWatch metrics and automatic discovery of new OPC-UA tags. This version requires stream manager and AWS IoT Greengrass Core software v1.10.0 or higher.	April 29, 2020
5	Fixed a compatibility issue with AWS IoT Greengrass Core software v1.9.4.	February 12, 2020
4	Fixed an issue with OPC-UA server reconnection.	February 7, 2020
3	Removed <code>iot:*</code> permissions requirement.	December 17, 2019
2	Added support for multiple OPC-UA secret resources.	December 10, 2019
1	Initial release.	December 2, 2019

A Greengrass group can contain only one version of the connector at a time. For information about upgrading a connector version, see [the section called “Upgrading connector versions”](#).

See also

- [Integrate with services and protocols using connectors](#)
- [the section called “Get started with connectors \(console\)”](#)
- [the section called “Get started with connectors \(CLI\)”](#)
- See the following topics in the *AWS IoT SiteWise User Guide*:

- [What is AWS IoT SiteWise?](#)
- [Using a gateway](#)
- [Gateway CloudWatch metrics](#)
- [Troubleshooting an IoT SiteWise gateway](#)

Kinesis Firehose

The Kinesis Firehose [connector](#) publishes data through an Amazon Data Firehose delivery stream to destinations such as Amazon S3, Amazon Redshift, or Amazon OpenSearch Service.

This connector is a data producer for a Kinesis delivery stream. It receives input data on an MQTT topic, and sends the data to a specified delivery stream. The delivery stream then sends the data record to the configured destination (for example, an S3 bucket).

This connector has the following versions.

Version	ARN
5	arn:aws:greengrass: <i>region</i> ::/connectors/KinesisFirehose/versions/5
4	arn:aws:greengrass: <i>region</i> ::/connectors/KinesisFirehose/versions/4
3	arn:aws:greengrass: <i>region</i> ::/connectors/KinesisFirehose/versions/3
2	arn:aws:greengrass: <i>region</i> ::/connectors/KinesisFirehose/versions/2
1	arn:aws:greengrass: <i>region</i> ::/connectors/KinesisFirehose/versions/1

For information about version changes, see the [Changelog](#).

Requirements

This connector has the following requirements:

Version 4 - 5

- AWS IoT Greengrass Core software v1.9.3 or later.
- [Python](#) version 3.7 or 3.8 installed on the core device and added to the PATH environment variable.

Note

To use Python 3.8, run the following command to create a symbolic link from the the default Python 3.7 installation folder to the installed Python 3.8 binaries.

```
sudo ln -s path-to-python-3.8/python3.8 /usr/bin/python3.7
```

This configures your device to meet the Python requirement for AWS IoT Greengrass.

- A configured Kinesis delivery stream. For more information, see [Creating an Amazon Data Firehose delivery stream](#) in the *Amazon Kinesis Firehose Developer Guide*.
- The [Greengrass group role](#) configured to allow the `firehose:PutRecord` and `firehose:PutRecordBatch` actions on the target delivery stream, as shown in the following example IAM policy.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "Stmt1528133056761",
      "Action": [
        "firehose:PutRecord",
        "firehose:PutRecordBatch"
      ],
      "Effect": "Allow",
      "Resource": [
        "arn:aws:firehose:region:account-id:deliverystream/stream-name"
      ]
    }
  ]
}
```

```
    }
  ]
}
```

This connector allows you to dynamically override the default delivery stream in the input message payload. If your implementation uses this feature, the IAM policy should include all target streams as resources. You can grant granular or conditional access to resources (for example, by using a wildcard * naming scheme).

For the group role requirement, you must configure the role to grant the required permissions and make sure the role has been added to the group. For more information, see [the section called “Manage the group role \(console\)”](#) or [the section called “Manage the group role \(CLI\)”](#).

Versions 2 - 3

- AWS IoT Greengrass Core software v1.7 or later.
- [Python](#) version 2.7 installed on the core device and added to the PATH environment variable.
- A configured Kinesis delivery stream. For more information, see [Creating an Amazon Data Firehose delivery stream](#) in the *Amazon Kinesis Firehose Developer Guide*.
- The [Greengrass group role](#) configured to allow the `firehose:PutRecord` and `firehose:PutRecordBatch` actions on the target delivery stream, as shown in the following example IAM policy.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "Stmt1528133056761",
      "Action": [
        "firehose:PutRecord",
        "firehose:PutRecordBatch"
      ],
      "Effect": "Allow",
      "Resource": [
        "arn:aws:firehose:region:account-id:deliverystream/stream-name"
      ]
    }
  ]
}
```

This connector allows you to dynamically override the default delivery stream in the input message payload. If your implementation uses this feature, the IAM policy should include all target streams as resources. You can grant granular or conditional access to resources (for example, by using a wildcard * naming scheme).

For the group role requirement, you must configure the role to grant the required permissions and make sure the role has been added to the group. For more information, see [the section called “Manage the group role \(console\)”](#) or [the section called “Manage the group role \(CLI\)”](#).

Version 1

- AWS IoT Greengrass Core software v1.7 or later.
- [Python](#) version 2.7 installed on the core device and added to the PATH environment variable.
- A configured Kinesis delivery stream. For more information, see [Creating an Amazon Data Firehose delivery stream](#) in the *Amazon Kinesis Firehose Developer Guide*.
- The [Greengrass group role](#) configured to allow the `firehose:PutRecord` action on the target delivery stream, as shown in the following example IAM policy.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "Stmt1528133056761",
      "Action": [
        "firehose:PutRecord"
      ],
      "Effect": "Allow",
      "Resource": [
        "arn:aws:firehose:region:account-id:deliverystream/stream-name"
      ]
    }
  ]
}
```

This connector allows you to dynamically override the default delivery stream in the input message payload. If your implementation uses this feature, the IAM policy should include all target streams as resources. You can grant granular or conditional access to resources (for example, by using a wildcard * naming scheme).

For the group role requirement, you must configure the role to grant the required permissions and make sure the role has been added to the group. For more information, see [the section called “Manage the group role \(console\)”](#) or [the section called “Manage the group role \(CLI\)”](#).

Connector Parameters

This connector provides the following parameters:

Versions 5

DefaultDeliveryStreamArn

The ARN of the default Firehose delivery stream to send data to. The destination stream can be overridden by the `delivery_stream_arn` property in the input message payload.

Note

The group role must allow the appropriate actions on all target delivery streams. For more information, see [the section called “Requirements”](#).

Display name in the AWS IoT console: **Default delivery stream ARN**

Required: true

Type: string

Valid pattern: `arn:aws:firehose:([a-z]{2}-[a-z]+-\d{1}):(\d{12}):deliverystream/([a-zA-Z0-9_\-\.]+)$`

DeliveryStreamQueueSize

The maximum number of records to retain in memory before new records for the same delivery stream are rejected. The minimum value is 2000.

Display name in the AWS IoT console: **Maximum number of records to buffer (per stream)**

Required: true

Type: string

Valid pattern: `^([2-9]\\d{3}|[1-9]\\d{4,})$`

MemorySize

The amount of memory (in KB) to allocate to this connector.

Display name in the AWS IoT console: **Memory size**

Required: true

Type: string

Valid pattern: `^[0-9]+$`

PublishInterval

The interval (in seconds) for publishing records to Firehose. To disable batching, set this value to 0.

Display name in the AWS IoT console: **Publish interval**

Required: true

Type: string

Valid values: 0 - 900

Valid pattern: `[0-9]|[1-9]\\d|[1-9]\\d\\d|900`

IsolationMode

The [containerization](#) mode for this connector. The default is `GreengrassContainer`, which means that the connector runs in an isolated runtime environment inside the AWS IoT Greengrass container.

Note

The default containerization setting for the group does not apply to connectors.

Display name in the AWS IoT console: **Container isolation mode**

Required: false

Type: string

Valid values: GreengrassContainer or NoContainer

Valid pattern: ^NoContainer\$|^GreengrassContainer\$

Versions 2 - 4

DefaultDeliveryStreamArn

The ARN of the default Firehose delivery stream to send data to. The destination stream can be overridden by the `delivery_stream_arn` property in the input message payload.

Note

The group role must allow the appropriate actions on all target delivery streams. For more information, see [the section called “Requirements”](#).

Display name in the AWS IoT console: **Default delivery stream ARN**

Required: true

Type: string

Valid pattern: `arn:aws:firehose:([a-z]{2}-[a-z]+-\d{1}):(\d{12}):deliverystream/([a-zA-Z0-9_\-.]+)$`

DeliveryStreamQueueSize

The maximum number of records to retain in memory before new records for the same delivery stream are rejected. The minimum value is 2000.

Display name in the AWS IoT console: **Maximum number of records to buffer (per stream)**

Required: true

Type: string

Valid pattern: `^([2-9]\\d{3}|[1-9]\\d{4,})$`

MemorySize

The amount of memory (in KB) to allocate to this connector.

Display name in the AWS IoT console: **Memory size**

Required: true

Type: string

Valid pattern: `^[0-9]+$`

PublishInterval

The interval (in seconds) for publishing records to Firehose. To disable batching, set this value to 0.

Display name in the AWS IoT console: **Publish interval**

Required: true

Type: string

Valid values: 0 - 900

Valid pattern: `[0-9] | [1-9]\\d | [1-9]\\d\\d | 900`

Version 1

DefaultDeliveryStreamArn

The ARN of the default Firehose delivery stream to send data to. The destination stream can be overridden by the `delivery_stream_arn` property in the input message payload.

Note

The group role must allow the appropriate actions on all target delivery streams. For more information, see [the section called "Requirements"](#).

Display name in the AWS IoT console: **Default delivery stream ARN**

Required: true

Type: string

Valid pattern: `arn:aws:firehose:([a-z]{2}-[a-z]+\d{1}):(\d{12}):deliverystream/([a-zA-Z0-9_\-.]+)$`

Example

Create Connector Example (AWS CLI)

The following CLI command creates a `ConnectorDefinition` with an initial version that contains the connector.

```
aws greengrass create-connector-definition --name MyGreengrassConnectors --initial-  
version '{  
  "Connectors": [  
    {  
      "Id": "MyKinesisFirehoseConnector",  
      "ConnectorArn": "arn:aws:greengrass:region::/connectors/KinesisFirehose/  
versions/5",  
      "Parameters": {  
        "DefaultDeliveryStreamArn": "arn:aws:firehose:region:account-  
id:deliverystream/stream-name",  
        "DeliveryStreamQueueSize": "5000",  
        "MemorySize": "65535",  
        "PublishInterval": "10",  
        "IsolationMode" : "GreengrassContainer"  
      }  
    }  
  ]  
}'
```

In the AWS IoT Greengrass console, you can add a connector from the group's **Connectors** page. For more information, see [the section called “Get started with connectors \(console\)”](#).

Input data

This connector accepts stream content on MQTT topics, and then sends the content to the target delivery stream. It accepts two types of input data:

- JSON data on the `kinesisfirehose/message` topic.
- Binary data on the `kinesisfirehose/message/binary/#` topic.

Versions 2 - 5

Topic filter: `kinesisfirehose/message`

Use this topic to send a message that contains JSON data.

Message properties

`request`

The data to send to the delivery stream and the target delivery stream, if different from the default stream.

Required: `true`

Type: object that includes the following properties:

`data`

The data to send to the delivery stream.

Required: `true`

Type: `string`

`delivery_stream_arn`

The ARN of the target Kinesis delivery stream. Include this property to override the default delivery stream.

Required: `false`

Type: `string`

Valid pattern: `arn:aws:firehose:([a-z]{2}-[a-z]+-\d{1}):(\d{12}):deliverystream/([a-zA-Z0-9_\-.]+)$`

`id`

An arbitrary ID for the request. This property is used to map an input request to an output response. When specified, the `id` property in the response object is set to this value. If you don't use this feature, you can omit this property or specify an empty string.

Required: `false`

Type: string

Valid pattern: .*

Example input

```
{
  "request": {
    "delivery_stream_arn": "arn:aws:firehose:region:account-
id:deliverystream/stream2-name",
    "data": "Data to send to the delivery stream."
  },
  "id": "request123"
}
```

Topic filter: kinesisfirehose/message/binary/#

Use this topic to send a message that contains binary data. The connector doesn't parse binary data. The data is streamed as is.

To map the input request to an output response, replace the # wildcard in the message topic with an arbitrary request ID. For example, if you publish a message to kinesisfirehose/message/binary/request123, the id property in the response object is set to request123.

If you don't want to map a request to a response, you can publish your messages to kinesisfirehose/message/binary/. Be sure to include the trailing slash.

Version 1

Topic filter: kinesisfirehose/message

Use this topic to send a message that contains JSON data.

Message properties

request

The data to send to the delivery stream and the target delivery stream, if different from the default stream.

Required: true

Type: object that includes the following properties:

data

The data to send to the delivery stream.

Required: true

Type: string

delivery_stream_arn

The ARN of the target Kinesis delivery stream. Include this property to override the default delivery stream.

Required: false

Type: string

Valid pattern: `arn:aws:firehose:([a-z]{2}-[a-z]+-\d{1}):(\d{12}):deliverystream/([a-zA-Z0-9_\-\.]+)$`

id

An arbitrary ID for the request. This property is used to map an input request to an output response. When specified, the `id` property in the response object is set to this value. If you don't use this feature, you can omit this property or specify an empty string.

Required: false

Type: string

Valid pattern: `.*`

Example input

```
{
  "request": {
    "delivery_stream_arn": "arn:aws:firehose:region:account-
id:deliverystream/stream2-name",
    "data": "Data to send to the delivery stream."
  },
  "id": "request123"
```



```
}
```

Topic filter: `kinesisfirehose/message/binary/#`

Use this topic to send a message that contains binary data. The connector doesn't parse binary data. The data is streamed as is.

To map the input request to an output response, replace the # wildcard in the message topic with an arbitrary request ID. For example, if you publish a message to `kinesisfirehose/message/binary/request123`, the `id` property in the response object is set to `request123`.

If you don't want to map a request to a response, you can publish your messages to `kinesisfirehose/message/binary/`. Be sure to include the trailing slash.

Output data

This connector publishes status information as output data on an MQTT topic.

Versions 2 - 5

Topic filter in subscription

`kinesisfirehose/message/status`

Example output

The response contains the status of each data record sent in the batch.

```
{
  "response": [
    {
      "ErrorCode": "error",
      "ErrorMessage": "test error",
      "id": "request123",
      "status": "fail"
    },
    {
      "firehose_record_id": "xyz2",
      "id": "request456",
```

```
        "status": "success"
      },
      {
        "firehose_record_id": "xyz3",
        "id": "request890",
        "status": "success"
      }
    ]
  }
}
```

Note

If the connector detects a retryable error (for example, connection errors), it retries the publish in the next batch. Exponential backoff is handled by the AWS SDK. Requests that fail with retryable errors are added back to the end of the queue for further publishing.

Version 1

Topic filter in subscription

kinesisfirehose/message/status

Example output: Success

```
{
  "response": {
    "firehose_record_id": "1lxfuuuFomkpJYzt/34ZU/r8JYPf8Wyf7AXqlXm",
    "status": "success"
  },
  "id": "request123"
}
```

Example output: Failure

```
{
  "response" : {
    "error": "ResourceNotFoundException",
    "error_message": "An error occurred (ResourceNotFoundException) when calling the PutRecord operation: Firehose test1 not found under account 123456789012.",
  }
}
```

```
    "status": "fail"
  },
  "id": "request123"
}
```

Usage Example

Use the following high-level steps to set up an example Python 3.7 Lambda function that you can use to try out the connector.

Note

- If you use other Python runtimes, you can create a symlink from Python3.x to Python 3.7.
- The [Get started with connectors \(console\)](#) and [Get started with connectors \(CLI\)](#) topics contain detailed steps that show you how to configure and deploy an example Twilio Notifications connector.

1. Make sure you meet the [requirements](#) for the connector.

For the group role requirement, you must configure the role to grant the required permissions and make sure the role has been added to the group. For more information, see [the section called “Manage the group role \(console\)”](#) or [the section called “Manage the group role \(CLI\)”](#).

2. Create and publish a Lambda function that sends input data to the connector.

Save the [example code](#) as a PY file. Download and unzip the [AWS IoT Greengrass Core SDK for Python](#). Then, create a zip package that contains the PY file and the greengrasssdk folder at the root level. This zip package is the deployment package that you upload to AWS Lambda.

After you create the Python 3.7 Lambda function, publish a function version and create an alias.

3. Configure your Greengrass group.
 - a. Add the Lambda function by its alias (recommended). Configure the Lambda lifecycle as long-lived (or "Pinned": true in the CLI).
 - b. Add the connector and configure its [parameters](#).

- c. Add subscriptions that allow the connector to receive [JSON input data](#) and send [output data](#) on supported topic filters.
 - Set the Lambda function as the source, the connector as the target, and use a supported input topic filter.
 - Set the connector as the source, AWS IoT Core as the target, and use a supported output topic filter. You use this subscription to view status messages in the AWS IoT console.
4. Deploy the group.
5. In the AWS IoT console, on the **Test** page, subscribe to the output data topic to view status messages from the connector. The example Lambda function is long-lived and starts sending messages immediately after the group is deployed.

When you're finished testing, you can set the Lambda lifecycle to on-demand (or "Pinned": `false` in the CLI) and deploy the group. This stops the function from sending messages.

Example

The following example Lambda function sends an input message to the connector. This message contains JSON data.

```
import greengrasssdk
import time
import json

iot_client = greengrasssdk.client('iot-data')
send_topic = 'kinesisfirehose/message'

def create_request_with_all_fields():
    return {
        "request": {
            "data": "Message from Firehose Connector Test"
        },
        "id" : "req_123"
    }

def publish_basic_message():
    messageToPublish = create_request_with_all_fields()
    print("Message To Publish: ", messageToPublish)
    iot_client.publish(topic=send_topic,
        payload=json.dumps(messageToPublish))
```

```
publish_basic_message()  
  
def lambda_handler(event, context):  
    return
```

Licenses

The Kinesis Firehose connector includes the following third-party software/licensing:

- [AWS SDK for Python \(Boto3\)](#)/Apache License 2.0
- [botocore](#)/Apache License 2.0
- [dateutil](#)/PSF License
- [docutils](#)/BSD License, GNU General Public License (GPL), Python Software Foundation License, Public Domain
- [jmespath](#)/MIT License
- [s3transfer](#)/Apache License 2.0
- [urllib3](#)/MIT License

This connector is released under the [Greengrass Core Software License Agreement](#).

Changelog

The following table describes the changes in each version of the connector.

Version	Changes
5	Added the <code>IsolationMode</code> parameter to configure the containerization mode for the connector.
4	Upgraded the Lambda runtime to Python 3.7, which changes the runtime requirement.
3	Fix to reduce excessive logging and other minor bug fixes.

Version	Changes
2	Added support for sending batched data records to Firehose at a specified interval. • Also requires the <code>firehose:PutRecordBatch</code> action in the group role. • New <code>MemorySize</code> , <code>DeliveryStreamQueueSize</code> , and <code>PublishInterval</code> parameters. • Output message contains an array of status responses for the published data records.
1	Initial release.

A Greengrass group can contain only one version of the connector at a time. For information about upgrading a connector version, see [the section called “Upgrading connector versions”](#).

See also

- [Integrate with services and protocols using connectors](#)
- [the section called “Get started with connectors \(console\)”](#)
- [the section called “Get started with connectors \(CLI\)”](#)
- [What is Amazon Kinesis Data Firehose?](#) in the *Amazon Kinesis Developer Guide*

ML Feedback connector

Warning

This connector has moved into the *extended life phase*, and AWS IoT Greengrass won't release updates that provide features, enhancements to existing features, security patches, or bug fixes. For more information, see [AWS IoT Greengrass Version 1 maintenance policy](#).

The ML Feedback connector makes it easier to access your machine learning (ML) model data for model retraining and analysis. The connector:

- Uploads input data (samples) used by your ML model to Amazon S3. Model input can be in any format, such as images, JSON, or audio. After samples are uploaded to the cloud, you can use them to retrain the model to improve the accuracy and precision of its predictions. For example, you can use [SageMaker AI Ground Truth](#) to label your samples and [SageMaker AI](#) to retrain the model.
- Publishes the prediction results from the model as MQTT messages. This lets you monitor and analyze the inference quality of your model in real time. You can also store prediction results and use them to analyze trends over time.
- Publishes metrics about sample uploads and sample data to Amazon CloudWatch.

To configure this connector, you describe your supported *feedback configurations* in JSON format. A feedback configuration defines properties such as the destination Amazon S3 bucket, content type, and [sampling strategy](#). (A sampling strategy is used to determine which samples to upload.)

You can use the ML Feedback connector in the following scenarios:

- With user-defined Lambda functions. Your local inference Lambda functions use the AWS IoT Greengrass Machine Learning SDK to invoke this connector and pass in the target feedback configuration, model input, and model output (prediction results). For an example, see [the section called “Usage Example”](#).
- With the [ML Image Classification connector](#) (v2). To use this connector with the ML Image Classification connector, configure the `MLFeedbackConnectorConfigId` parameter for the ML Image Classification connector.
- With the [ML Object Detection connector](#). To use this connector with the ML Object Detection connector, configure the `MLFeedbackConnectorConfigId` parameter for the ML Object Detection connector.

ARN: `arn:aws:greengrass:region::/connectors/MLFeedback/versions/1`

Requirements

This connector has the following requirements:

- AWS IoT Greengrass Core Software v1.9.3 or later.
- [Python](#) version 3.7 or 3.8 installed on the core device and added to the PATH environment variable.

Note

To use Python 3.8, run the following command to create a symbolic link from the the default Python 3.7 installation folder to the installed Python 3.8 binaries.

```
sudo ln -s path-to-python-3.8/python3.8 /usr/bin/python3.7
```

This configures your device to meet the Python requirement for AWS IoT Greengrass.

- One or more Amazon S3 buckets. The number of buckets you use depends on your sampling strategy.
- The [Greengrass group role](#) configured to allow the `s3:PutObject` action on objects in the destination Amazon S3 bucket, as shown in the following example IAM policy.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": "s3:PutObject",
      "Resource": [
        "arn:aws:s3:::bucket-name/*"
      ]
    }
  ]
}
```

The policy should include all destination buckets as resources. You can grant granular or conditional access to resources (for example, by using a wildcard `*` naming scheme).

For the group role requirement, you must configure the role to grant the required permissions and make sure the role has been added to the group. For more information, see [the section called “Manage the group role \(console\)”](#) or [the section called “Manage the group role \(CLI\)”](#).

- The [CloudWatch Metrics connector](#) added to the Greengrass group and configured. This is required only if you want to use the metrics reporting feature.
- [AWS IoT Greengrass Machine Learning SDK](#) v1.1.0 is required to interact with this connector.

Parameters

FeedbackConfigurationMap

A set of one or more feedback configurations that the connector can use to upload samples to Amazon S3. A feedback configuration defines parameters such as the destination bucket, content type, and [sampling strategy](#). When this connector is invoked, the calling Lambda function or connector specifies a target feedback configuration.

Display name in the AWS IoT console: **Feedback configuration map**

Required: true

Type: A well-formed JSON string that defines the set of supported feedback configurations. For an example, see [the section called “FeedbackConfigurationMap example”](#).

The ID of a feedback configuration object has the following requirements.

The ID:

- Must be unique across configuration objects.
- Must begin with a letter or number. Can contain lowercase and uppercase letters, numbers, and hyphens.
- Must be 2 - 63 characters in length.

Required: true

Type: string

Valid pattern: `^[a-zA-Z0-9][a-zA-Z0-9-]{1,62}$`

Examples: MyConfig0, config-a, 12id

The body of a feedback configuration object contains the following properties.

s3-bucket-name

The name of the destination Amazon S3 bucket.

Note

The group role must allow the `s3:PutObject` action on all destination buckets. For more information, see [the section called “Requirements”](#).

Required: true

Type: string

Valid pattern: `^[a-z0-9\.\-]{3,63}$`

content-type

The content type of the samples to upload. All content for an individual feedback configuration must be of the same type.

Required: true

Type: string

Examples: image/jpeg, application/json, audio/ogg

s3-prefix

The key prefix to use for uploaded samples. A prefix is similar to a directory name. It allows you to store similar data under the same directory in a bucket. For more information, see [Object key and metadata](#) in the *Amazon Simple Storage Service User Guide*.

Required: false

Type: string

file-ext

The file extension to use for uploaded samples. Must be a valid file extension for the content type.

Required: false

Type: string

Examples: jpg, json, ogg

sampling-strategy

The [sampling strategy](#) to use to filter which samples to upload. If omitted, the connector tries to upload all the samples that it receives.

Required: false

Type: A well-formed JSON string that contains the following properties.

strategy-name

The name of the sampling strategy.

Required: true

Type: string

Valid values: RANDOM_SAMPLING, LEAST_CONFIDENCE, MARGIN, or ENTROPY

rate

The rate for the [Random](#) sampling strategy.

Required: true if strategy-name is RANDOM_SAMPLING.

Type: number

Valid values: 0.0 - 1.0

threshold

The threshold for the [Least Confidence](#), [Margin](#), or [Entropy](#) sampling strategy.

Required: true if strategy-name is LEAST_CONFIDENCE, MARGIN, or ENTROPY.

Type: number

Valid values:

- 0.0 - 1.0 for the LEAST_CONFIDENCE or MARGIN strategy.
- 0.0 - no limit for the ENTROPY strategy.

RequestLimit

The maximum number of requests that the connector can process at a time.

You can use this parameter to restrict memory consumption by limiting the number of requests that the connector processes at the same time. Requests that exceed this limit are ignored.

Display name in the AWS IoT console: **Request limit**

Required: false

Type: string

Valid values: 0 - 999

Valid pattern: `^$|^([0-9]){1,3}$`

Create Connector Example (AWS CLI)

The following CLI command creates a `ConnectorDefinition` with an initial version that contains the ML Feedback connector.

```
aws greengrass create-connector-definition --name MyGreengrassConnectors --initial-
version '{
  "Connectors": [
    {
      "Id": "MyMLFeedbackConnector",
      "ConnectorArn": "arn:aws:greengrass:region::/connectors/MLFeedback/
versions/1",
      "Parameters": {
        "FeedbackConfigurationMap": "{ \"RandomSamplingConfiguration\":
{ \"s3-bucket-name\": \"my-aws-bucket-random-sampling\", \"content-type\":
\"image/png\", \"file-ext\": \"png\", \"sampling-strategy\": { \"strategy-name
\": \"RANDOM_SAMPLING\", \"rate\": 0.5 } }, \"LeastConfidenceConfiguration\": {
  \"s3-bucket-name\": \"my-aws-bucket-least-confidence-sampling\", \"content-type\":
  \"image/png\", \"file-ext\": \"png\", \"sampling-strategy\": { \"strategy-name\":
  \"LEAST_CONFIDENCE\", \"threshold\": 0.4 } } }",
        "RequestLimit": "10"
      }
    }
  ]
}'
```

FeedbackConfigurationMap example

The following is an expanded example value for the `FeedbackConfigurationMap` parameter. This example includes several feedback configurations that use different sampling strategies.

```
{
  "ConfigID1": {
    "s3-bucket-name": "my-aws-bucket-random-sampling",
    "content-type": "image/png",
    "file-ext": "png",
    "sampling-strategy": {
      "strategy-name": "RANDOM_SAMPLING",
```

```
        "rate": 0.5
      }
    },
    "ConfigID2": {
      "s3-bucket-name": "my-aws-bucket-margin-sampling",
      "content-type": "image/png",
      "file-ext": "png",
      "sampling-strategy": {
        "strategy-name": "MARGIN",
        "threshold": 0.4
      }
    },
    "ConfigID3": {
      "s3-bucket-name": "my-aws-bucket-least-confidence-sampling",
      "content-type": "image/png",
      "file-ext": "png",
      "sampling-strategy": {
        "strategy-name": "LEAST_CONFIDENCE",
        "threshold": 0.4
      }
    },
    "ConfigID4": {
      "s3-bucket-name": "my-aws-bucket-entropy-sampling",
      "content-type": "image/png",
      "file-ext": "png",
      "sampling-strategy": {
        "strategy-name": "ENTROPY",
        "threshold": 2
      }
    },
    "ConfigID5": {
      "s3-bucket-name": "my-aws-bucket-no-sampling",
      "s3-prefix": "DeviceA",
      "content-type": "application/json"
    }
  }
}
```

Sampling strategies

The connector supports four sampling strategies that determine whether to upload samples that are passed to the connector. Samples are discrete instances of data that a model uses for a prediction. You can use sampling strategies to filter for the samples that are most likely to improve model accuracy.

RANDOM_SAMPLING

Randomly uploads samples based on the supplied rate. It uploads a sample if a randomly generated value is less than the rate. The higher the rate, the more samples are uploaded.

Note

This strategy disregards any model prediction that is supplied.

LEAST_CONFIDENCE

Uploads samples whose maximum confidence probability falls below the supplied threshold.

Example scenario:

Threshold: .6

Model prediction: [.2, .2, .4, .2]

Maximum confidence probability: .4

Result:

Use the sample because maximum confidence probability (.4) $\leq$ threshold (.6).

MARGIN

Uploads samples if the margin between the top two confidence probabilities falls within the supplied threshold. The margin is the difference between the top two probabilities.

Example scenario:

Threshold: .02

Model prediction: [.3, .35, .34, .01]

Top two confidence probabilities: [.35, .34]

Margin: .01 (.35 - .34)

Result:

Use the sample because margin (.01) $\leq$ threshold (.02).

ENTROPY

Uploads samples whose entropy is greater than the supplied threshold. Uses the model prediction's normalized entropy.

Example scenario:

Threshold: 0.75

Model prediction: [.5, .25, .25]

Entropy for prediction: 1.03972

Result:

Use sample because entropy (1.03972) > threshold (0.75).

Input data

User-defined Lambda functions use the `publish` function of the feedback client in the AWS IoT Greengrass Machine Learning SDK to invoke the connector. For an example, see [the section called "Usage Example"](#).

Note

This connector doesn't accept MQTT messages as input data.

The `publish` function takes the following arguments:

ConfigId

The ID of the target feedback configuration. This must match the ID of a feedback configuration defined in the [FeedbackConfigurationMap](#) parameter for the ML Feedback connector.

Required: true

Type: string

ModelInput

The input data that was passed to a model for inference. This input data is uploaded using the target configuration unless it is filtered out based on the sampling strategy.

Required: true

Type: bytes

ModelPrediction

The prediction results from the model. The result type can be a dictionary or a list. For example, the prediction results from the ML Image Classification connector is a list of probabilities (such as [0.25, 0.60, 0.15]). This data is published to the `/feedback/message/prediction` topic.

Required: true

Type: dictionary or list of float values

Metadata

Customer-defined, application-specific metadata that is attached to the uploaded sample and published to the `/feedback/message/prediction` topic. The connector also inserts a `publish-ts` key with a timestamp value into the metadata.

Required: false

Type: dictionary

Example: `{"some-key": "some value"}`

Output data

This connector publishes data to three MQTT topics:

- Status information from the connector on the `feedback/message/status` topic.
- Prediction results on the `feedback/message/prediction` topic.
- Metrics destined for CloudWatch on the `cloudwatch/metric/put` topic.

You must configure subscriptions to allow the connector to communicate on MQTT topics. For more information, see [the section called "Inputs and outputs"](#).

Topic filter: `feedback/message/status`

Use this topic to monitor the status of sample uploads and dropped samples. The connector publishes to this topic every time that it receives a request.

Example output: Sample upload succeeded

```
{
  "response": {
    "status": "success",
    "s3_response": {
      "ResponseMetadata": {
        "HostId": "IOWQ4fDEXAMPLEQM+ey7N9WgVhSnQ6JEXAMPLEZb7hSQDASK
+Jd1vEXAMPLEEa3Km",
        "RetryAttempts": 1,
        "HTTPStatusCode": 200,
        "RequestId": "79104EXAMPLEB723",
        "HTTPHeaders": {
          "content-length": "0",
          "x-amz-id-2":
"lbbqaDVF0hMlyU3gRvAX1ZIdg8P0WkGkCSSFsYFvSwLZk3j7QZhG5EXAMPLEdd4/pEXAMPLEUqU=",
          "server": "AmazonS3",
          "x-amz-expiration": "expiry-date=\\"Wed, 17 Jul 2019 00:00:00 GMT\\",
rule-id=\\"OGZjYWY3OTgtYWI2Zi00ZDl1LWE4YmQtNzMyYzEXAMPLEoUw\\\"",
          "x-amz-request-id": "79104EXAMPLEB723",
          "etag": "\\\"b9c4f172e64458a5fd674EXAMPLE5628\\\"",
          "date": "Thu, 11 Jul 2019 00:12:50 GMT",
          "x-amz-server-side-encryption": "AES256"
        }
      },
      "bucket": "greengrass-feedback-connector-data-us-west-2",
      "ETag": "\\\"b9c4f172e64458a5fd674EXAMPLE5628\\\"",
      "Expiration": "expiry-date=\\"Wed, 17 Jul 2019 00:00:00 GMT\\", rule-id=
\\"OGZjYWY3OTgtYWI2Zi00ZDl1LWE4YmQtNzMyYzEXAMPLEoUw\\\"",
      "key": "s3-key-prefix/UUID.file_ext",
      "ServerSideEncryption": "AES256"
    }
  },
  "id": "5aaa913f-97a3-48ac-5907-18cd96b89eeb"
}
```

The connector adds the bucket and key fields to the response from Amazon S3. For more information about the Amazon S3 response, see [PUT object](#) in the *Amazon Simple Storage Service API Reference*.

Example output: Sample dropped because of the sampling strategy

```
{
```

```
"response": {
  "status": "sample_dropped_by_strategy"
},
"id": "4bf5aeb0-d1e4-4362-5bb4-87c05de78ba3"
}
```

Example output: Sample upload failed

A failure status includes the error message as the `error_message` value and the exception class as the `error` value.

```
{
  "response": {
    "status": "fail",
    "error_message": "[RequestId: 4bf5aeb0-d1e4-4362-5bb4-87c05de78ba3] Failed to upload model input data due to exception. Model prediction will not be published. Exception type: NoSuchBucket, error: An error occurred (NoSuchBucket) when calling the PutObject operation: The specified bucket does not exist",
    "error": "NoSuchBucket"
  },
  "id": "4bf5aeb0-d1e4-4362-5bb4-87c05de78ba3"
}
```

Example output: Request throttled because of the request limit

```
{
  "response": {
    "status": "fail",
    "error_message": "Request limit has been reached (max request: 10 ). Dropping request.",
    "error": "Queue.Full"
  },
  "id": "4bf5aeb0-d1e4-4362-5bb4-87c05de78ba3"
}
```

Topic filter: feedback/message/prediction

Use this topic to listen for predictions based on uploaded sample data. This lets you analyze your model performance in real time. Model predictions are published to this topic only if data is successfully uploaded to Amazon S3. Messages published on this topic are in JSON format. They contain the link to the uploaded data object, the model's prediction, and the metadata included in the request.

You can also store prediction results and use them to report and analyze trends over time. Trends can provide valuable insights. For example, a *decreasing accuracy over time* trend can help you to decide whether the model needs to be retrained.

Example output

```
{
  "source-ref": "s3://greengrass-feedback-connector-data-us-west-2/s3-key-prefix/
UUID.file_ext",
  "model-prediction": [
    0.5,
    0.2,
    0.2,
    0.1
  ],
  "config-id": "ConfigID2",
  "metadata": {
    "publish-ts": "2019-07-11 00:12:48.816752"
  }
}
```

Tip

You can configure the [IoT Analytics connector](#) to subscribe to this topic and send the information to AWS IoT Analytics for further or historical analysis.

Topic filter: cloudwatch/metric/put

This is the output topic used to publish metrics to CloudWatch. This feature requires that you install and configure the [CloudWatch Metrics connector](#).

Metrics include:

- The number of uploaded samples.
- The size of uploaded samples.
- The number of errors from uploads to Amazon S3.
- The number of dropped samples based on the sampling strategy.
- The number of throttled requests.

Example output: Size of the data sample (published before the actual upload)

```
{
  "request": {
    "namespace": "GreengrassFeedbackConnector",
    "metricData": {
      "value": 47592,
      "unit": "Bytes",
      "metricName": "SampleSize"
    }
  }
}
```

Example output: Sample upload succeeded

```
{
  "request": {
    "namespace": "GreengrassFeedbackConnector",
    "metricData": {
      "value": 1,
      "unit": "Count",
      "metricName": "SampleUploadSuccess"
    }
  }
}
```

Example output: Sample upload succeeded and prediction result published

```
{
  "request": {
    "namespace": "GreengrassFeedbackConnector",
    "metricData": {
      "value": 1,
      "unit": "Count",
      "metricName": "SampleAndPredictionPublished"
    }
  }
}
```

Example output: Sample upload failed

```
{
```

```
"request": {
  "namespace": "GreengrassFeedbackConnector",
  "metricData": {
    "value": 1,
    "unit": "Count",
    "metricName": "SampleUploadFailure"
  }
}
```

Example output: Sample dropped because of the sampling strategy

```
{
  "request": {
    "namespace": "GreengrassFeedbackConnector",
    "metricData": {
      "value": 1,
      "unit": "Count",
      "metricName": "SampleNotUsed"
    }
  }
}
```

Example output: Request throttled because of the request limit

```
{
  "request": {
    "namespace": "GreengrassFeedbackConnector",
    "metricData": {
      "value": 1,
      "unit": "Count",
      "metricName": "ErrorRequestThrottled"
    }
  }
}
```

Usage Example

The following example is a user-defined Lambda function that uses the [AWS IoT Greengrass Machine Learning SDK](#) to send data to the ML Feedback connector.

Note

You can download the AWS IoT Greengrass Machine Learning SDK from the AWS IoT Greengrass [downloads page](#).

```
import json
import logging
import os
import sys
import greengrass_machine_learning_sdk as ml

client = ml.client('feedback')

try:
    feedback_config_id = os.environ["FEEDBACK_CONFIG_ID"]
    model_input_data_dir = os.environ["MODEL_INPUT_DIR"]
    model_prediction_str = os.environ["MODEL_PREDICTIONS"]
    model_prediction = json.loads(model_prediction_str)
except Exception as e:
    logging.info("Failed to open environment variables. Failed with exception:
{}".format(e))
    sys.exit(1)

try:
    with open(os.path.join(model_input_data_dir, os.listdir(model_input_data_dir)[0]),
'rb') as f:
        content = f.read()
except Exception as e:
    logging.info("Failed to open model input directory. Failed with exception:
{}".format(e))
    sys.exit(1)

def invoke_feedback_connector():
    logging.info("Invoking feedback connector.")
    try:
        client.publish(
            ConfigId=feedback_config_id,
            ModelInput=content,
            ModelPrediction=model_prediction
        )
    except Exception as e:
```

```
logging.info("Exception raised when invoking feedback connector:{}".format(e))
sys.exit(1)

invoke_feedback_connector()

def function_handler(event, context):
    return
```

Licenses

The ML Feedback connector includes the following third-party software/licensing:

- [AWS SDK for Python \(Boto3\)](#)/Apache License 2.0
- [botocore](#)/Apache License 2.0
- [dateutil](#)/PSF License
- [docutils](#)/BSD License, GNU General Public License (GPL), Python Software Foundation License, Public Domain
- [jmespath](#)/MIT License
- [s3transfer](#)/Apache License 2.0
- [urllib3](#)/MIT License
- [six](#)/MIT

This connector is released under the [Greengrass Core Software License Agreement](#).

See also

- [Integrate with services and protocols using connectors](#)
- [the section called “Get started with connectors \(console\)”](#)
- [the section called “Get started with connectors \(CLI\)”](#)

ML Image Classification connector

Warning

This connector has moved into the *extended life phase*, and AWS IoT Greengrass won't release updates that provide features, enhancements to existing features, security patches, or bug fixes. For more information, see [AWS IoT Greengrass Version 1 maintenance policy](#).

The ML Image Classification [connectors](#) provide a machine learning (ML) inference service that runs on the AWS IoT Greengrass core. This local inference service performs image classification using a model trained by the SageMaker AI image classification algorithm.

User-defined Lambda functions use the AWS IoT Greengrass Machine Learning SDK to submit inference requests to the local inference service. The service runs inference locally and returns probabilities that the input image belongs to specific categories.

AWS IoT Greengrass provides the following versions of this connector, which is available for multiple platforms.

Version 2

Connector	Description and ARN
ML Image Classification Aarch64 JTX2	Image classification inference service for NVIDIA Jetson TX2. Supports GPU acceleration. ARN: <code>arn:aws:greengrass:<i>region</i>::/connectors/ImageClassificationAarch64JTX2/versions/2</code>
ML Image Classification x86_64	Image classification inference service for x86_64 platforms. ARN: <code>arn:aws:greengrass:<i>region</i>::/connectors/ImageClassificationx86_64/versions/2</code>

Connector	Description and ARN
	eClassificationx86-64/versions/2
ML Image Classification ARMv7	Image classification inference service for ARMv7 platforms. ARN: <code>arn:aws:greengrass: <i>region</i> :/connectors/ImageClassificationARMv7/versions/2</code>

Version 1

Connector	Description and ARN
ML Image Classification Aarch64 JTX2	Image classification inference service for NVIDIA Jetson TX2. Supports GPU acceleration. ARN: <code>arn:aws:greengrass: <i>region</i> :/connectors/ImageClassificationAarch64JTX2/versions/1</code>
ML Image Classification x86_64	Image classification inference service for x86_64 platforms. ARN: <code>arn:aws:greengrass: <i>region</i> :/connectors/ImageClassificationx86-64/versions/1</code>
ML Image Classification Armv7	Image classification inference service for Armv7 platforms.

Connector	Description and ARN
	ARN: arn:aws:greengrass : <i>region</i> ::/connectors/ImageClassificationARMv7/versions/1

For information about version changes, see the [Changelog](#).

Requirements

These connectors have the following requirements:

Version 2

- AWS IoT Greengrass Core Software v1.9.3 or later.
- [Python](#) version 3.7 or 3.8 installed on the core device and added to the PATH environment variable.

Note

To use Python 3.8, run the following command to create a symbolic link from the the default Python 3.7 installation folder to the installed Python 3.8 binaries.

```
sudo ln -s path-to-python-3.8/python3.8 /usr/bin/python3.7
```

This configures your device to meet the Python requirement for AWS IoT Greengrass.

- Dependencies for the Apache MXNet framework installed on the core device. For more information, see [the section called “Installing MXNet dependencies”](#).
- An [ML resource](#) in the Greengrass group that references an SageMaker AI model source. This model must be trained by the SageMaker AI image classification algorithm. For more information, see [Image classification algorithm](#) in the *Amazon SageMaker AI Developer Guide*.
- The [ML Feedback connector](#) added to the Greengrass group and configured. This is required only if you want to use the connector to upload model input data and publish predictions to an MQTT topic.

- The [Greengrass group role](#) configured to allow the `sagemaker:DescribeTrainingJob` action on the target training job, as shown in the following example IAM policy.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "sagemaker:DescribeTrainingJob"
      ],
      "Resource": "arn:aws:sagemaker:region:account-id:training-  
job:training-job-name"
    }
  ]
}
```

For the group role requirement, you must configure the role to grant the required permissions and make sure the role has been added to the group. For more information, see [the section called “Manage the group role \(console\)”](#) or [the section called “Manage the group role \(CLI\)”](#).

You can grant granular or conditional access to resources (for example, by using a wildcard * naming scheme). If you change the target training job in the future, make sure to update the group role.

- [AWS IoT Greengrass Machine Learning SDK](#) v1.1.0 is required to interact with this connector.

Version 1

- AWS IoT Greengrass Core Software v1.7 or later.
- [Python](#) version 2.7 installed on the core device and added to the PATH environment variable.
- Dependencies for the Apache MXNet framework installed on the core device. For more information, see [the section called “Installing MXNet dependencies”](#).
- An [ML resource](#) in the Greengrass group that references an SageMaker AI model source. This model must be trained by the SageMaker AI image classification algorithm. For more information, see [Image classification algorithm](#) in the *Amazon SageMaker AI Developer Guide*.
- The [Greengrass group role](#) configured to allow the `sagemaker:DescribeTrainingJob` action on the target training job, as shown in the following example IAM policy.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "sagemaker:DescribeTrainingJob"
      ],
      "Resource": "arn:aws:sagemaker:region:account-id:training-  
job:training-job-name"
    }
  ]
}
```

For the group role requirement, you must configure the role to grant the required permissions and make sure the role has been added to the group. For more information, see [the section called “Manage the group role \(console\)”](#) or [the section called “Manage the group role \(CLI\)”](#).

You can grant granular or conditional access to resources (for example, by using a wildcard * naming scheme). If you change the target training job in the future, make sure to update the group role.

- [AWS IoT Greengrass Machine Learning SDK](#) v1.0.0 or later is required to interact with this connector.

Connector Parameters

These connectors provide the following parameters.

Version 2

MLModelDestinationPath

The absolute local path of the ML resource inside the Lambda environment. This is the destination path that's specified for the ML resource.

Note

If you created the ML resource in the console, this is the local path.

Display name in the AWS IoT console: **Model destination path**

Required: true

Type: string

Valid pattern: .+

`MLModelResourceId`

The ID of the ML resource that references the source model.

Display name in the AWS IoT console: **SageMaker job ARN resource**

Required: true

Type: string

Valid pattern: [a-zA-Z0-9:_-]+

`MLModelSageMakerJobArn`

The ARN of the SageMaker AI training job that represents the SageMaker AI model source. The model must be trained by the SageMaker AI image classification algorithm.

Display name in the AWS IoT console: **SageMaker job ARN**

Required: true

Type: string

Valid pattern: ^arn:aws:sagemaker:[a-zA-Z0-9-]+:[0-9]+:training-job/[a-zA-Z0-9][a-zA-Z0-9-]+\$

`LocalInferenceServiceName`

The name for the local inference service. User-defined Lambda functions invoke the service by passing the name to the `invoke_inference_service` function of the AWS IoT Greengrass Machine Learning SDK. For an example, see [the section called "Usage Example"](#).

Display name in the AWS IoT console: **Local inference service name**

Required: true

Type: string

Valid pattern: `[a-zA-Z0-9][a-zA-Z0-9-]{1,62}`

LocalInferenceServiceTimeoutSeconds

The amount of time (in seconds) before the inference request is terminated. The minimum value is 1.

Display name in the AWS IoT console: **Timeout (second)**

Required: true

Type: string

Valid pattern: `[1-9][0-9]*`

LocalInferenceServiceMemoryLimitKB

The amount of memory (in KB) that the service has access to. The minimum value is 1.

Display name in the AWS IoT console: **Memory limit (KB)**

Required: true

Type: string

Valid pattern: `[1-9][0-9]*`

GPUAcceleration

The CPU or GPU (accelerated) computing context. This property applies to the ML Image Classification Aarch64 JTX2 connector only.

Display name in the AWS IoT console: **GPU acceleration**

Required: true

Type: string

Valid values: CPU or GPU

MLFeedbackConnectorConfigId

The ID of the feedback configuration to use to upload model input data. This must match the ID of a feedback configuration defined for the [ML Feedback connector](#).

This parameter is required only if you want to use the ML Feedback connector to upload model input data and publish predictions to an MQTT topic.

Display name in the AWS IoT console: **ML Feedback connector configuration ID**

Required: false

Type: string

Valid pattern: `^[a-zA-Z0-9]{1,62}$`

Version 1

`MLModelDestinationPath`

The absolute local path of the ML resource inside the Lambda environment. This is the destination path that's specified for the ML resource.



Note

If you created the ML resource in the console, this is the local path.

Display name in the AWS IoT console: **Model destination path**

Required: true

Type: string

Valid pattern: `.+`

`MLModelResourceId`

The ID of the ML resource that references the source model.

Display name in the AWS IoT console: **SageMaker job ARN resource**

Required: true

Type: string

Valid pattern: `[a-zA-Z0-9:_-]+`

`MLModelSageMakerJobArn`

The ARN of the SageMaker AI training job that represents the SageMaker AI model source. The model must be trained by the SageMaker AI image classification algorithm.

Display name in the AWS IoT console: **SageMaker job ARN**

Required: true

Type: string

Valid pattern: `^arn:aws:sagemaker:[a-zA-Z0-9-]+:[0-9]+:training-job/[a-zA-Z0-9][a-zA-Z0-9-]+$`

`LocalInferenceServiceName`

The name for the local inference service. User-defined Lambda functions invoke the service by passing the name to the `invoke_inference_service` function of the AWS IoT Greengrass Machine Learning SDK. For an example, see [the section called "Usage Example"](#).

Display name in the AWS IoT console: **Local inference service name**

Required: true

Type: string

Valid pattern: `[a-zA-Z0-9][a-zA-Z0-9-]{1,62}`

`LocalInferenceServiceTimeoutSeconds`

The amount of time (in seconds) before the inference request is terminated. The minimum value is 1.

Display name in the AWS IoT console: **Timeout (second)**

Required: true

Type: string

Valid pattern: `[1-9][0-9]*`

`LocalInferenceServiceMemoryLimitKB`

The amount of memory (in KB) that the service has access to. The minimum value is 1.

Display name in the AWS IoT console: **Memory limit (KB)**

Required: true

Type: string

Valid pattern: `[1-9][0-9]*`

GPUAcceleration

The CPU or GPU (accelerated) computing context. This property applies to the ML Image Classification Aarch64 JTX2 connector only.

Display name in the AWS IoT console: **GPU acceleration**

Required: true

Type: string

Valid values: CPU or GPU

Create Connector Example (AWS CLI)

The following CLI commands create a ConnectorDefinition with an initial version that contains an ML Image Classification connector.

Example: CPU Instance

This example creates an instance of the ML Image Classification Armv7l connector.

```
aws greengrass create-connector-definition --name MyGreengrassConnectors --initial-  
version '{  
  "Connectors": [  
    {  
      "Id": "MyImageClassificationConnector",  
      "ConnectorArn": "arn:aws:greengrass:region::/connectors/  
ImageClassificationARMv7/versions/2",  
      "Parameters": {  
        "MLModelDestinationPath": "/path-to-model",  
        "MLModelResourceId": "my-ml-resource",  
        "MLModelSageMakerJobArn": "arn:aws:sagemaker:us-  
west-2:123456789012:training-job:MyImageClassifier",  
        "LocalInferenceServiceName": "imageClassification",  
        "LocalInferenceServiceTimeoutSeconds": "10",  
        "LocalInferenceServiceMemoryLimitKB": "500000",  
        "MLFeedbackConnectorConfigId": "MyConfig0"  
      }  
    }  
  ]  
}'
```

Example: GPU Instance

This example creates an instance of the ML Image Classification Aarch64 JTX2 connector, which supports GPU acceleration on an NVIDIA Jetson TX2 board.

```
aws greengrass create-connector-definition --name MyGreengrassConnectors --initial-  
version '{  
  "Connectors": [  
    {  
      "Id": "MyImageClassificationConnector",  
      "ConnectorArn": "arn:aws:greengrass:region::/connectors/  
ImageClassificationAarch64JTX2/versions/2",  
      "Parameters": {  
        "MLModelDestinationPath": "/path-to-model",  
        "MLModelResourceId": "my-ml-resource",  
        "MLModelSageMakerJobArn": "arn:aws:sagemaker:us-  
west-2:123456789012:training-job:MyImageClassifier",  
        "LocalInferenceServiceName": "imageClassification",  
        "LocalInferenceServiceTimeoutSeconds": "10",  
        "LocalInferenceServiceMemoryLimitKB": "500000",  
        "GPUAcceleration": "GPU",  
        "MLFeedbackConnectorConfigId": "MyConfig0"  
      }  
    }  
  ]  
}'
```

Note

The Lambda function in these connectors have a [long-lived](#) lifecycle.

In the AWS IoT Greengrass console, you can add a connector from the group's **Connectors** page. For more information, see [the section called "Get started with connectors \(console\)"](#).

Input data

These connectors accept an image file as input. Input image files must be in jpeg or png format. For more information, see [the section called "Usage Example"](#).

These connectors don't accept MQTT messages as input data.

Output data

These connectors return a formatted prediction for the object identified in the input image:

```
[0.3,0.1,0.04,...]
```

The prediction contains a list of values that correspond with the categories used in the training dataset during model training. Each value represents the probability that the image falls under the corresponding category. The category with the highest probability is the dominant prediction.

These connectors don't publish MQTT messages as output data.

Usage Example

The following example Lambda function uses the [AWS IoT Greengrass Machine Learning SDK](#) to interact with an ML Image Classification connector.

Note

You can download the SDK from the [AWS IoT Greengrass Machine Learning SDK](#) downloads page.

The example initializes an SDK client and synchronously calls the SDK's `invoke_inference_service` function to invoke the local inference service. It passes in the algorithm type, service name, image type, and image content. Then, the example parses the service response to get the probability results (predictions).

Python 3.7

```
import logging
from threading import Timer

import numpy as np

import greengrass_machine_learning_sdk as ml

# We assume the inference input image is provided as a local file
# to this inference client Lambda function.
with open('/test_img/test.jpg', 'rb') as f:
    content = bytearray(f.read())
```

```

client = ml.client('inference')

def infer():
    logging.info('invoking Greengrass ML Inference service')

    try:
        resp = client.invoke_inference_service(
            AlgoType='image-classification',
            ServiceName='imageClassification',
            ContentType='image/jpeg',
            Body=content
        )
    except ml.GreengrassInferenceException as e:
        logging.info('inference exception {}'.format(e.__class__.__name__, e))
        return
    except ml.GreengrassDependencyException as e:
        logging.info('dependency exception {}'.format(e.__class__.__name__,
e))
        return

    logging.info('resp: {}'.format(resp))
    predictions = resp['Body'].read().decode("utf-8")
    logging.info('predictions: {}'.format(predictions))

    # The connector output is in the format: [0.3,0.1,0.04,...]
    # Remove the '[' and ']' at the beginning and end.
    predictions = predictions[1:-1]
    count = len(predictions.split(','))
    predictions_arr = np.fromstring(predictions, count=count, sep=',')

    # Perform business logic that relies on the predictions_arr, which is an array
    # of probabilities.

    # Schedule the infer() function to run again in one second.
    Timer(1, infer).start()
    return

infer()

def function_handler(event, context):
    return

```

Python 2.7

```
import logging
from threading import Timer

import numpy

import greengrass_machine_learning_sdk as gg_ml

# The inference input image.
with open("/test_img/test.jpg", "rb") as f:
    content = f.read()

client = gg_ml.client("inference")

def infer():
    logging.info("Invoking Greengrass ML Inference service")

    try:
        resp = client.invoke_inference_service(
            AlgoType="image-classification",
            ServiceName="imageClassification",
            ContentType="image/jpeg",
            Body=content,
        )
    except gg_ml.GreengrassInferenceException as e:
        logging.info('Inference exception %s("%s")', e.__class__.__name__, e)
        return
    except gg_ml.GreengrassDependencyException as e:
        logging.info('Dependency exception %s("%s")', e.__class__.__name__, e)
        return

    logging.info("Response: %s", resp)
    predictions = resp["Body"].read()
    logging.info("Predictions: %s", predictions)

    # The connector output is in the format: [0.3,0.1,0.04,...]
    # Remove the '[' and ']' at the beginning and end.
    predictions = predictions[1:-1]
    predictions_arr = numpy.fromstring(predictions, sep=",")
    logging.info("Split into %s predictions.", len(predictions_arr))

    # Perform business logic that relies on predictions_arr, which is an array
```

```
# of probabilities.

# Schedule the infer() function to run again in one second.
Timer(1, infer).start()

infer()

# In this example, the required AWS Lambda handler is never called.
def function_handler(event, context):
    return
```

The `invoke_inference_service` function in the AWS IoT Greengrass Machine Learning SDK accepts the following arguments.

Argument	Description
AlgoType	The name of the algorithm type to use for inference. Currently, only <code>image-classification</code> is supported. Required: true Type: string Valid values: <code>image-classification</code>
ServiceName	The name of the local inference service. Use the name that you specified for the <code>LocalInferenceServiceName</code> parameter when you configured the connector. Required: true Type: string
ContentType	The mime type of the input image.

Argument	Description
	Required: true
	Type: string
	Valid values: image/jpeg, image/png
Body	The content of the input image file.
	Required: true
	Type: binary

Installing MXNet dependencies on the AWS IoT Greengrass core

To use an ML Image Classification connector, you must install the dependencies for the Apache MXNet framework on the core device. The connectors use the framework to serve the ML model.

Note

These connectors are bundled with a precompiled MXNet library, so you don't need to install the MXNet framework on the core device.

AWS IoT Greengrass provides scripts to install the dependencies for the following common platforms and devices (or to use as a reference for installing them). If you're using a different platform or device, see the [MXNet documentation](#) for your configuration.

Before installing the MXNet dependencies, make sure that the required [system libraries](#) (with the specified minimum versions) are present on the device.

NVIDIA Jetson TX2

1. Install CUDA Toolkit 9.0 and cuDNN 7.0. You can follow the instructions in [the section called "Setting up other devices"](#) in the Getting Started tutorial.
2. Enable universe repositories so the connector can install community-maintained open software. For more information, see [Repositories/Ubuntu](#) in the Ubuntu documentation.
 - a. Open the `/etc/apt/sources.list` file.

- b. Make sure that the following lines are uncommented.

```
deb http://ports.ubuntu.com/ubuntu-ports/ xenial universe
deb-src http://ports.ubuntu.com/ubuntu-ports/ xenial universe
deb http://ports.ubuntu.com/ubuntu-ports/ xenial-updates universe
deb-src http://ports.ubuntu.com/ubuntu-ports/ xenial-updates universe
```

3. Save a copy of the following installation script to a file named `nvidiajtx2.sh` on the core device.

Python 3.7

```
#!/bin/bash
set -e

echo "Installing dependencies on the system..."
echo 'Assuming that universe repos are enabled and checking dependencies...'
apt-get -y update
apt-get -y dist-upgrade
apt-get install -y liblapack3 libopenblas-dev liblapack-dev libatlas-base-dev
apt-get install -y python3.7 python3.7-dev

python3.7 -m pip install --upgrade pip
python3.7 -m pip install numpy==1.15.0
python3.7 -m pip install opencv-python || echo 'Error: Unable to install
OpenCV with pip on this platform. Try building the latest OpenCV from source
(https://github.com/opencv/opencv).'

echo 'Dependency installation/upgrade complete.'
```

Note

If [OpenCV](#) does not install successfully using this script, you can try building from source. For more information, see [Installation in Linux](#) in the OpenCV documentation, or refer to other online resources for your platform.

Python 2.7

```
#!/bin/bash
```



```
set -e

echo "Installing dependencies on the system..."
echo 'Assuming that universe repos are enabled and checking dependencies...'
apt-get -y update
apt-get -y dist-upgrade
apt-get install -y liblapack3 libopenblas-dev liblapack-dev libatlas-base-dev
python-dev

echo 'Install latest pip...'
wget https://bootstrap.pypa.io/get-pip.py
python get-pip.py
rm get-pip.py

pip install numpy==1.15.0 scipy

echo 'Dependency installation/upgrade complete.'
```

4. From the directory where you saved the file, run the following command:

```
sudo nvidiajtx2.sh
```

x86_64 (Ubuntu or Amazon Linux)

1. Save a copy of the following installation script to a file named `x86_64.sh` on the core device.

Python 3.7

```
#!/bin/bash
set -e

echo "Installing dependencies on the system..."

release=$(awk -F= '/^NAME/{print $2}' /etc/os-release)

if [ "$release" == "Ubuntu" ]; then
    # Ubuntu. Supports EC2 and DeepLens. DeepLens has all the dependencies
    # installed, so
    # this is mostly to prepare dependencies on Ubuntu EC2 instance.
    apt-get -y update
    apt-get -y dist-upgrade
```

```

apt-get install -y libgfortran3 libsm6 libxext6 libxrender1
apt-get install -y python3.7 python3.7-dev
elif [ "$release" == "Amazon Linux" ]; then
    # Amazon Linux. Expect python to be installed already
    yum -y update
    yum -y upgrade

    yum install -y compat-gcc-48-libgfortran libSM libXrender libXext
else
    echo "OS Release not supported: $release"
    exit 1
fi

python3.7 -m pip install --upgrade pip
python3.7 -m pip install numpy==1.15.0
python3.7 -m pip install opencv-python || echo 'Error: Unable to install
OpenCV with pip on this platform. Try building the latest OpenCV from source
(https://github.com/opencv/opencv).'

echo 'Dependency installation/upgrade complete.'
```

Note

If [OpenCV](#) does not install successfully using this script, you can try building from source. For more information, see [Installation in Linux](#) in the OpenCV documentation, or refer to other online resources for your platform.

Python 2.7

```

#!/bin/bash
set -e

echo "Installing dependencies on the system..."

release=$(awk -F= '/^NAME/{print $2}' /etc/os-release)

if [ "$release" == "Ubuntu" ]; then
    # Ubuntu. Supports EC2 and DeepLens. DeepLens has all the dependencies
    # installed, so
    # this is mostly to prepare dependencies on Ubuntu EC2 instance.
```

```
apt-get -y update
apt-get -y dist-upgrade

apt-get install -y libgfortran3 libsm6 libxext6 libxrender1 python-dev
python-pip
elif [ "$release" == '"Amazon Linux"' ]; then
    # Amazon Linux. Expect python to be installed already
    yum -y update
    yum -y upgrade

    yum install -y compat-gcc-48-libgfortran libSM libXrender libXext python-
    pip
else
    echo "OS Release not supported: $release"
    exit 1
fi

pip install numpy==1.15.0 scipy opencv-python

echo 'Dependency installation/upgrade complete.'
```

2. From the directory where you saved the file, run the following command:

```
sudo x86_64.sh
```

Armv7 (Raspberry Pi)

1. Save a copy of the following installation script to a file named `armv7l.sh` on the core device.

Python 3.7

```
#!/bin/bash
set -e

echo "Installing dependencies on the system..."

apt-get update
apt-get -y upgrade

apt-get install -y liblapack3 libopenblas-dev liblapack-dev
apt-get install -y python3.7 python3.7-dev
```

```
python3.7 -m pip install --upgrade pip
python3.7 -m pip install numpy==1.15.0
python3.7 -m pip install opencv-python || echo 'Error: Unable to install
OpenCV with pip on this platform. Try building the latest OpenCV from source
(https://github.com/opencv/opencv).'

echo 'Dependency installation/upgrade complete.'
```

Note

If [OpenCV](#) does not install successfully using this script, you can try building from source. For more information, see [Installation in Linux](#) in the OpenCV documentation, or refer to other online resources for your platform.

Python 2.7

```
#!/bin/bash
set -e

echo "Installing dependencies on the system..."

apt-get update
apt-get -y upgrade

apt-get install -y liblapack3 libopenblas-dev liblapack-dev python-dev

# python-opencv depends on python-numpy. The latest version in the APT
# repository is python-numpy-1.8.2
# This script installs python-numpy first so that python-opencv can be
# installed, and then install the latest
# numpy-1.15.x with pip
apt-get install -y python-numpy python-opencv
dpkg --remove --force-depends python-numpy

echo 'Install latest pip...'
wget https://bootstrap.pypa.io/get-pip.py
python get-pip.py
rm get-pip.py

pip install --upgrade numpy==1.15.0 picamera scipy
```

```
echo 'Dependency installation/upgrade complete.'
```

2. From the directory where you saved the file, run the following command:

```
sudo bash armv7l.sh
```

Note

On a Raspberry Pi, using pip to install machine learning dependencies is a memory-intensive operation that can cause the device to run out of memory and become unresponsive. As a workaround, you can temporarily increase the swap size: In `/etc/dphys-swapfile`, increase the value of the `CONF_SWAPSIZE` variable and then run the following command to restart `dphys-swapfile`.

```
/etc/init.d/dphys-swapfile restart
```

Logging and troubleshooting

Depending on your group settings, event and error logs are written to CloudWatch Logs, the local file system, or both. Logs from this connector use the prefix `LocalInferenceServiceName`. If the connector behaves unexpectedly, check the connector's logs. These usually contain useful debugging information, such as a missing ML library dependency or the cause of a connector startup failure.

If the AWS IoT Greengrass group is configured to write local logs, the connector writes log files to `greengrass-root/ggc/var/log/user/region/aws/`. For more information about Greengrass logging, see [the section called “Monitoring with AWS IoT Greengrass logs”](#).

Use the following information to help troubleshoot issues with the ML Image Classification connectors.

Required system libraries

The following tabs list the system libraries required for each ML Image Classification connector.

ML Image Classification Aarch64 JTX2

Library	Minimum version
ld-linux-aarch64.so.1	GLIBC_2.17
libc.so.6	GLIBC_2.17
libcublas.so.9.0	<i>not applicable</i>
libcudart.so.9.0	<i>not applicable</i>
libcudnn.so.7	<i>not applicable</i>
libcufft.so.9.0	<i>not applicable</i>
libcurand.so.9.0	<i>not applicable</i>
libcusolver.so.9.0	<i>not applicable</i>
libgcc_s.so.1	GCC_4.2.0
libgomp.so.1	GOMP_4.0, OMP_1.0
libm.so.6	GLIBC_2.23
libpthread.so.0	GLIBC_2.17
librt.so.1	GLIBC_2.17
libstdc++.so.6	GLIBCXX_3.4.21, CXXABI_1.3.8

ML Image Classification x86_64

Library	Minimum version
ld-linux-x86-64.so.2	GCC_4.0.0
libc.so.6	GLIBC_2.4

Library	Minimum version
libgfortran.so.3	GFORTRAN_1.0
libm.so.6	GLIBC_2.23
libpthread.so.0	GLIBC_2.2.5
librt.so.1	GLIBC_2.2.5
libstdc++.so.6	CXXABI_1.3.8, GLIBCXX_3.4.21

ML Image Classification Armv7

Library	Minimum version
ld-linux-armhf.so.3	GLIBC_2.4
libc.so.6	GLIBC_2.7
libgcc_s.so.1	GCC_4.0.0
libgfortran.so.3	GFORTRAN_1.0
libm.so.6	GLIBC_2.4
libpthread.so.0	GLIBC_2.4
librt.so.1	GLIBC_2.4
libstdc++.so.6	CXXABI_1.3.8, CXXABI_ARM_1.3.3, GLIBCXX_3.4.20

Issues

Symptom	Solution
On a Raspberry Pi, the following error message is logged and you are not using	Run the following command to disable the driver:

Symptom	Solution
the camera: Failed to initialize libdc1394	<pre>sudo ln /dev/null /dev/raw1394</pre> This operation is ephemeral and the symbolic link will disappear after rebooting. Consult the manual of your OS distribution to learn how to automatically create the link up on reboot.

Licenses

The ML Image Classification connectors includes the following third-party software/licensing:

- [AWS SDK for Python \(Boto3\)](#)/Apache License 2.0
- [botocore](#)/Apache License 2.0
- [dateutil](#)/PSF License
- [docutils](#)/BSD License, GNU General Public License (GPL), Python Software Foundation License, Public Domain
- [jmespath](#)/MIT License
- [s3transfer](#)/Apache License 2.0
- [urllib3](#)/MIT License
- [Deep Neural Network Library \(DNNL\)](#)/Apache License 2.0
- [OpenMP* Runtime Library](#)/See [Intel OpenMP Runtime Library licensing](#).
- [mxnet](#)/Apache License 2.0
- [six](#)/MIT

Intel OpenMP Runtime Library licensing. The Intel® OpenMP* runtime is dual-licensed, with a commercial (COM) license as part of the Intel® Parallel Studio XE Suite products, and a BSD open source (OSS) license.

This connector is released under the [Greengrass Core Software License Agreement](#).

Changelog

The following table describes the changes in each version of the connector.

Version	Changes
2	Added the <code>MLFeedbackConnectorConfigId</code> parameter to support the use of the ML Feedback connector to upload model input data, publish predictions to an MQTT topic, and publish metrics to Amazon CloudWatch.
1	Initial release.

A Greengrass group can contain only one version of the connector at a time. For information about upgrading a connector version, see [the section called “Upgrading connector versions”](#).

See also

- [Integrate with services and protocols using connectors](#)
- [the section called “Get started with connectors \(console\)”](#)
- [the section called “Get started with connectors \(CLI\)”](#)
- [Perform machine learning inference](#)
- [Image classification algorithm](#) in the *Amazon SageMaker AI Developer Guide*

ML Object Detection connector

Warning

This connector has moved into the *extended life phase*, and AWS IoT Greengrass won't release updates that provide features, enhancements to existing features, security patches, or bug fixes. For more information, see [AWS IoT Greengrass Version 1 maintenance policy](#).

The ML Object Detection [connectors](#) provide a machine learning (ML) inference service that runs on the AWS IoT Greengrass core. This local inference service performs object detection using an object detection model compiled by the SageMaker AI Neo deep learning compiler. Two types of object detection models are supported: Single Shot Multibox Detector (SSD) and You Only Look Once (YOLO) v3. For more information, see [Object Detection Model Requirements](#).

User-defined Lambda functions use the AWS IoT Greengrass Machine Learning SDK to submit inference requests to the local inference service. The service performs local inference on an input image and returns a list of predictions for each object detected in the image. Each prediction contains an object category, a prediction confidence score, and pixel coordinates that specify a bounding box around the predicted object.

AWS IoT Greengrass provides ML Object Detection connectors for multiple platforms:

Connector	Description and ARN
ML Object Detection Aarch64 JTX2	Object detection inference service for NVIDIA Jetson TX2. Supports GPU acceleration. ARN: <code>arn:aws:greengrass: <i>region</i>::/connectors/ObjectDetectionAarch64JTX2/versions/1</code>
ML Object Detection x86_64	Object detection inference service for x86_64 platforms. ARN: <code>arn:aws:greengrass: <i>region</i>::/connectors/ObjectDetectiononx86-64/versions/1</code>
ML Object Detection ARMv7	Object detection inference service for ARMv7 platforms. ARN: <code>arn:aws:greengrass: <i>region</i>::/connectors/ObjectDetectionARMv7/versions/1</code>

Requirements

These connectors have the following requirements:

- AWS IoT Greengrass Core Software v1.9.3 or later.
- [Python](#) version 3.7 or 3.8 installed on the core device and added to the PATH environment variable.

Note

To use Python 3.8, run the following command to create a symbolic link from the the default Python 3.7 installation folder to the installed Python 3.8 binaries.

```
sudo ln -s path-to-python-3.8/python3.8 /usr/bin/python3.7
```

This configures your device to meet the Python requirement for AWS IoT Greengrass.

- Dependencies for the SageMaker AI Neo deep learning runtime installed on the core device. For more information, see [the section called “Installing Neo deep learning runtime dependencies”](#).
- An [ML resource](#) in the Greengrass group. The ML resource must reference an Amazon S3 bucket that contains an object detection model. For more information, see [Amazon S3 model sources](#).

Note

The model must be a Single Shot Multibox Detector or You Only Look Once v3 object detection model type. It must be compiled using the SageMaker AI Neo deep learning compiler. For more information, see [Object Detection Model Requirements](#).

- The [ML Feedback connector](#) added to the Greengrass group and configured. This is required only if you want to use the connector to upload model input data and publish predictions to an MQTT topic.
- [AWS IoT Greengrass Machine Learning SDK](#) v1.1.0 is required to interact with this connector.

Object detection model requirements

The ML Object Detection connectors support Single Shot multibox Detector (SSD) and You Only Look Once (YOLO) v3 object detection model types. You can use the object detection components

provided by [GluonCV](#) to train the model with your own dataset. Or, you can use pre-trained models from the GluonCV Model Zoo:

- [Pre-trained SSD model](#)
- [Pre-trained YOLO v3 model](#)

Your object detection model must be trained with 512 x 512 input images. The pre-trained models from the GluonCV Model Zoo already meet this requirement.

Trained object detection models must be compiled with the SageMaker AI Neo deep learning compiler. When compiling, make sure the target hardware matches the hardware of your Greengrass core device. For more information, see [SageMaker AI Neo](#) in the *Amazon SageMaker AI Developer Guide*.

The compiled model must be added as an ML resource ([Amazon S3 model source](#)) to the same Greengrass group as the connector.

Connector Parameters

These connectors provide the following parameters.

MLModelDestinationPath

The absolute path to the the Amazon S3 bucket that contains the Neo-compatible ML model. This is the destination path that's specified for the ML model resource.

Display name in the AWS IoT console: **Model destination path**

Required: true

Type: string

Valid pattern: .+

MLModelResourceId

The ID of the ML resource that references the source model.

Display name in the AWS IoT console: **Greengrass group ML resource**

Required: true

Type: S3MachineLearningModelResource

Valid pattern: `^[a-zA-Z0-9:_-]+$`

`LocalInferenceServiceName`

The name for the local inference service. User-defined Lambda functions invoke the service by passing the name to the `invoke_inference_service` function of the AWS IoT Greengrass Machine Learning SDK. For an example, see [the section called “Usage Example”](#).

Display name in the AWS IoT console: **Local inference service name**

Required: true

Type: string

Valid pattern: `^[a-zA-Z0-9][a-zA-Z0-9-]{1,62}$`

`LocalInferenceServiceTimeoutSeconds`

The time (in seconds) before the inference request is terminated. The minimum value is 1. The default value is 10.

Display name in the AWS IoT console: **Timeout (second)**

Required: true

Type: string

Valid pattern: `^[1-9][0-9]*$`

`LocalInferenceServiceMemoryLimitKB`

The amount of memory (in KB) that the service has access to. The minimum value is 1.

Display name in the AWS IoT console: **Memory limit**

Required: true

Type: string

Valid pattern: `^[1-9][0-9]*$`

`GPUAcceleration`

The CPU or GPU (accelerated) computing context. This property applies to the ML Image Classification Aarch64 JTX2 connector only.

Display name in the AWS IoT console: **GPU acceleration**

Required: true

Type: string

Valid values: CPU or GPU

MLFeedbackConnectorConfigId

The ID of the feedback configuration to use to upload model input data. This must match the ID of a feedback configuration defined for the [ML Feedback connector](#).

This parameter is required only if you want to use the ML Feedback connector to upload model input data and publish predictions to an MQTT topic.

Display name in the AWS IoT console: **ML Feedback connector configuration ID**

Required: false

Type: string

Valid pattern: `^[a-zA-Z0-9][a-zA-Z0-9-]{1,62}$`

Create Connector Example (AWS CLI)

The following CLI command creates a `ConnectorDefinition` with an initial version that contains an ML Object Detection connector. This example creates an instance of the ML Object Detection ARMv7l connector.

```
aws greengrass create-connector-definition --name MyGreengrassConnectors --initial-  
version '{  
  "Connectors": [  
    {  
      "Id": "MyObjectDetectionConnector",  
      "ConnectorArn": "arn:aws:greengrass:region::/connectors/  
ObjectDetectionARMv7/versions/1",  
      "Parameters": {  
        "MLModelDestinationPath": "/path-to-model",  
        "MLModelResourceId": "my-ml-resource",  
        "LocalInferenceServiceName": "objectDetection",  
        "LocalInferenceServiceTimeoutSeconds": "10",  
        "LocalInferenceServiceMemoryLimitKB": "500000",  
        "MLFeedbackConnectorConfigId" : "object-detector-random-sampling"  
      }  
    }  
  ]  
}
```

```
}
  ]
}'
```

Note

The Lambda function in these connectors have a [long-lived](#) lifecycle.

In the AWS IoT Greengrass console, you can add a connector from the group's **Connectors** page. For more information, see [the section called "Get started with connectors \(console\)"](#).

Input data

These connectors accept an image file as input. Input image files must be in jpeg or png format. For more information, see [the section called "Usage Example"](#).

These connectors don't accept MQTT messages as input data.

Output data

These connectors return a formatted list of prediction results for the identified objects in the input image:

```
{
  "prediction": [
    [
      14,
      0.9384938478469849,
      0.37763649225234985,
      0.5110225081443787,
      0.6697432398796082,
      0.8544386029243469
    ],
    [
      14,
      0.8859519958496094,
      0,
      0.43536216020584106,
      0.3314110040664673,
      0.9538808465003967
    ]
  ],
}
```

```
[
  12,
  0.04128098487854004,
  0.5976729989051819,
  0.5747185945510864,
  0.704264223575592,
  0.857937216758728
],
...
]
```

Each prediction in the list is contained in square brackets and contains six values:

- The first value represents the predicted object category for the identified object. Object categories and their corresponding values are determined when training your object detection machine learning model in the Neo deep learning compiler.
- The second value is the confidence score for the object category prediction. This represents the probability that the prediction was correct.
- The last four values correspond to pixel dimensions that represent a bounding box around the predicted object in the image.

These connectors don't publish MQTT messages as output data.

Usage Example

The following example Lambda function uses the [AWS IoT Greengrass Machine Learning SDK](#) to interact with an ML Object Detection connector.

Note

You can download the SDK from the [AWS IoT Greengrass Machine Learning SDK](#) downloads page.

The example initializes an SDK client and synchronously calls the SDK's `invoke_inference_service` function to invoke the local inference service. It passes in the algorithm type, service name, image type, and image content. Then, the example parses the service response to get the probability results (predictions).


```
import logging
from threading import Timer

import numpy as np

import greengrass_machine_learning_sdk as ml

# We assume the inference input image is provided as a local file
# to this inference client Lambda function.
with open('/test_img/test.jpg', 'rb') as f:
    content = bytearray(f.read())

client = ml.client('inference')

def infer():
    logging.info('invoking Greengrass ML Inference service')

    try:
        resp = client.invoke_inference_service(
            AlgoType='object-detection',
            ServiceName='objectDetection',
            ContentType='image/jpeg',
            Body=content
        )
    except ml.GreengrassInferenceException as e:
        logging.info('inference exception {}'.format(e.__class__.__name__, e))
        return
    except ml.GreengrassDependencyException as e:
        logging.info('dependency exception {}'.format(e.__class__.__name__, e))
        return

    logging.info('resp: {}'.format(resp))
    predictions = resp['Body'].read().decode("utf-8")
    logging.info('predictions: {}'.format(predictions))
    predictions = eval(predictions)

    # Perform business logic that relies on the predictions.

    # Schedule the infer() function to run again in ten second.
    Timer(10, infer).start()
    return

infer()
```

```
def function_handler(event, context):  
    return
```

The `invoke_inference_service` function in the AWS IoT Greengrass Machine Learning SDK accepts the following arguments.

Argument	Description
AlgoType	The name of the algorithm type to use for inference. Currently, only <code>object-detection</code> is supported. Required: true Type: string Valid values: <code>object-detection</code>
ServiceName	The name of the local inference service. Use the name that you specified for the <code>LocalInferenceServiceName</code> parameter when you configured the connector. Required: true Type: string
ContentType	The mime type of the input image. Required: true Type: string Valid values: <code>image/jpeg</code>, <code>image/png</code>
Body	The content of the input image file. Required: true Type: binary

Installing Neo deep learning runtime dependencies on the AWS IoT Greengrass core

The ML Object Detection connectors are bundled with the SageMaker AI Neo deep learning runtime (DLR). The connectors use the runtime to serve the ML model. To use these connectors, you must install the dependencies for the DLR on your core device.

Before you install the DLR dependencies, make sure that the required [system libraries](#) (with the specified minimum versions) are present on the device.

NVIDIA Jetson TX2

1. Install CUDA Toolkit 9.0 and cuDNN 7.0. You can follow the instructions in [the section called "Setting up other devices"](#) in the Getting Started tutorial.
2. Enable universe repositories so the connector can install community-maintained open software. For more information, see [Repositories/Ubuntu](#) in the Ubuntu documentation.
 - a. Open the `/etc/apt/sources.list` file.
 - b. Make sure that the following lines are uncommented.

```
deb http://ports.ubuntu.com/ubuntu-ports/ xenial universe
deb-src http://ports.ubuntu.com/ubuntu-ports/ xenial universe
deb http://ports.ubuntu.com/ubuntu-ports/ xenial-updates universe
deb-src http://ports.ubuntu.com/ubuntu-ports/ xenial-updates universe
```

3. Save a copy of the following installation script to a file named `nvidiajtx2.sh` on the core device.

```
#!/bin/bash
set -e

echo "Installing dependencies on the system..."
echo 'Assuming that universe repos are enabled and checking dependencies...'
apt-get -y update
apt-get -y dist-upgrade
apt-get install -y liblapack3 libopenblas-dev liblapack-dev libatlas-base-dev
apt-get install -y python3.7 python3.7-dev

python3.7 -m pip install --upgrade pip
python3.7 -m pip install numpy==1.15.0
```

```
python3.7 -m pip install opencv-python || echo 'Error: Unable to install OpenCV
with pip on this platform. Try building the latest OpenCV from source (https://
github.com/opencv/opencv).'

echo 'Dependency installation/upgrade complete.'
```

Note

If [OpenCV](#) does not install successfully using this script, you can try building from source. For more information, see [Installation in Linux](#) in the OpenCV documentation, or refer to other online resources for your platform.

4. From the directory where you saved the file, run the following command:

```
sudo nvidiajtx2.sh
```

x86_64 (Ubuntu or Amazon Linux)

1. Save a copy of the following installation script to a file named `x86_64.sh` on the core device.

```
#!/bin/bash
set -e

echo "Installing dependencies on the system..."

release=$(awk -F= '/^NAME/{print $2}' /etc/os-release)

if [ "$release" == '"Ubuntu"' ]; then
    # Ubuntu. Supports EC2 and DeepLens. DeepLens has all the dependencies
    installed, so
    # this is mostly to prepare dependencies on Ubuntu EC2 instance.
    apt-get -y update
    apt-get -y dist-upgrade

    apt-get install -y libgfortran3 libsm6 libxext6 libxrender1
    apt-get install -y python3.7 python3.7-dev
elif [ "$release" == '"Amazon Linux"' ]; then
    # Amazon Linux. Expect python to be installed already
    yum -y update
```

```
yum -y upgrade

yum install -y compat-gcc-48-libgfortran libSM libXrender libXext
else
    echo "OS Release not supported: $release"
    exit 1
fi

python3.7 -m pip install --upgrade pip
python3.7 -m pip install numpy==1.15.0
python3.7 -m pip install opencv-python || echo 'Error: Unable to install OpenCV
with pip on this platform. Try building the latest OpenCV from source (https://
github.com/opencv/opencv).'

echo 'Dependency installation/upgrade complete.'
```

 Note

If [OpenCV](#) does not install successfully using this script, you can try building from source. For more information, see [Installation in Linux](#) in the OpenCV documentation, or refer to other online resources for your platform.

2. From the directory where you saved the file, run the following command:

```
sudo x86_64.sh
```

ARMv7 (Raspberry Pi)

1. Save a copy of the following installation script to a file named `armv7l.sh` on the core device.

```
#!/bin/bash
set -e

echo "Installing dependencies on the system..."

apt-get update
apt-get -y upgrade

apt-get install -y liblapack3 libopenblas-dev liblapack-dev
```

```
apt-get install -y python3.7 python3.7-dev

python3.7 -m pip install --upgrade pip
python3.7 -m pip install numpy==1.15.0
python3.7 -m pip install opencv-python || echo 'Error: Unable to install OpenCV
with pip on this platform. Try building the latest OpenCV from source (https://
github.com/opencv/opencv).'

echo 'Dependency installation/upgrade complete.'
```

Note

If [OpenCV](#) does not install successfully using this script, you can try building from source. For more information, see [Installation in Linux](#) in the OpenCV documentation, or refer to other online resources for your platform.

2. From the directory where you saved the file, run the following command:

```
sudo bash armv7l.sh
```

Note

On a Raspberry Pi, using pip to install machine learning dependencies is a memory-intensive operation that can cause the device to run out of memory and become unresponsive. As a workaround, you can temporarily increase the swap size. In `/etc/dphys-swapfile`, increase the value of the `CONF_SWAPSIZE` variable and then run the following command to restart `dphys-swapfile`.

```
/etc/init.d/dphys-swapfile restart
```

Logging and troubleshooting

Depending on your group settings, event and error logs are written to CloudWatch Logs, the local file system, or both. Logs from this connector use the prefix `LocalInferenceServiceName`. If the connector behaves unexpectedly, check the connector's logs. These usually contain useful debugging information, such as a missing ML library dependency or the cause of a connector startup failure.

If the AWS IoT Greengrass group is configured to write local logs, the connector writes log files to *greengrass-root*/ggc/var/log/user/*region*/aws/. For more information about Greengrass logging, see [the section called “Monitoring with AWS IoT Greengrass logs”](#).

Use the following information to help troubleshoot issues with the ML Object Detection connectors.

Required system libraries

The following tabs list the system libraries required for each ML Object Detection connector.

ML Object Detection Aarch64 JTX2

Library	Minimum version
ld-linux-aarch64.so.1	GLIBC_2.17
libc.so.6	GLIBC_2.17
libcublas.so.9.0	<i>not applicable</i>
libcudart.so.9.0	<i>not applicable</i>
libcudnn.so.7	<i>not applicable</i>
libcufft.so.9.0	<i>not applicable</i>
libcurand.so.9.0	<i>not applicable</i>
libcusolver.so.9.0	<i>not applicable</i>
libgcc_s.so.1	GCC_4.2.0
libgomp.so.1	GOMP_4.0, OMP_1.0
libm.so.6	GLIBC_2.23
libnvinfer.so.4	<i>not applicable</i>
libnvrm_gpu.so	<i>not applicable</i>
libnvrm.so	<i>not applicable</i>

Library	Minimum version
libnvidia-fatbinaryloader.so.28.2.1	<i>not applicable</i>
libnvos.so	<i>not applicable</i>
libpthread.so.0	GLIBC_2.17
librt.so.1	GLIBC_2.17
libstdc++.so.6	GLIBCXX_3.4.21, CXXABI_1.3.8

ML Object Detection x86_64

Library	Minimum version
ld-linux-x86-64.so.2	GCC_4.0.0
libc.so.6	GLIBC_2.4
libgfortran.so.3	GFORTTRAN_1.0
libm.so.6	GLIBC_2.23
libpthread.so.0	GLIBC_2.2.5
librt.so.1	GLIBC_2.2.5
libstdc++.so.6	CXXABI_1.3.8, GLIBCXX_3.4.21

ML Object Detection ARMv7

Library	Minimum version
ld-linux-armhf.so.3	GLIBC_2.4
libc.so.6	GLIBC_2.7
libgcc_s.so.1	GCC_4.0.0

Library	Minimum version
libgfortran.so.3	GFORTRAN_1.0
libm.so.6	GLIBC_2.4
libpthread.so.0	GLIBC_2.4
librt.so.1	GLIBC_2.4
libstdc++.so.6	CXXABI_1.3.8, CXXABI_ARM_1.3.3, GLIBCXX_3.4.20

Issues

Symptom	Solution
On a Raspberry Pi, the following error message is logged and you are not using the camera: Failed to initialize libdc1394	Run the following command to disable the driver: <pre>sudo ln /dev/null /dev/raw1394</pre> This operation is ephemeral. The symbolic link disappears after you reboot. Consult the manual of your OS distribution to learn how to create the link automatically upon reboot.

Licenses

The ML Object Detection connectors include the following third-party software/licensing:

- [AWS SDK for Python \(Boto3\)](#)/Apache License 2.0
- [botocore](#)/Apache License 2.0
- [dateutil](#)/PSF License
- [docutils](#)/BSD License, GNU General Public License (GPL), Python Software Foundation License, Public Domain

- [jmespath](#)/MIT License
- [s3transfer](#)/Apache License 2.0
- [urllib3](#)/MIT License
- [Deep Learning Runtime](#)/Apache License 2.0
- [six](#)/MIT

This connector is released under the [Greengrass Core Software License Agreement](#).

See also

- [Integrate with services and protocols using connectors](#)
- [the section called “Get started with connectors \(console\)”](#)
- [the section called “Get started with connectors \(CLI\)”](#)
- [Perform machine learning inference](#)
- [Object detection algorithm](#) in the *Amazon SageMaker AI Developer Guide*

Modbus-RTU Protocol Adapter connector

The Modbus-RTU Protocol Adapter [connector](#) polls information from Modbus RTU devices that are in the AWS IoT Greengrass group.

This connector receives parameters for a Modbus RTU request from a user-defined Lambda function. It sends the corresponding request, and then publishes the response from the target device as an MQTT message.

This connector has the following versions.

Version	ARN
3	arn:aws:greengrass: <i>region</i> ::/connectors/ModbusRTUProtocolAdapter/versions/3

Version	ARN
2	arn:aws:greengrass: <i>region</i> ::/connectors/ModbusRTUProtocolAdapter/versions/2
1	arn:aws:greengrass: <i>region</i> ::/connectors/ModbusRTUProtocolAdapter/versions/1

For information about version changes, see the [Changelog](#).

Requirements

This connector has the following requirements:

Version 3

- AWS IoT Greengrass Core software v1.9.3 or later.
- [Python](#) version 3.7 or 3.8 installed on the core device and added to the PATH environment variable.

Note

To use Python 3.8, run the following command to create a symbolic link from the the default Python 3.7 installation folder to the installed Python 3.8 binaries.

```
sudo ln -s path-to-python-3.8/python3.8 /usr/bin/python3.7
```

This configures your device to meet the Python requirement for AWS IoT Greengrass.

- A physical connection between the AWS IoT Greengrass core and the Modbus devices. The core must be physically connected to the Modbus RTU network through a serial port; for example, a USB port.
- A [local device resource](#) in the Greengrass group that points to the physical Modbus serial port.
- A user-defined Lambda function that sends Modbus RTU request parameters to this connector. The request parameters must conform to expected patterns and include the IDs

and addresses of the target devices on the Modbus RTU network. For more information, see [the section called "Input data"](#).

Versions 1 - 2

- AWS IoT Greengrass Core software v1.7 or later.
- [Python](#) version 2.7 installed on the core device and added to the PATH environment variable.
- A physical connection between the AWS IoT Greengrass core and the Modbus devices. The core must be physically connected to the Modbus RTU network through a serial port; for example, a USB port.
- A [local device resource](#) in the Greengrass group that points to the physical Modbus serial port.
- A user-defined Lambda function that sends Modbus RTU request parameters to this connector. The request parameters must conform to expected patterns and include the IDs and addresses of the target devices on the Modbus RTU network. For more information, see [the section called "Input data"](#).

Connector Parameters

This connector supports the following parameters:

ModbusSerialPort-ResourceId

The ID of the local device resource that represents the physical Modbus serial port.

Note

This connector is granted read-write access to the resource.

Display name in the AWS IoT console: **Modbus serial port resource**

Required: true

Type: string

Valid pattern: .+

ModbusSerialPort

The absolute path to the physical Modbus serial port on the device. This is the source path that's specified for the Modbus local device resource.

Display name in the AWS IoT console: **Source path of Modbus serial port resource**

Required: true

Type: string

Valid pattern: .+

Create Connector Example (AWS CLI)

The following CLI command creates a ConnectorDefinition with an initial version that contains the Modbus-RTU Protocol Adapter connector.

```
aws greengrass create-connector-definition --name MyGreengrassConnectors --initial-  
version '{  
  "Connectors": [  
    {  
      "Id": "MyModbusRTUProtocolAdapterConnector",  
      "ConnectorArn": "arn:aws:greengrass:region::/connectors/  
ModbusRTUProtocolAdapter/versions/3",  
      "Parameters": {  
        "ModbusSerialPort-ResourceId": "MyLocalModbusSerialPort",  
        "ModbusSerialPort": "/path-to-port"  
      }  
    }  
  ]  
}'
```

Note

The Lambda function in this connector has a [long-lived](#) lifecycle.

In the AWS IoT Greengrass console, you can add a connector from the group's **Connectors** page. For more information, see [the section called "Get started with connectors \(console\)"](#).

Note

After you deploy the Modbus-RTU Protocol Adapter connector, you can use AWS IoT Things Graph to orchestrate interactions between devices in your group. For more information, see [Modbus](#) in the *AWS IoT Things Graph User Guide*.

Input data

This connector accepts Modbus RTU request parameters from a user-defined Lambda function on an MQTT topic. Input messages must be in JSON format.

Topic filter in subscription

```
modbus/adapter/request
```

Message properties

The request message varies based on the type of Modbus RTU request that it represents. The following properties are required for all requests:

- In the request object:
 - `operation`. The name of the operation to execute. For example, specify `"operation": "ReadCoilsRequest"` to read coils. This value must be a Unicode string. For supported operations, see [the section called “Modbus RTU requests and responses”](#).
 - `device`. The target device of the request. This value must be between 0 - 247.
- The `id` property. An ID for the request. This value is used for data deduplication and is returned as is in the `id` property of all responses, including error responses. This value must be a Unicode string.

Note

If your request includes an address field, you must specify the value as an integer. For example, `"address": 1`.

The other parameters to include in the request depend on the operation. All request parameters are required except the CRC, which is handled separately. For examples, see [the section called “Example requests and responses”](#).

Example input: Read coils request

```
{
  "request": {
    "operation": "ReadCoilsRequest",
    "device": 1,
    "address": 1,
    "count": 1
  },
  "id": "TestRequest"
}
```

Output data

This connector publishes responses to incoming Modbus RTU requests.

Topic filter in subscription

modbus/adapter/response

Message properties

The format of the response message varies based on the corresponding request and the response status. For examples, see [the section called "Example requests and responses"](#).

Note

A response for a write operation is simply an echo of the request. Although no meaningful information is returned for write responses, it's a good practice to check the status of the response.

Every response includes the following properties:

- In the response object:
 - status. The status of the request. The status can be one of the following values:
 - Success. The request was valid, sent to the Modbus RTU network, and a response was returned.

- **Exception.** The request was valid, sent to the Modbus RTU network, and an exception response was returned. For more information, see [the section called “Response status: Exception”](#).
- **No Response.** The request was invalid, and the connector caught the error before the request was sent over the Modbus RTU network. For more information, see [the section called “Response status: No response”](#).
- **device.** The device that the request was sent to.
- **operation.** The request type that was sent.
- **payload.** The response content that was returned. If the status is No Response, this object contains only an error property with the error description (for example, "error": "[Input/Output] No Response received from the remote unit").
- **The id property.** The ID of the request, used for data deduplication.

Example output: Success

```
{
  "response" : {
    "status" : "success",
    "device": 1,
    "operation": "ReadCoilsRequest",
    "payload": {
      "function_code": 1,
      "bits": [1]
    }
  },
  "id" : "TestRequest"
}
```

Example output: Failure

```
{
  "response" : {
    "status" : "fail",
    "error_message": "Internal Error",
    "error": "Exception",
    "device": 1,
    "operation": "ReadCoilsRequest",
    "payload": {
      "function_code": 129,
      "exception_code": 2
    }
  }
}
```



```
}  
},  
"id" : "TestRequest"  
}
```

For more examples, see [the section called “Example requests and responses”](#).

Modbus RTU requests and responses

This connector accepts Modbus RTU request parameters as [input data](#) and publishes responses as [output data](#).

The following common operations are supported.

Operation name in request	Function code in response
ReadCoilsRequest	01
ReadDiscreteInputsRequest	02
ReadHoldingRegistersRequest	03
ReadInputRegistersRequest	04
WriteSingleCoilRequest	05
WriteSingleRegisterRequest	06
WriteMultipleCoilsRequest	15
WriteMultipleRegistersRequest	16
MaskWriteRegisterRequest	22
ReadWriteMultipleRegistersRequest	23

Example requests and responses

The following are example requests and responses for supported operations.

Read Coils

Request example:

```
{
  "request": {
    "operation": "ReadCoilsRequest",
    "device": 1,
    "address": 1,
    "count": 1
  },
  "id": "TestRequest"
}
```

Response example:

```
{
  "response": {
    "status": "success",
    "device": 1,
    "operation": "ReadCoilsRequest",
    "payload": {
      "function_code": 1,
      "bits": [1]
    }
  },
  "id" : "TestRequest"
}
```

Read Discrete Inputs

Request example:

```
{
  "request": {
    "operation": "ReadDiscreteInputsRequest",
    "device": 1,
    "address": 1,
    "count": 1
  },
  "id": "TestRequest"
}
```

Response example:

```
{
  "response": {
    "status": "success",
    "device": 1,
    "operation": "ReadDiscreteInputsRequest",
    "payload": {
      "function_code": 2,
      "bits": [1]
    }
  },
  "id" : "TestRequest"
}
```

Read Holding Registers**Request example:**

```
{
  "request": {
    "operation": "ReadHoldingRegistersRequest",
    "device": 1,
    "address": 1,
    "count": 1
  },
  "id": "TestRequest"
}
```

Response example:

```
{
  "response": {
    "status": "success",
    "device": 1,
    "operation": "ReadHoldingRegistersRequest",
    "payload": {
      "function_code": 3,
      "registers": [20,30]
    }
  },
  "id" : "TestRequest"
}
```

```
}
```

Read Input Registers

Request example:

```
{
  "request": {
    "operation": "ReadInputRegistersRequest",
    "device": 1,
    "address": 1,
    "value": 1
  },
  "id": "TestRequest"
}
```

Write Single Coil

Request example:

```
{
  "request": {
    "operation": "WriteSingleCoilRequest",
    "device": 1,
    "address": 1,
    "value": 1
  },
  "id": "TestRequest"
}
```

Response example:

```
{
  "response": {
    "status": "success",
    "device": 1,
    "operation": "WriteSingleCoilRequest",
    "payload": {
      "function_code": 5,
      "address": 1,
      "value": true
    }
  },
}
```

```
"id" : "TestRequest"
```

Write Single Register

Request example:

```
{
  "request": {
    "operation": "WriteSingleRegisterRequest",
    "device": 1,
    "address": 1,
    "value": 1
  },
  "id": "TestRequest"
}
```

Write Multiple Coils

Request example:

```
{
  "request": {
    "operation": "WriteMultipleCoilsRequest",
    "device": 1,
    "address": 1,
    "values": [1,0,0,1]
  },
  "id": "TestRequest"
}
```

Response example:

```
{
  "response": {
    "status": "success",
    "device": 1,
    "operation": "WriteMultipleCoilsRequest",
    "payload": {
      "function_code": 15,
      "address": 1,
      "count": 4
    }
  },
}
```

```
    "id" : "TestRequest"
  }
```

Write Multiple Registers

Request example:

```
{
  "request": {
    "operation": "WriteMultipleRegistersRequest",
    "device": 1,
    "address": 1,
    "values": [20,30,10]
  },
  "id": "TestRequest"
}
```

Response example:

```
{
  "response": {
    "status": "success",
    "device": 1,
    "operation": "WriteMultipleRegistersRequest",
    "payload": {
      "function_code": 23,
      "address": 1,
      "count": 3
    }
  },
  "id" : "TestRequest"
}
```

Mask Write Register

Request example:

```
{
  "request": {
    "operation": "MaskWriteRegisterRequest",
    "device": 1,
    "address": 1,
    "and_mask": 175,
```

```
    "or_mask": 1
  },
  "id": "TestRequest"
}
```

Response example:

```
{
  "response": {
    "status": "success",
    "device": 1,
    "operation": "MaskWriteRegisterRequest",
    "payload": {
      "function_code": 22,
      "and_mask": 0,
      "or_mask": 8
    }
  },
  "id" : "TestRequest"
}
```

Read Write Multiple Registers**Request example:**

```
{
  "request": {
    "operation": "ReadWriteMultipleRegistersRequest",
    "device": 1,
    "read_address": 1,
    "read_count": 2,
    "write_address": 3,
    "write_registers": [20,30,40]
  },
  "id": "TestRequest"
}
```

Response example:

```
{
  "response": {
    "status": "success",
```

```
    "device": 1,
    "operation": "ReadWriteMultipleRegistersRequest",
    "payload": {
        "function_code": 23,
        "registers": [10,20,10,20]
    },
    "id" : "TestRequest"
}
```

Note

The registers returned in this response are the registers that are read from.

Response status: Exception

Exceptions can occur when the request format is valid, but the request is not completed successfully. In this case, the response contains the following information:

- The status is set to Exception.
- The function_code equals the function code of the request + 128.
- The exception_code contains the exception code. For more information, see Modbus exception codes.

Example:

```
{
  "response" : {
    "status" : "fail",
    "error_message": "Internal Error",
    "error": "Exception",
    "device": 1,
    "operation": "ReadCoilsRequest",
    "payload": {
      "function_code": 129,
      "exception_code": 2
    }
  },
  "id" : "TestRequest"
}
```



```
}
```

Response status: No response

This connector performs validation checks on the Modbus request. For example, it checks for invalid formats and missing fields. If the validation fails, the connector doesn't send the request. Instead, it returns a response that contains the following information:

- The status is set to No Response.
- The error contains the reason for the error.
- The error_message contains the error message.

Examples:

```
{
  "response" : {
    "status" : "fail",
    "error_message": "Invalid address field. Expected <type 'int'>, got <type 'str'>",
    "error": "No Response",
    "device": 1,
    "operation": "ReadCoilsRequest",
    "payload": {
      "error": "Invalid address field. Expected <type 'int'>, got <type 'str'>"
    }
  },
  "id" : "TestRequest"
}
```

If the request targets a nonexistent device or if the Modbus RTU network is not working, you might get a `ModbusIOException`, which uses the No Response format.

```
{
  "response" : {
    "status" : "fail",
    "error_message": "[Input/Output] No Response received from the remote unit",
    "error": "No Response",
    "device": 1,
    "operation": "ReadCoilsRequest",
    "payload": {
      "error": "[Input/Output] No Response received from the remote unit"
    }
  }
}
```

```
}  
},  
"id" : "TestRequest"  
}
```

Usage Example

Use the following high-level steps to set up an example Python 3.7 Lambda function that you can use to try out the connector.

Note

- If you use other Python runtimes, you can create a symlink from Python3.x to Python 3.7.
- The [Get started with connectors \(console\)](#) and [Get started with connectors \(CLI\)](#) topics contain detailed steps that show you how to configure and deploy an example Twilio Notifications connector.

1. Make sure you meet the [requirements](#) for the connector.
2. Create and publish a Lambda function that sends input data to the connector.

Save the [example code](#) as a PY file. Download and unzip the [AWS IoT Greengrass Core SDK for Python](#). Then, create a zip package that contains the PY file and the greengrasssdk folder at the root level. This zip package is the deployment package that you upload to AWS Lambda.

After you create the Python 3.7 Lambda function, publish a function version and create an alias.

3. Configure your Greengrass group.
 - a. Add the Lambda function by its alias (recommended). Configure the Lambda lifecycle as long-lived (or "Pinned": true in the CLI).
 - b. Add the required local device resource and grant read/write access to the Lambda function.
 - c. Add the connector and configure its [parameters](#).
 - d. Add subscriptions that allow the connector to receive [input data](#) and send [output data](#) on supported topic filters.

- Set the Lambda function as the source, the connector as the target, and use a supported input topic filter.
 - Set the connector as the source, AWS IoT Core as the target, and use a supported output topic filter. You use this subscription to view status messages in the AWS IoT console.
4. Deploy the group.
 5. In the AWS IoT console, on the **Test** page, subscribe to the output data topic to view status messages from the connector. The example Lambda function is long-lived and starts sending messages immediately after the group is deployed.

When you're finished testing, you can set the Lambda lifecycle to on-demand (or "Pinned": false in the CLI) and deploy the group. This stops the function from sending messages.

Example

The following example Lambda function sends an input message to the connector.

```
import greengrasssdk
import json

TOPIC_REQUEST = 'modbus/adapter/request'

# Creating a greengrass core sdk client
iot_client = greengrasssdk.client('iot-data')

def create_read_coils_request():
    request = {
        "request": {
            "operation": "ReadCoilsRequest",
            "device": 1,
            "address": 1,
            "count": 1
        },
        "id": "TestRequest"
    }
    return request

def publish_basic_request():
    iot_client.publish(payload=json.dumps(create_read_coils_request()),
                      topic=TOPIC_REQUEST)
```

```
publish_basic_request()

def lambda_handler(event, context):
    return
```

Licenses

The Modbus-RTU Protocol Adapter connector includes the following third-party software/licensing:

- [pymodbus](#)/BSD
- [pyserial](#)/BSD

This connector is released under the [Greengrass Core Software License Agreement](#).

Changelog

The following table describes the changes in each version of the connector.

Version	Changes
3	Upgraded the Lambda runtime to Python 3.7, which changes the runtime requirement.
2	Updated connector ARN for AWS Region support. Improved error logging.
1	Initial release.

A Greengrass group can contain only one version of the connector at a time. For information about upgrading a connector version, see [the section called “Upgrading connector versions”](#).

See also

- [Integrate with services and protocols using connectors](#)
- [the section called “Get started with connectors \(console\)”](#)

- [the section called “Get started with connectors \(CLI\)”](#)

Modbus-TCP Protocol Adapter connector

The Modbus-TCP Protocol Adapter [connector](#) collects data from local devices through the ModbusTCP protocol and publishes it to the selected StreamManager streams.

You can also use this connector with the IoT SiteWise connector and your IoT SiteWise gateway. Your gateway must supply the configuration for the connector. For more information, see [Configure a Modbus TCP source](#) in the IoT SiteWise user guide.

Note

This connector runs in [No container](#) isolation mode, so you can deploy it to a AWS IoT Greengrass group running in a Docker container.

This connector has the following versions.

Version	ARN
3	arn:aws:greengrass: <i>region</i> ::/connectors/ModbusTCPConnector/versions/3
2	arn:aws:greengrass: <i>region</i> ::/connectors/ModbusTCPConnector/versions/2
1	arn:aws:greengrass: <i>region</i> ::/connectors/ModbusTCPConnector/versions/1

For information about version changes, see the [Changelog](#).

Requirements

This connector has the following requirements:

Version 1 - 3

- AWS IoT Greengrass Core software v1.10.2 or later.
- Stream manager enabled on the AWS IoT Greengrass group.
- Java 8 installed on the core device and added to the PATH environment variable.

Note

This connector is only available in the following regions:

- ap-southeast-1
- ap-southeast-2
- eu-central-1
- eu-west-1
- us-east-1
- us-west-2
- cn-north-1

Connector Parameters

This connector supports the following parameters:

LocalStoragePath

The directory on the AWS IoT Greengrass host that the IoT SiteWise connector can write persistent data to. The default directory is `/var/sitewise`.

Display name in the AWS IoT console: **Local storage path**

Required: false

Type: string

Valid pattern: `^\s*$|\/`.

MaximumBufferSize

The maximum size in GB for IoT SiteWise disk usage. The default size is 10GB.

Display name in the AWS IoT console: **Maximum disk buffer size**

Required: false

Type: string

Valid pattern: `^\s*$| [0-9]+`

CapabilityConfiguration

The set of Modbus TCP collector configurations that the connector collects data from and connects to.

Display name in the AWS IoT console: **CapabilityConfiguration**

Required: false

Type: A well-formed JSON string that defines the set of supported feedback configurations.

The following is an example of a CapabilityConfiguration:

```
{
  "sources": [
    {
      "type": "ModBusTCPSource",
      "name": "SourceName1",
      "measurementDataStreamPrefix": "SourceName1_Prefix",
      "destination": {
        "type": "StreamManager",
        "streamName": "SiteWise_Stream_1",
        "streamBufferSize": 8
      },
      "endpoint": {
        "ipAddress": "127.0.0.1",
        "port": 8081,
        "unitId": 1
      },
      "propertyGroups": [
        {
          "name": "GroupName",
          "tagPathDefinitions": [
            {
              "type": "ModBusTCPAddress",
              "tag": "TT-001",
            }
          ]
        }
      ]
    }
  ]
}
```

```

        "address": "30001",
        "size": 2,
        "srcDataType": "float",
        "transformation": "byteWordSwap",
        "dstDataType": "double"
      }
    ],
    "scanMode": {
      "type": "POLL",
      "rate": 100
    }
  }
]
}

```

Create Connector Example (AWS CLI)

The following CLI command creates a ConnectorDefinition with an initial version that contains the Modbus-TCP Protocol Adapter connector.

```

aws greengrass create-connector-definition --name MyGreengrassConnectors --initial-
version '
{
  "Connectors": [
    {
      "Id": "MyModbusTCPConnector",
      "ConnectorArn": "arn:aws:greengrass:region::/connectors/ModbusTCP/
versions/3",
      "Parameters": {
        "capability_configuration": "{\"version\":1,\"namespace\":
        \"iotsitewise:modbuscollector:1\", \"configuration\": {\"sources\": [{\"type
        \": \"ModBusTCPSource\", \"name\": \"SourceName1\", \"measurementDataStreamPrefix
        \": \"\", \"endpoint\": {\"ipAddress\": \"127.0.0.1\", \"port\": 8081, \"unitId\": 1},
        \"propertyGroups\": [{\"name\": \"PropertyGroupName\", \"tagPathDefinitions\": [{\"type
        \": \"ModBusTCPAddress\", \"tag\": \"TT-001\", \"address\": \"30001\", \"size\": 2,
        \"srcDataType\": \"hexdump\", \"transformation\": \"noSwap\", \"dstDataType\": \"string
        \"}], \"scanMode\": {\"rate\": 200, \"type\": \"POLL\"}}], \"destination\": {\"type\":
        \"StreamManager\", \"streamName\": \"SiteWise_Stream\", \"streamBufferSize\": 10},
        \"minimumInterRequestDuration\": 200}}]}\"
      }
    }
  ]
}
'

```



```
    }
  ]
}'
```

Note

The Lambda function in this connector has a [long-lived](#) lifecycle.

Input data

This connector doesn't accept MQTT messages as input data.

Output data

This connector publishes data to `StreamManager`. You must configure the destination message stream. The output messages are of the following structure:

```
{
  "alias": "string",
  "messages": [
    {
      "name": "string",
      "value": boolean|double|integer|string,
      "timestamp": number,
      "quality": "string"
    }
  ]
}
```

Licenses

The Modbus-TCP Protocol Adapter connector includes the following third-party software/licensing:

- [Digital Petri](#) Modbus

This connector is released under the [Greengrass Core Software License Agreement](#).

Changelog

The following table describes the changes in each version of the connector.

Version	Changes	Date
3 (recommended)	This version contains bug fixes.	December 22, 2021
2	Added support for ASCII, UTF8, and ISO8859 encoded source strings.	May 24, 2021
1	Initial release.	December 15, 2020

A Greengrass group can contain only one version of the connector at a time. For information about upgrading a connector version, see [the section called “Upgrading connector versions”](#).

See also

- [Integrate with services and protocols using connectors](#)
- [the section called “Get started with connectors \(console\)”](#)
- [the section called “Get started with connectors \(CLI\)”](#)

Raspberry Pi GPIO connector

Warning

This connector has moved into the *extended life phase*, and AWS IoT Greengrass won't release updates that provide features, enhancements to existing features, security patches, or bug fixes. For more information, see [AWS IoT Greengrass Version 1 maintenance policy](#).

The Raspberry Pi GPIO [connector](#) controls general-purpose input/output (GPIO) pins on a Raspberry Pi core device.

This connector polls input pins at a specified interval and publishes state changes to MQTT topics. It also accepts read and write requests as MQTT messages from user-defined Lambda functions. Write requests are used to set the pin to high or low voltage.

The connector provides parameters that you use to designate input and output pins. This behavior is configured before group deployment. It can't be changed at runtime.

- Input pins can be used to receive data from peripheral devices.
- Output pins can be used to control peripherals or send data to peripherals.

You can use this connector for many scenarios, such as:

- Controlling green, yellow, and red LED lights for a traffic light.
- Controlling a fan (attached to an electrical relay) based on data from a humidity sensor.
- Alerting employees in a retail store when customers press a button.
- Using a smart light switch to control other IoT devices.

 Note

This connector is not suitable for applications that have real-time requirements. Events with short durations might be missed.

This connector has the following versions.

Version	ARN
3	arn:aws:greengrass: <i>region</i> ::/ connectors/RaspberryPiGPIO/ versions/3
2	arn:aws:greengrass: <i>region</i> ::/ connectors/RaspberryPiGPIO/ versions/2
1	arn:aws:greengrass: <i>region</i> ::/ connectors/RaspberryPiGPIO/ versions/1

For information about version changes, see the [Changelog](#).

Requirements

This connector has the following requirements:

Version 3

- AWS IoT Greengrass Core software v1.9.3 or later.
- [Python](#) version 3.7 installed on the core device and added to the PATH environment variable.
- Raspberry Pi 4 Model B, or Raspberry Pi 3 Model B/B+. You must know the pin sequence of your Raspberry Pi. For more information, see [the section called “GPIO Pin sequence”](#).
- A [local device resource](#) in the Greengrass group that points to /dev/gpiomem on the Raspberry Pi. If you create the resource in the console, you must select the **Automatically add OS group permissions of the Linux group that owns the resource** option. In the API, set the `GroupOwnerSetting.AutoAddGroupOwner` property to `true`.
- The [RPi.GPIO](#) module installed on the Raspberry Pi. In Raspbian, this module is installed by default. You can use the following command to reinstall it:

```
sudo pip install RPi.GPIO
```

Versions 1 - 2

- AWS IoT Greengrass Core software v1.7 or later.
- [Python](#) version 2.7 installed on the core device and added to the PATH environment variable.
- Raspberry Pi 4 Model B, or Raspberry Pi 3 Model B/B+. You must know the pin sequence of your Raspberry Pi. For more information, see [the section called “GPIO Pin sequence”](#).
- A [local device resource](#) in the Greengrass group that points to /dev/gpiomem on the Raspberry Pi. If you create the resource in the console, you must select the **Automatically add OS group permissions of the Linux group that owns the resource** option. In the API, set the `GroupOwnerSetting.AutoAddGroupOwner` property to `true`.
- The [RPi.GPIO](#) module installed on the Raspberry Pi. In Raspbian, this module is installed by default. You can use the following command to reinstall it:

```
sudo pip install RPi.GPIO
```

GPIO Pin sequence

The Raspberry Pi GPIO connector references GPIO pins by the numbering scheme of the underlying System on Chip (SoC), not by the physical layout of GPIO pins. The physical ordering of pins might vary in Raspberry Pi versions. For more information, see [GPIO](#) in the Raspberry Pi documentation.

The connector can't validate that the input and output pins you configure map correctly to the underlying hardware of your Raspberry Pi. If the pin configuration is invalid, the connector returns a runtime error when it attempts to start on the device. To resolve this issue, reconfigure the connector and then redeploy.

Note

Make sure that peripherals for GPIO pins are properly wired to prevent component damage.

Connector Parameters

This connector provides the following parameters:

InputGpios

A comma-separated list of GPIO pin numbers to configure as inputs. Optionally append U to set a pin's pull-up resistor, or D to set the pull-down resistor. Example: "5, 6U, 7D".

Display name in the AWS IoT console: **Input GPIO pins**

Required: false. You must specify input pins, output pins, or both.

Type: string

Valid pattern: `^$|^[0-9]+[UD]?([,][0-9]+[UD]?)*$`

InputPollPeriod

The interval (in milliseconds) between each polling operation, which checks input GPIO pins for state changes. The minimum value is 1.

This value depends on your scenario and the type of devices that are polled. For example, a value of 50 should be fast enough to detect a button press.

Display name in the AWS IoT console: **Input GPIO polling period**

Required: false

Type: string

Valid pattern: `^$|^[1-9][0-9]*$`

OutputGpios

A comma-separated list of GPIO pin numbers to configure as outputs. Optionally append H to set a high state (1), or L to set a low state (0). Example: "8H, 9, 27L".

Display name in the AWS IoT console: **Output GPIO pins**

Required: false. You must specify input pins, output pins, or both.

Type: string

Valid pattern: `^$|^[0-9]+[HL]?([,][0-9]+[HL]?)*$`

GpioMem-ResourceId

The ID of the local device resource that represents `/dev/gpiomem`.

Note

This connector is granted read-write access to the resource.

Display name in the AWS IoT console: **Resource for `/dev/gpiomem` device**

Required: true

Type: string

Valid pattern: `.+`

Create Connector Example (AWS CLI)

The following CLI command creates a `ConnectorDefinition` with an initial version that contains the Raspberry Pi GPIO connector.

```
aws greengrass create-connector-definition --name MyGreengrassConnectors --initial-  
version '{
```

```
"Connectors": [  
  {  
    "Id": "MyRaspberryPiGPIOConnector",  
    "ConnectorArn": "arn:aws:greengrass:region::/connectors/RaspberryPiGPIO/  
versions/3",  
    "Parameters": {  
      "GpioMem-ResourceId": "my-gpio-resource",  
      "InputGpios": "5,6U,7D",  
      "InputPollPeriod": 50,  
      "OutputGpios": "8H,9,27L"  
    }  
  }  
]
```

Note

The Lambda function in this connector has a [long-lived](#) lifecycle.

In the AWS IoT Greengrass console, you can add a connector from the group's **Connectors** page. For more information, see [the section called “Get started with connectors \(console\)”](#).

Input data

This connector accepts read or write requests for GPIO pins on two MQTT topics.

- Read requests on the `gpio/+ / + / read` topic.
- Write requests on the `gpio/+ / + / write` topic.

To publish to these topics, replace the + wildcards with the core thing name and the target pin number, respectively. For example:

```
gpio/core-thing-name/gpio-number/read
```

Note

Currently, when you create a subscription that uses the Raspberry Pi GPIO connector, you must specify a value for at least one of the + wildcards in the topic.

Topic filter: `gpio/+/+/read`

Use this topic to direct the connector to read the state of the GPIO pin that's specified in the topic.

The connector publishes the response to the corresponding output topic (for example, `gpio/core-thing-name/gpio-number/state`).

Message properties

None. Messages that are sent to this topic are ignored.

Topic filter: `gpio/+/+/write`

Use this topic to send write requests to a GPIO pin. This directs the connector to set the GPIO pin that's specified in the topic to a low or high voltage.

- 0 sets the pin to low voltage.
- 1 sets the pin to high voltage.

The connector publishes the response to the corresponding output `/state` topic (for example, `gpio/core-thing-name/gpio-number/state`).

Message properties

The value 0 or 1, as an integer or string.

Example input

```
0
```

Output data

This connector publishes data to two topics:

- High or low state changes on the `gpio/+/+/state` topic.
- Errors on the `gpio/+/+/error` topic.

Topic filter: `gpio/+/+/state`

Use this topic to listen for state changes on input pins and responses for read requests. The connector returns the string "0" if the pin is in a low state, or "1" if it's in a high state.

When publishing to this topic, the connector replaces the + wildcards with the core thing name and the target pin, respectively. For example:

```
gpio/core-thing-name/gpio-number/state
```

Note

Currently, when you create a subscription that uses the Raspberry Pi GPIO connector, you must specify a value for at least one of the + wildcards in the topic.

Example output

```
0
```

Topic filter: gpio/+ /error

Use this topic to listen for errors. The connector publishes to this topic as a result of an invalid request (for example, when a state change is requested on an input pin).

When publishing to this topic, the connector replaces the + wildcard with the core thing name.

Example output

```
{
  "topic": "gpio/my-core-thing/22/write",
  "error": "Invalid GPIO operation",
  "long_description": "GPIO 22 is configured as an INPUT GPIO. Write operations
are not permitted."
}
```

Usage Example

Use the following high-level steps to set up an example Python 3.7 Lambda function that you can use to try out the connector.

Note

- If you use other Python runtimes, you can create a symlink from Python3.x to Python 3.7.

- The [Get started with connectors \(console\)](#) and [Get started with connectors \(CLI\)](#) topics contain detailed steps that show you how to configure and deploy an example Twilio Notifications connector.

1. Make sure you meet the [requirements](#) for the connector.
2. Create and publish a Lambda function that sends input data to the connector.

Save the [example code](#) as a PY file. Download and unzip the [AWS IoT Greengrass Core SDK for Python](#). Then, create a zip package that contains the PY file and the greengrasssdk folder at the root level. This zip package is the deployment package that you upload to AWS Lambda.

After you create the Python 3.7 Lambda function, publish a function version and create an alias.

3. Configure your Greengrass group.
 - a. Add the Lambda function by its alias (recommended). Configure the Lambda lifecycle as long-lived (or "Pinned": `true` in the CLI).
 - b. Add the required local device resource and grant read/write access to the Lambda function.
 - c. Add the connector and configure its [parameters](#).
 - d. Add subscriptions that allow the connector to receive [input data](#) and send [output data](#) on supported topic filters.
 - Set the Lambda function as the source, the connector as the target, and use a supported input topic filter.
 - Set the connector as the source, AWS IoT Core as the target, and use a supported output topic filter. You use this subscription to view status messages in the AWS IoT console.
4. Deploy the group.
5. In the AWS IoT console, on the **Test** page, subscribe to the output data topic to view status messages from the connector. The example Lambda function is long-lived and starts sending messages immediately after the group is deployed.

When you're finished testing, you can set the Lambda lifecycle to on-demand (or "Pinned": `false` in the CLI) and deploy the group. This stops the function from sending messages.

Example

The following example Lambda function sends an input message to the connector. This example sends read requests for a set of input GPIO pins. It shows how to construct topics using the core thing name and pin number.

```
import greengrasssdk
import json
import os

iot_client = greengrasssdk.client('iot-data')
INPUT_GPIO_S = [6, 17, 22]

thingName = os.environ['AWS_IOT_THING_NAME']

def get_read_topic(gpio_num):
    return '/'.join(['gpio', thingName, str(gpio_num), 'read'])

def get_write_topic(gpio_num):
    return '/'.join(['gpio', thingName, str(gpio_num), 'write'])

def send_message_to_connector(topic, message=''):
    iot_client.publish(topic=topic, payload=str(message))

def set_gpio_state(gpio, state):
    send_message_to_connector(get_write_topic(gpio), str(state))

def read_gpio_state(gpio):
    send_message_to_connector(get_read_topic(gpio))

def publish_basic_message():
    for i in INPUT_GPIO_S:
        read_gpio_state(i)

publish_basic_message()

def lambda_handler(event, context):
    return
```

Licenses

The Raspberry Pi GPIO; connector includes the following third-party software/licensing:

- [RPi.GPIO](#)/MIT

This connector is released under the [Greengrass Core Software License Agreement](#).

Changelog

The following table describes the changes in each version of the connector.

Version	Changes
3	Upgraded the Lambda runtime to Python 3.7, which changes the runtime requirement.
2	Updated connector ARN for AWS Region support.
1	Initial release.

A Greengrass group can contain only one version of the connector at a time. For information about upgrading a connector version, see [the section called “Upgrading connector versions”](#).

See also

- [Integrate with services and protocols using connectors](#)
- [the section called “Get started with connectors \(console\)”](#)
- [the section called “Get started with connectors \(CLI\)”](#)
- [GPIO](#) in the Raspberry Pi documentation

Serial Stream connector

Warning

This connector has moved into the *extended life phase*, and AWS IoT Greengrass won't release updates that provide features, enhancements to existing features, security patches, or bug fixes. For more information, see [AWS IoT Greengrass Version 1 maintenance policy](#).

The Serial Stream [connector](#) reads and writes to a serial port on an AWS IoT Greengrass core device.

This connector supports two modes of operation:

- **Read-On-Demand.** Receives read and write requests on MQTT topics and publishes the response of the read operation or the status of the write operation.
- **Polling-Read.** Reads from the serial port at regular intervals. This mode also supports Read-On-Demand requests.

 Note

Read requests are limited to a maximum read length of 63994 bytes. Write requests are limited to a maximum data length of 128000 bytes.

This connector has the following versions.

Version	ARN
3	arn:aws:greengrass: <i>region</i> ::/connectors/SerialStream/versions/3
2	arn:aws:greengrass: <i>region</i> ::/connectors/SerialStream/versions/2
1	arn:aws:greengrass: <i>region</i> ::/connectors/SerialStream/versions/1

For information about version changes, see the [Changelog](#).

Requirements

This connector has the following requirements:

Version 3

- AWS IoT Greengrass Core software v1.9.3 or later.
- [Python](#) version 3.7 or 3.8 installed on the core device and added to the PATH environment variable.

Note

To use Python 3.8, run the following command to create a symbolic link from the the default Python 3.7 installation folder to the installed Python 3.8 binaries.

```
sudo ln -s path-to-python-3.8/python3.8 /usr/bin/python3.7
```

This configures your device to meet the Python requirement for AWS IoT Greengrass.

- A [local device resource](#) in the Greengrass group that points to the target serial port.

Note

Before you deploy this connector, we recommend that you set up the serial port and verify that you can read and write to it.

Versions 1 - 2

- AWS IoT Greengrass Core software v1.7 or later.
- [Python](#) version 2.7 installed on the core device and added to the PATH environment variable.
- A [local device resource](#) in the Greengrass group that points to the target serial port.

Note

Before you deploy this connector, we recommend that you set up the serial port and verify that you can read and write to it.

Connector Parameters

This connector provides the following parameters:

BaudRate

The baud rate of the serial connection.

Display name in the AWS IoT console: **Baud rate**

Required: true

Type: string

Valid values: 110, 300, 600, 1200, 2400, 4800, 9600, 14400, 19200, 28800, 38400, 56000, 57600, 115200, 230400

Valid pattern: `^110$|^300$|^600$|^1200$|^2400$|^4800$|^9600$|^14400$|^19200$|^28800$|^38400$|^56000$|^57600$|^115200$|^230400$`

Timeout

The timeout (in seconds) for a read operation.

Display name in the AWS IoT console: **Timeout**

Required: true

Type: string

Valid values: 1 - 59

Valid pattern: `^([1-9]|[1-5][0-9])$`

SerialPort

The absolute path to the physical serial port on the device. This is the source path that's specified for the local device resource.

Display name in the AWS IoT console: **Serial port**

Required: true

Type: string

Valid pattern: `[/a-zA-Z0-9_-]+`

SerialPort-ResourceId

The ID of the local device resource that represents the physical serial port.

 Note

This connector is granted read-write access to the resource.

Display name in the AWS IoT console: **Serial port resource**

Required: true

Type: string

Valid pattern: [a-zA-Z0-9_-]+

PollingRead

Sets the read mode: Polling-Read or Read-On-Demand.

- For Polling-Read mode, specify `true`. In this mode, the `PollingInterval`, `PollingReadType`, and `PollingReadLength` properties are required.
- For Read-On-Demand mode, specify `false`. In this mode, the type and length values are specified in the read request.

Display name in the AWS IoT console: **Read mode**

Required: true

Type: string

Valid values: `true`, `false`

Valid pattern: ^([Tt][Rr][Uu][Ee]|[Ff][Aa][Ll][Ss][Ee])\$

PollingReadLength

The length of data (in bytes) to read in each polling read operation. This applies only when using Polling-Read mode.

Display name in the AWS IoT console: **Polling read length**

Required: `false`. This property is required when `PollingRead` is `true`.

Type: string

Valid pattern: `^([1-9][0-9]{0,3}|[1-5][0-9]{4}|6[0-2][0-9]{3}|63[0-8][0-9]{2}|639[0-8][0-9]|6399[0-4])$`

PollingReadInterval

The interval (in seconds) at which the polling read takes place. This applies only when using Polling-Read mode.

Display name in the AWS IoT console: **Polling read interval**

Required: false. This property is required when PollingRead is true.

Type: string

Valid values: 1 - 999

Valid pattern: `^([1-9]|[1-9][0-9]|[1-9][0-9][0-9])$`

PollingReadType

The type of data that the polling thread reads. This applies only when using Polling-Read mode.

Display name in the AWS IoT console: **Polling read type**

Required: false. This property is required when PollingRead is true.

Type: string

Valid values: `ascii`, `hex`

Valid pattern: `^([Aa][Ss][Cc][Ii][Ii]|[Hh][Ee][Xx])$`

RtsCts

Indicates whether to enable the RTS/CTS flow control. The default value is false. For more information, see [RTS, CTS, and RTR](#).

Display name in the AWS IoT console: **RTS/CTS flow control**

Required: false

Type: string

Valid values: `true`, `false`

Valid pattern: `^([Tt][Rr][Uu][Ee]|[Ff][Aa][Ll][Ss][Ee])$`

XonXoff

Indicates whether to enable the software flow control. The default value is `false`. For more information, see [Software flow control](#).

Display name in the AWS IoT console: **Software flow control**

Required: `false`

Type: `string`

Valid values: `true`, `false`

Valid pattern: `^([Tt][Rr][Uu][Ee]|[Ff][Aa][Ll][Ss][Ee])$`

Parity

The parity of the serial port. The default value is `N`. For more information, see [Parity](#).

Display name in the AWS IoT console: **Serial port parity**

Required: `false`

Type: `string`

Valid values: `N`, `E`, `O`, `S`, `M`

Valid pattern: `^([NEOSMneosm])$`

Create Connector Example (AWS CLI)

The following CLI command creates a `ConnectorDefinition` with an initial version that contains the Serial Stream connector. It configures the connector for Polling-Read mode.

```
aws greengrass create-connector-definition --name MyGreengrassConnectors --initial-
version '{
  "Connectors": [
    {
      "Id": "MySerialStreamConnector",
      "ConnectorArn": "arn:aws:greengrass:region::/connectors/SerialStream/
versions/3",
```

```
    "Parameters": {
      "BaudRate" : "9600",
      "Timeout" : "25",
      "SerialPort" : "/dev/serial1",
      "SerialPort-ResourceId" : "my-serial-port-resource",
      "PollingRead" : "true",
      "PollingReadLength" : "30",
      "PollingReadInterval" : "30",
      "PollingReadType" : "hex"
    }
  ]
}'
```

In the AWS IoT Greengrass console, you can add a connector from the group's **Connectors** page. For more information, see [the section called “Get started with connectors \(console\)”](#).

Input data

This connector accepts read or write requests for serial ports on two MQTT topics. Input messages must be in JSON format.

- Read requests on the `serial/+/read/#` topic.
- Write requests on the `serial/+/write/#` topic.

To publish to these topics, replace the `+` wildcard with the core thing name and `#` wildcard with the path to the serial port. For example:

```
serial/core-thing-name/read/dev/serial-port
```

Topic filter: `serial/+/read/#`

Use this topic to send on-demand read requests to a serial pin. Read requests are limited to a maximum read length of 63994 bytes.

Message properties

`readLength`

The length of data to read from the serial port.

Required: true

Type: string

Valid pattern: `^[1-9][0-9]*$`

type

The type of data to read.

Required: true

Type: string

Valid values: `ascii`, `hex`

Valid pattern: `(?i)^(ascii|hex)$`

id

An arbitrary ID for the request. This property is used to map an input request to an output response.

Required: false

Type: string

Valid pattern: `.+`

Example input

```
{
  "readLength": "30",
  "type": "ascii",
  "id": "abc123"
}
```

Topic filter: `serial/+/write/#`

Use this topic to send write requests to a serial pin. Write requests are limited to a maximum data length of 128000 bytes.

Message properties

data

The string to write to the serial port.

Required: true

Type: string

Valid pattern: `^[1-9][0-9]*$`

type

The type of data to read.

Required: true

Type: string

Valid values: `ascii`, `hex`

Valid pattern: `^(ascii|hex|ASCII|HEX)$`

id

An arbitrary ID for the request. This property is used to map an input request to an output response.

Required: false

Type: string

Valid pattern: `.+`

Example input: ASCII request

```
{
  "data": "random serial data",
  "type": "ascii",
  "id": "abc123"
}
```

Example input: hex request

```
{
  "data": "base64 encoded data",
  "type": "hex",
  "id": "abc123"
}
```

```
}
```

Output data

The connector publishes output data on two topics:

- Status information from the connector on the `serial/+/status/#` topic.
- Responses from read requests on the `serial/+/read_response/#` topic.

When publishing to this topic, the connector replaces the `+` wildcard with the core thing name and `#` wildcard with the path to the serial port. For example:

```
serial/core-thing-name/status/dev/serial-port
```

Topic filter: `serial/+/status/#`

Use this topic to listen for the status of read and write requests. If an `id` property is included in the request, it's returned in the response.

Example output: Success

```
{
  "response": {
    "status": "success"
  },
  "id": "abc123"
}
```

Example output: Failure

A failure response includes an `error_message` property that describes the error or timeout encountered while performing the read or write operation.

```
{
  "response": {
    "status": "fail",
    "error_message": "Could not write to port"
  },
  "id": "abc123"
}
```

Topic filter: serial/+/read_response/#

Use this topic to receive response data from a read operation. The response data is Base64 encoded if the type is hex.

Example output

```
{
  "data": "output of serial read operation"
  "id": "abc123"
}
```

Usage Example

Use the following high-level steps to set up an example Python 3.7 Lambda function that you can use to try out the connector.

Note

- If you use other Python runtimes, you can create a symlink from Python3.x to Python 3.7.
- The [Get started with connectors \(console\)](#) and [Get started with connectors \(CLI\)](#) topics contain detailed steps that show you how to configure and deploy an example Twilio Notifications connector.

1. Make sure you meet the [requirements](#) for the connector.
2. Create and publish a Lambda function that sends input data to the connector.

Save the [example code](#) as a PY file. Download and unzip the [AWS IoT Greengrass Core SDK for Python](#). Then, create a zip package that contains the PY file and the greengrasssdk folder at the root level. This zip package is the deployment package that you upload to AWS Lambda.

After you create the Python 3.7 Lambda function, publish a function version and create an alias.

3. Configure your Greengrass group.
 - a. Add the Lambda function by its alias (recommended). Configure the Lambda lifecycle as long-lived (or "Pinned": true in the CLI).

- b. Add the required local device resource and grant read/write access to the Lambda function.
- c. Add the connector to your group and configure its [parameters](#).
- d. Add subscriptions to the group that allow the connector to receive [input data](#) and send [output data](#) on supported topic filters.
 - Set the Lambda function as the source, the connector as the target, and use a supported input topic filter.
 - Set the connector as the source, AWS IoT Core as the target, and use a supported output topic filter. You use this subscription to view status messages in the AWS IoT console.
4. Deploy the group.
5. In the AWS IoT console, on the **Test** page, subscribe to the output data topic to view status messages from the connector. The example Lambda function is long-lived and starts sending messages immediately after the group is deployed.

When you're finished testing, you can set the Lambda lifecycle to on-demand (or "Pinned": false in the CLI) and deploy the group. This stops the function from sending messages.

Example

The following example Lambda function sends an input message to the connector.

```
import greengrasssdk
import json

TOPIC_REQUEST = 'serial/CORE_THING_NAME/write/dev/serial11'

# Creating a greengrass core sdk client
iot_client = greengrasssdk.client('iot-data')

def create_serial_stream_request():
    request = {
        "data": "TEST",
        "type": "ascii",
        "id": "abc123"
    }
    return request

def publish_basic_request():
```



```
iot_client.publish(payload=json.dumps(create_serial_stream_request()),
topic=TOPIC_REQUEST)

publish_basic_request()

def lambda_handler(event, context):
    return
```

Licenses

The Serial Stream connector includes the following third-party software/licensing:

- [pyserial](#)/BSD

This connector is released under the [Greengrass Core Software License Agreement](#).

Changelog

The following table describes the changes in each version of the connector.

Version	Changes
3	Upgraded the Lambda runtime to Python 3.7, which changes the runtime requirement.
2	Updated connector ARN for AWS Region support.
1	Initial release.

A Greengrass group can contain only one version of the connector at a time. For information about upgrading a connector version, see [the section called “Upgrading connector versions”](#).

See also

- [Integrate with services and protocols using connectors](#)
- [the section called “Get started with connectors \(console\)”](#)
- [the section called “Get started with connectors \(CLI\)”](#)

ServiceNow MetricBase Integration connector

Warning

This connector has moved into the *extended life phase*, and AWS IoT Greengrass won't release updates that provide features, enhancements to existing features, security patches, or bug fixes. For more information, see [AWS IoT Greengrass Version 1 maintenance policy](#).

The ServiceNow MetricBase Integration [connector](#) publishes time series metrics from Greengrass devices to ServiceNow MetricBase. This allows you to store, analyze, and visualize time series data from the Greengrass core environment, and act on local events.

This connector receives time series data on an MQTT topic, and publishes the data to the ServiceNow API at regular intervals.

You can use this connector to support scenarios such as:

- Create threshold-based alerts and alarms based on time series data collected from Greengrass devices.
- Use time services data from Greengrass devices with custom applications built on the ServiceNow platform.

This connector has the following versions.

Version	ARN
4	arn:aws:greengrass: <i>region</i> ::/connectors/ServiceNowMetricBaseIntegration/versions/4
3	arn:aws:greengrass: <i>region</i> ::/connectors/ServiceNowMetricBaseIntegration/versions/3
2	arn:aws:greengrass: <i>region</i> ::/connectors/ServiceNowMetricBaseIntegration/versions/2

Version	ARN
1	arn:aws:greengrass: <i>region</i> ::/connectors/ServiceNowMetricBaseIntegration/versions/1

For information about version changes, see the [Changelog](#).

Requirements

This connector has the following requirements:

Version 3 - 4

- AWS IoT Greengrass Core software v1.9.3 or later. AWS IoT Greengrass must be configured to support local secrets, as described in [Secrets Requirements](#).

Note

This requirement includes allowing access to your Secrets Manager secrets. If you're using the default Greengrass service role, Greengrass has permission to get the values of secrets with names that start with *greengrass-*.

- [Python](#) version 3.7 or 3.8 installed on the core device and added to the PATH environment variable.

Note

To use Python 3.8, run the following command to create a symbolic link from the the default Python 3.7 installation folder to the installed Python 3.8 binaries.

```
sudo ln -s path-to-python-3.8/python3.8 /usr/bin/python3.7
```

This configures your device to meet the Python requirement for AWS IoT Greengrass.

- A ServiceNow account with an activated subscription to MetricBase. In addition, a metric and metric table must be created in the account. For more information, see [MetricBase](#) in the ServiceNow documentation.

- A text type secret in AWS Secrets Manager that stores the user name and password to log in to your ServiceNow instance with basic authentication. The secret must contain "user" and "password" keys with corresponding values. For more information, see [Creating a basic secret](#) in the *AWS Secrets Manager User Guide*.
- A secret resource in the Greengrass group that references the Secrets Manager secret. For more information, see [Deploy secrets to the core](#).

Versions 1 - 2

- AWS IoT Greengrass Core software v1.7 or later. AWS IoT Greengrass must be configured to support local secrets, as described in [Secrets Requirements](#).

Note

This requirement includes allowing access to your Secrets Manager secrets. If you're using the default Greengrass service role, Greengrass has permission to get the values of secrets with names that start with *greengrass-*.

- [Python](#) version 2.7 installed on the core device and added to the PATH environment variable.
- A ServiceNow account with an activated subscription to MetricBase. In addition, a metric and metric table must be created in the account. For more information, see [MetricBase](#) in the ServiceNow documentation.
- A text type secret in AWS Secrets Manager that stores the user name and password to log in to your ServiceNow instance with basic authentication. The secret must contain "user" and "password" keys with corresponding values. For more information, see [Creating a basic secret](#) in the *AWS Secrets Manager User Guide*.
- A secret resource in the Greengrass group that references the Secrets Manager secret. For more information, see [Deploy secrets to the core](#).

Connector Parameters

This connector provides the following parameters:

Version 4

PublishInterval

The maximum number of seconds to wait between publish events to ServiceNow. The maximum value is 900.

The connector publishes to ServiceNow when PublishBatchSize is reached or PublishInterval expires.

Display name in the AWS IoT console: **Publish interval in seconds**

Required: true

Type: string

Valid values: 1 - 900

Valid pattern: [1-9] | [1-9]\d | [1-9]\d\d | 900

PublishBatchSize

The maximum number of metric values that can be batched before they are published to ServiceNow.

The connector publishes to ServiceNow when PublishBatchSize is reached or PublishInterval expires.

Display name in the AWS IoT console: **Publish batch size**

Required: true

Type: string

Valid pattern: ^[0-9]+\$

InstanceName

The name of the instance used to connect to ServiceNow.

Display name in the AWS IoT console: **Name of ServiceNow instance**

Required: true

Type: string

Valid pattern: .+

DefaultTableName

The name of the table that contains the `GlideRecord` associated with the time series `MetricBase` database. The `table` property in the input message payload can be used to override this value.

Display name in the AWS IoT console: **Name of the table to contain the metric**

Required: true

Type: string

Valid pattern: .+

MaxMetricsToRetain

The maximum number of metrics to save in memory before they are replaced with new metrics.

This limit applies when there's no connection to the internet and the connector starts to buffer the metrics to publish later. When the buffer is full, the oldest metrics are replaced by new metrics.

Note

Metrics are not saved if the host process for the connector is interrupted. For example, this can happen during group deployment or when the device restarts.

This value should be greater than the batch size and large enough to hold messages based on the incoming rate of the MQTT messages.

Display name in the AWS IoT console: **Maximum metrics to retain in memory**

Required: true

Type: string

Valid pattern: ^[0-9]+\$

AuthSecretArn

The secret in AWS Secrets Manager that stores the ServiceNow user name and password. This must be a text type secret. The secret must contain "user" and "password" keys with corresponding values.

Display name in the AWS IoT console: **ARN of auth secret**

Required: true

Type: string

Valid pattern: `arn:aws:secretsmanager:[a-z0-9\-\-]+:[0-9]{12}:secret:([a-zA-Z0-9\-\-]+/)*[a-zA-Z0-9/_+=,.\@-\-]+-[a-zA-Z0-9]+`

AuthSecretArn-ResourceId

The secret resource in the group that references the Secrets Manager secret for the ServiceNow credentials.

Display name in the AWS IoT console: **Auth token resource**

Required: true

Type: string

Valid pattern: `.+`

IsolationMode

The [containerization](#) mode for this connector. The default is GreengrassContainer, which means that the connector runs in an isolated runtime environment inside the AWS IoT Greengrass container.

Note

The default containerization setting for the group does not apply to connectors.

Display name in the AWS IoT console: **Container isolation mode**

Required: false

Type: string

Valid values: GreengrassContainer or NoContainer

Valid pattern: ^NoContainer\$|^GreengrassContainer\$

Version 1 - 3

PublishInterval

The maximum number of seconds to wait between publish events to ServiceNow. The maximum value is 900.

The connector publishes to ServiceNow when PublishBatchSize is reached or PublishInterval expires.

Display name in the AWS IoT console: **Publish interval in seconds**

Required: true

Type: string

Valid values: 1 - 900

Valid pattern: [1-9] | [1-9]\d | [1-9]\d\d | 900

PublishBatchSize

The maximum number of metric values that can be batched before they are published to ServiceNow.

The connector publishes to ServiceNow when PublishBatchSize is reached or PublishInterval expires.

Display name in the AWS IoT console: **Publish batch size**

Required: true

Type: string

Valid pattern: ^[0-9]+\$

InstanceName

The name of the instance used to connect to ServiceNow.

Display name in the AWS IoT console: **Name of ServiceNow instance**

Required: true

Type: string

Valid pattern: .+

DefaultTableName

The name of the table that contains the `GlideRecord` associated with the time series `MetricBase` database. The `table` property in the input message payload can be used to override this value.

Display name in the AWS IoT console: **Name of the table to contain the metric**

Required: true

Type: string

Valid pattern: .+

MaxMetricsToRetain

The maximum number of metrics to save in memory before they are replaced with new metrics.

This limit applies when there's no connection to the internet and the connector starts to buffer the metrics to publish later. When the buffer is full, the oldest metrics are replaced by new metrics.

Note

Metrics are not saved if the host process for the connector is interrupted. For example, this can happen during group deployment or when the device restarts.

This value should be greater than the batch size and large enough to hold messages based on the incoming rate of the MQTT messages.

Display name in the AWS IoT console: **Maximum metrics to retain in memory**

Required: true

Type: string

Valid pattern: ^[0-9]+\$

AuthSecretArn

The secret in AWS Secrets Manager that stores the ServiceNow user name and password. This must be a text type secret. The secret must contain "user" and "password" keys with corresponding values.

Display name in the AWS IoT console: **ARN of auth secret**

Required: true

Type: string

Valid pattern: `arn:aws:secretsmanager:[a-z0-9\-\+]:[0-9]{12}:secret:([a-zA-Z0-9\-\+\/])*[a-zA-Z0-9/_+=,.\@-\+]-[a-zA-Z0-9]+`

AuthSecretArn-ResourceId

The secret resource in the group that references the Secrets Manager secret for the ServiceNow credentials.

Display name in the AWS IoT console: **Auth token resource**

Required: true

Type: string

Valid pattern: `.\+`

Create Connector Example (AWS CLI)

The following CLI command creates a `ConnectorDefinition` with an initial version that contains the ServiceNow MetricBase Integration connector.

```
aws greengrass create-connector-definition --name MyGreengrassConnectors --initial-
version '{
  "Connectors": [
    {
      "Id": "MyServiceNowMetricBaseIntegrationConnector",
      "ConnectorArn": "arn:aws:greengrass:region::/connectors/
ServiceNowMetricBaseIntegration/versions/4",
      "Parameters": {
        "PublishInterval" : "10",
        "PublishBatchSize" : "50",
        "InstanceName" : "myinstance",
```

```

        "DefaultTableName" : "u_greengrass_app",
        "MaxMetricsToRetain" : "20000",
        "AuthSecretArn" : "arn:aws:secretsmanager:region:account-
id:secret:greengrass-secret-hash",
        "AuthSecretArn-ResourceId" : "MySecretResource",
        "IsolationMode" : "GreengrassContainer"
    }
}
]
}'

```

Note

The Lambda function in this connector has a [long-lived](#) lifecycle.

In the AWS IoT Greengrass console, you can add a connector from the group's **Connectors** page. For more information, see [the section called "Get started with connectors \(console\)"](#).

Input data

This connector accepts time series metrics on an MQTT topic and publishes the metrics to ServiceNow. Input messages must be in JSON format.

Topic filter in subscription

```
servicenow/metricbase/metric
```

Message properties

```
request
```

Information about the table, record, and metric. This request represents the `seriesRef` object in a time series POST request. For more information, see [Clotho Time Series API - POST](#).

Required: true

Type: object that includes the following properties:

```
subject
```

The `sys_id` of the specific record in the table.

Required: true

Type: string

`metric_name`

The metric field name.

Required: true

Type: string

`table`

The name of the table to store the record in. Specify this value to override the `DefaultTableName` parameter.

Required: false

Type: string

`value`

The value of the individual data point.

Required: true

Type: float

`timestamp`

The timestamp of the individual data point. The default value is the current time.

Required: false

Type: string

Example input

```
{
  "request": {
    "subject": "ef43c6d40a0a0b5700c77f9bf387afe3",
    "metric_name": "u_count",
    "table": "u_greengrass_app"
    "value": 1.0,
    "timestamp": "2018-10-14T10:30:00"
```

```
}  
}
```

Output data

This connector publishes status information as output data on an MQTT topic.

Topic filter in subscription

```
servicenow/metricbase/metric/status
```

Example output: Success

```
{  
  "response": {  
    "metric_name": "Errors",  
    "table_name": "GliderProd",  
    "processed_on": "2018-10-14T10:35:00",  
    "response_id": "khjKSkj132qwr23fcba",  
    "status": "success",  
    "values": [  
      {  
        "timestamp": "2016-10-14T10:30:00",  
        "value": 1.0  
      },  
      {  
        "timestamp": "2016-10-14T10:31:00",  
        "value": 1.1  
      }  
    ]  
  }  
}
```

Example output: Failure

```
{  
  "response": {  
    "error": "InvalidInputException",  
    "error_message": "metric value is invalid",  
    "status": "fail"  
  }  
}
```

Note

If the connector detects a retryable error (for example, connection errors), it retries the publish in the next batch.

Usage Example

Use the following high-level steps to set up an example Python 3.7 Lambda function that you can use to try out the connector.

Note

- If you use other Python runtimes, you can create a symlink from Python3.x to Python 3.7.
- The [Get started with connectors \(console\)](#) and [Get started with connectors \(CLI\)](#) topics contain detailed steps that show you how to configure and deploy an example Twilio Notifications connector.

1. Make sure you meet the [requirements](#) for the connector.
2. Create and publish a Lambda function that sends input data to the connector.

Save the [example code](#) as a PY file. Download and unzip the [AWS IoT Greengrass Core SDK for Python](#). Then, create a zip package that contains the PY file and the greengrasssdk folder at the root level. This zip package is the deployment package that you upload to AWS Lambda.

After you create the Python 3.7 Lambda function, publish a function version and create an alias.

3. Configure your Greengrass group.
 - a. Add the Lambda function by its alias (recommended). Configure the Lambda lifecycle as long-lived (or "Pinned": `true` in the CLI).
 - b. Add the required secret resource and grant read access to the Lambda function.
 - c. Add the connector and configure its [parameters](#).
 - d. Add subscriptions that allow the connector to receive [input data](#) and send [output data](#) on supported topic filters.

- Set the Lambda function as the source, the connector as the target, and use a supported input topic filter.
 - Set the connector as the source, AWS IoT Core as the target, and use a supported output topic filter. You use this subscription to view status messages in the AWS IoT console.
4. Deploy the group.
 5. In the AWS IoT console, on the **Test** page, subscribe to the output data topic to view status messages from the connector. The example Lambda function is long-lived and starts sending messages immediately after the group is deployed.

When you're finished testing, you can set the Lambda lifecycle to on-demand (or "Pinned": false in the CLI) and deploy the group. This stops the function from sending messages.

Example

The following example Lambda function sends an input message to the connector.

```
import greengrasssdk
import json

iot_client = greengrasssdk.client('iot-data')
SEND_TOPIC = 'servicenow/metricbase/metric'

def create_request_with_all_fields():
    return {
        "request": {
            "subject": '2efdf6badbd523803acfae441b961961',
            "metric_name": 'u_count',
            "value": 1234,
            "timestamp": '2018-10-20T20:22:20',
            "table": 'u_greengrass_metricbase_test'
        }
    }

def publish_basic_message():
    messageToPublish = create_request_with_all_fields()
    print("Message To Publish: ", messageToPublish)
    iot_client.publish(topic=SEND_TOPIC,
        payload=json.dumps(messageToPublish))

publish_basic_message()
```

```
def lambda_handler(event, context):  
    return
```

Licenses

The ServiceNow MetricBase Integration connector includes the following third-party software/licensing:

- [pysnow](#)/MIT

This connector is released under the [Greengrass Core Software License Agreement](#).

Changelog

The following table describes the changes in each version of the connector.

Version	Changes
4	Added the <code>IsolationMode</code> parameter to configure the containerization mode for the connector.
3	Upgraded the Lambda runtime to Python 3.7, which changes the runtime requirement.
2	Fix to reduce excessive logging.
1	Initial release.

A Greengrass group can contain only one version of the connector at a time. For information about upgrading a connector version, see [the section called “Upgrading connector versions”](#).

See also

- [Integrate with services and protocols using connectors](#)
- [the section called “Get started with connectors \(console\)”](#)
- [the section called “Get started with connectors \(CLI\)”](#)

SNS connector

The SNS [connector](#) publishes messages to an Amazon SNS topic. This enables web servers, email addresses, and other message subscribers to respond to events in the Greengrass group.

This connector receives SNS message information on an MQTT topic, and then sends the message to a specified SNS topic. You can optionally use custom Lambda functions to implement filtering or formatting logic on messages before they are published to this connector.

This connector has the following versions.

Version	ARN
4	arn:aws:greengrass: <i>region</i> ::/connectors/SNS/versions/4
3	arn:aws:greengrass: <i>region</i> ::/connectors/SNS/versions/3
2	arn:aws:greengrass: <i>region</i> ::/connectors/SNS/versions/2
1	arn:aws:greengrass: <i>region</i> ::/connectors/SNS/versions/1

For information about version changes, see the [Changelog](#).

Requirements

This connector has the following requirements:

Version 3 - 4

- AWS IoT Greengrass Core software v1.9.3 or later.
- [Python](#) version 3.7 or 3.8 installed on the core device and added to the PATH environment variable.

Note

To use Python 3.8, run the following command to create a symbolic link from the the default Python 3.7 installation folder to the installed Python 3.8 binaries.

```
sudo ln -s path-to-python-3.8/python3.8 /usr/bin/python3.7
```

This configures your device to meet the Python requirement for AWS IoT Greengrass.

- A configured SNS topic. For more information, see [Creating an Amazon SNS topic](#) in the *Amazon Simple Notification Service Developer Guide*.
- The [Greengrass group role](#) configured to allow the `sns:Publish` action on the target Amazon SNS topic, as shown in the following example IAM policy.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "Stmt1528133056761",
      "Action": [
        "sns:Publish"
      ],
      "Effect": "Allow",
      "Resource": [
        "arn:aws:sns:region:account-id:topic-name"
      ]
    }
  ]
}
```

This connector allows you to dynamically override the default topic in the input message payload. If your implementation uses this feature, the IAM policy must allow `sns:Publish` permission on all target topics. You can grant granular or conditional access to resources (for example, by using a wildcard `*` naming scheme).

For the group role requirement, you must configure the role to grant the required permissions and make sure the role has been added to the group. For more information, see [the section called “Manage the group role \(console\)”](#) or [the section called “Manage the group role \(CLI\)”](#).

Versions 1 - 2

- AWS IoT Greengrass Core software v1.7 or later.
- [Python](#) version 2.7 installed on the core device and added to the PATH environment variable.
- A configured SNS topic. For more information, see [Creating an Amazon SNS topic](#) in the *Amazon Simple Notification Service Developer Guide*.
- The [Greengrass group role](#) configured to allow the `sns:Publish` action on the target Amazon SNS topic, as shown in the following example IAM policy.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "Stmt1528133056761",
      "Action": [
        "sns:Publish"
      ],
      "Effect": "Allow",
      "Resource": [
        "arn:aws:sns:region:account-id:topic-name"
      ]
    }
  ]
}
```

This connector allows you to dynamically override the default topic in the input message payload. If your implementation uses this feature, the IAM policy must allow `sns:Publish` permission on all target topics. You can grant granular or conditional access to resources (for example, by using a wildcard `*` naming scheme).

For the group role requirement, you must configure the role to grant the required permissions and make sure the role has been added to the group. For more information, see [the section called “Manage the group role \(console\)”](#) or [the section called “Manage the group role \(CLI\)”](#).

Connector Parameters

This connector provides the following parameters:

Version 4

DefaultSNSArn

The ARN of the default SNS topic to publish messages to. The destination topic can be overridden by the `sns_topic_arn` property in the input message payload.

Note

The group role must allow `sns:Publish` permission to all target topics. For more information, see [the section called “Requirements”](#).

Display name in the AWS IoT console: **Default SNS topic ARN**

Required: true

Type: string

Valid pattern: `arn:aws:sns:([a-z]{2}-[a-z]+-\d{1}):(\d{12}):([a-zA-Z0-9-_\])+)`

IsolationMode

The [containerization](#) mode for this connector. The default is `GreengrassContainer`, which means that the connector runs in an isolated runtime environment inside the AWS IoT Greengrass container.

Note

The default containerization setting for the group does not apply to connectors.

Display name in the AWS IoT console: **Container isolation mode**

Required: false

Type: string

Valid values: `GreengrassContainer` or `NoContainer`

Valid pattern: `^NoContainer$|^GreengrassContainer$`

Versions 1 - 3

DefaultSNSArn

The ARN of the default SNS topic to publish messages to. The destination topic can be overridden by the `sns_topic_arn` property in the input message payload.

Note

The group role must allow `sns:Publish` permission to all target topics. For more information, see [the section called "Requirements"](#).

Display name in the AWS IoT console: **Default SNS topic ARN**

Required: true

Type: string

Valid pattern: `arn:aws:sns:([a-z]{2}-[a-z]+\d{1}):(\d{12}):([a-zA-Z0-9-_\d]{1,256})$`

Create Connector Example (AWS CLI)

The following CLI command creates a `ConnectorDefinition` with an initial version that contains the SNS connector.

```
aws greengrass create-connector-definition --name MyGreengrassConnectors --initial-  
version '{  
  "Connectors": [  
    {  
      "Id": "MySNSConnector",  
      "ConnectorArn": "arn:aws:greengrass:region::/connectors/SNS/versions/4",  
      "Parameters": {  
        "DefaultSNSArn": "arn:aws:sns:region:account-id:topic-name",  
        "IsolationMode" : "GreengrassContainer"  
      }  
    }  
  ]  
}'
```

In the AWS IoT Greengrass console, you can add a connector from the group's **Connectors** page. For more information, see [the section called "Get started with connectors \(console\)"](#).

Input data

This connector accepts SNS message information on an MQTT topic, and then publishes the message as is to the target SNS topic. Input messages must be in JSON format.

Topic filter in subscription

sns/message

Message properties

request

Information about the message to send to the SNS topic.

Required: true

Type: object that includes the following properties:

message

The content of the message as a string or in JSON format. For examples, see [Example input](#).

To send JSON, the `message_structure` property must be set to `json` and the message must be a string-encoded JSON object that contains a default key.

Required: true

Type: string

Valid pattern: `.*`

subject

The subject of the message.

Required: false

Type: ASCII text, up to 100 characters. This must begin with a letter, number, or punctuation mark. This must not include line breaks or control characters.

Valid pattern: `.*`

`sns_topic_arn`

The ARN of the SNS topic to publish messages to. If specified, the connector publishes to this topic instead of the default topic.

 Note

The group role must allow `sns:Publish` permission to any target topics. For more information, see [the section called “Requirements”](#).

Required: `false`

Type: `string`

Valid pattern: `arn:aws:sns:([a-z]{2}-[a-z]+-\d{1}):(\d{12}):([a-zA-Z0-9-_\+]*)$`

`message_structure`

The structure of the message.

Required: `false`. This must be specified to send a JSON message.

Type: `string`

Valid values: `json`

`id`

An arbitrary ID for the request. This property is used to map an input request to an output response. When specified, the `id` property in the response object is set to this value. If you don't use this feature, you can omit this property or specify an empty string.

Required: `false`

Type: `string`

Valid pattern: `.*`

Limits

The message size is bounded by a maximum SNS message size of 256 KB.

Example input: String message

This example sends a string message. It specifies the optional `sns_topic_arn` property, which overrides the default destination topic.

```
{
  "request": {
    "subject": "Message subject",
    "message": "Message data",
    "sns_topic_arn": "arn:aws:sns:region:account-id:topic2-name"
  },
  "id": "request123"
}
```

Example input: JSON message

This example sends a message as a string encoded JSON object that includes the default key.

```
{
  "request": {
    "subject": "Message subject",
    "message": "{ \"default\": \"Message data\" }",
    "message_structure": "json"
  },
  "id": "request123"
}
```

Output data

This connector publishes status information as output data on an MQTT topic.

Topic filter in subscription

`sns/message/status`

Example output: Success

```
{
  "response": {
    "sns_message_id": "f80a81bc-f44c-56f2-a0f0-d5af6a727c8a",
    "status": "success"
  },
}
```



```
{
  "id": "request123"
}
```

Example output: Failure

```
{
  "response" : {
    "error": "InvalidInputException",
    "error_message": "SNS Topic Arn is invalid",
    "status": "fail"
  },
  "id": "request123"
}
```

Usage Example

Use the following high-level steps to set up an example Python 3.7 Lambda function that you can use to try out the connector.

Note

- If you use other Python runtimes, you can create a symlink from Python3.x to Python 3.7.
- The [Get started with connectors \(console\)](#) and [Get started with connectors \(CLI\)](#) topics contain detailed steps that show you how to configure and deploy an example Twilio Notifications connector.

1. Make sure you meet the [requirements](#) for the connector.

For the group role requirement, you must configure the role to grant the required permissions and make sure the role has been added to the group. For more information, see [the section called “Manage the group role \(console\)”](#) or [the section called “Manage the group role \(CLI\)”](#).

2. Create and publish a Lambda function that sends input data to the connector.

Save the [example code](#) as a PY file. Download and unzip the [AWS IoT Greengrass Core SDK for Python](#). Then, create a zip package that contains the PY file and the greengrasssdk folder at the root level. This zip package is the deployment package that you upload to AWS Lambda.

After you create the Python 3.7 Lambda function, publish a function version and create an alias.

3. Configure your Greengrass group.
 - a. Add the Lambda function by its alias (recommended). Configure the Lambda lifecycle as long-lived (or "Pinned": `true` in the CLI).
 - b. Add the connector and configure its [parameters](#).
 - c. Add subscriptions that allow the connector to receive [input data](#) and send [output data](#) on supported topic filters.
 - Set the Lambda function as the source, the connector as the target, and use a supported input topic filter.
 - Set the connector as the source, AWS IoT Core as the target, and use a supported output topic filter. You use this subscription to view status messages in the AWS IoT console.
4. Deploy the group.
5. In the AWS IoT console, on the **Test** page, subscribe to the output data topic to view status messages from the connector. The example Lambda function is long-lived and starts sending messages immediately after the group is deployed.

When you're finished testing, you can set the Lambda lifecycle to on-demand (or "Pinned": `false` in the CLI) and deploy the group. This stops the function from sending messages.

Example

The following example Lambda function sends an input message to the connector.

```
import greengrasssdk
import time
import json

iot_client = greengrasssdk.client('iot-data')
send_topic = 'sns/message'

def create_request_with_all_fields():
    return {
        "request": {
            "message": "Message from SNS Connector Test"
        },
    },
```

```
        "id" : "req_123"
    }

def publish_basic_message():
    messageToPublish = create_request_with_all_fields()
    print("Message To Publish: ", messageToPublish)
    iot_client.publish(topic=send_topic,
        payload=json.dumps(messageToPublish))

publish_basic_message()

def lambda_handler(event, context):
    return
```

Licenses

The SNS connector includes the following third-party software/licensing:

- [AWS SDK for Python \(Boto3\)](#)/Apache License 2.0
- [botocore](#)/Apache License 2.0
- [dateutil](#)/PSF License
- [docutils](#)/BSD License, GNU General Public License (GPL), Python Software Foundation License, Public Domain
- [jmespath](#)/MIT License
- [s3transfer](#)/Apache License 2.0
- [urllib3](#)/MIT License

This connector is released under the [Greengrass Core Software License Agreement](#).

Changelog

The following table describes the changes in each version of the connector.

Version	Changes
4	Added the <code>IsolationMode</code> parameter to configure the containerization mode for the connector.

Version	Changes
3	Upgraded the Lambda runtime to Python 3.7, which changes the runtime requirement.
2	Fix to reduce excessive logging.
1	Initial release.

A Greengrass group can contain only one version of the connector at a time. For information about upgrading a connector version, see [the section called “Upgrading connector versions”](#).

See also

- [Integrate with services and protocols using connectors](#)
- [the section called “Get started with connectors \(console\)”](#)
- [the section called “Get started with connectors \(CLI\)”](#)
- [Publish action](#) in the Boto 3 documentation
- [What is Amazon Simple Notification Service?](#) in the *Amazon Simple Notification Service Developer Guide*

Splunk Integration connector

Warning

This connector has moved into the *extended life phase*, and AWS IoT Greengrass won't release updates that provide features, enhancements to existing features, security patches, or bug fixes. For more information, see [AWS IoT Greengrass Version 1 maintenance policy](#).

The Splunk Integration [connector](#) publishes data from Greengrass devices to Splunk. This allows you to use Splunk to monitor and analyze the Greengrass core environment, and act on local events. The connector integrates with HTTP Event Collector (HEC). For more information, see [Introduction to Splunk HTTP Event Collector](#) in the Splunk documentation.

This connector receives logging and event data on an MQTT topic and publishes the data as is to the Splunk API.

You can use this connector to support industrial scenarios, such as:

- Operators can use periodic data from actuators and sensors (for example, temperature, pressure, and water readings) to initiate alarms when values exceed certain thresholds.
- Developers use data collected from industrial machinery to build ML models that can monitor the equipment for potential issues.

This connector has the following versions.

Version	ARN
4	arn:aws:greengrass: <i>region</i> ::/connectors/SplunkIntegration/versions/4
3	arn:aws:greengrass: <i>region</i> ::/connectors/SplunkIntegration/versions/3
2	arn:aws:greengrass: <i>region</i> ::/connectors/SplunkIntegration/versions/2
1	arn:aws:greengrass: <i>region</i> ::/connectors/SplunkIntegration/versions/1

For information about version changes, see the [Changelog](#).

Requirements

This connector has the following requirements:

Version 3 - 4

- AWS IoT Greengrass Core software v1.9.3 or later. AWS IoT Greengrass must be configured to support local secrets, as described in [Secrets Requirements](#).

Note

This requirement includes allowing access to your Secrets Manager secrets. If you're using the default Greengrass service role, Greengrass has permission to get the values of secrets with names that start with *greengrass-*.

- [Python](#) version 3.7 or 3.8 installed on the core device and added to the PATH environment variable.

Note

To use Python 3.8, run the following command to create a symbolic link from the the default Python 3.7 installation folder to the installed Python 3.8 binaries.

```
sudo ln -s path-to-python-3.8/python3.8 /usr/bin/python3.7
```

This configures your device to meet the Python requirement for AWS IoT Greengrass.

- The HTTP Event Collector functionality must be enabled in Splunk. For more information, see [Set up and use HTTP eEvent Collector in Splunk Web](#) in the Splunk documentation.
- A text type secret in AWS Secrets Manager that stores your Splunk HTTP Event Collector token. For more information, see [About event collector tokens](#) in the Splunk documentation and [Creating a basic secret](#) in the *AWS Secrets Manager User Guide*.

Note

To create the secret in the Secrets Manager console, enter your token on the **Plaintext** tab. Don't include quotation marks or other formatting. In the API, specify the token as the value for the `SecretString` property.

- A secret resource in the Greengrass group that references the Secrets Manager secret. For more information, see [Deploy secrets to the core](#).

Versions 1 - 2

- AWS IoT Greengrass Core software v1.7 or later. AWS IoT Greengrass must be configured to support local secrets, as described in [Secrets Requirements](#).

Note

This requirement includes allowing access to your Secrets Manager secrets. If you're using the default Greengrass service role, Greengrass has permission to get the values of secrets with names that start with *greengrass-*.

- [Python](#) version 2.7 installed on the core device and added to the PATH environment variable.
- The HTTP Event Collector functionality must be enabled in Splunk. For more information, see [Set up and use HTTP eEvent Collector in Splunk Web](#) in the Splunk documentation.
- A text type secret in AWS Secrets Manager that stores your Splunk HTTP Event Collector token. For more information, see [About event collector tokens](#) in the Splunk documentation and [Creating a basic secret](#) in the *AWS Secrets Manager User Guide*.

Note

To create the secret in the Secrets Manager console, enter your token on the **Plaintext** tab. Don't include quotation marks or other formatting. In the API, specify the token as the value for the `SecretString` property.

- A secret resource in the Greengrass group that references the Secrets Manager secret. For more information, see [Deploy secrets to the core](#).

Connector Parameters

This connector provides the following parameters:

Version 4

`SplunkEndpoint`

The endpoint of your Splunk instance. This value must contain the protocol, hostname, and port.

Display name in the AWS IoT console: **Splunk endpoint**

Required: `true`

Type: `string`

Valid pattern: `^(http:\\\\|https:\\\\)?[a-z0-9]+([-\\.]{1}[a-z0-9]+)*.[a-z]{2,5}(:[0-9]{1,5})?(\\/.*)?$`

MemorySize

The amount of memory (in KB) to allocate to the connector.

Display name in the AWS IoT console: **Memory size**

Required: true

Type: string

Valid pattern: `^[0-9]+$`

SplunkQueueSize

The maximum number of items to save in memory before the items are submitted or discarded. When this limit is met, the oldest items in the queue are replaced with newer items. This limit typically applies when there's no connection to the internet.

Display name in the AWS IoT console: **Maximum items to retain**

Required: true

Type: string

Valid pattern: `^[0-9]+$`

SplunkFlushIntervalSeconds

The interval (in seconds) for publishing received data to Splunk HEC. The maximum value is 900. To configure the connector to publish items as they are received (without batching), specify 0.

Display name in the AWS IoT console: **Splunk publish interval**

Required: true

Type: string

Valid pattern: `[0-9] | [1-9]\\d | [1-9]\\d\\d | 900`

SplunkTokenSecretArn

The secret in AWS Secrets Manager that stores the Splunk token. This must be a text type secret.

Display name in the AWS IoT console: **ARN of Splunk auth token secret**

Required: true

Type: string

Valid pattern: `arn:aws:secretsmanager:[a-z]{2}-[a-z]+-\d{1}:\d{12}?:secret:[a-zA-Z0-9-_\d]{1,64}`

SplunkTokenSecretArn-ResourceId

The secret resource in the Greengrass group that references the Splunk secret.

Display name in the AWS IoT console: **Splunk auth token resource**

Required: true

Type: string

Valid pattern: `.*`

SplunkCustomCALocation

The file path of the custom certificate authority (CA) for Splunk (for example, `/etc/ssl/certs/splunk.crt`).

Display name in the AWS IoT console: **Splunk custom certificate authority location**

Required: false

Type: string

Valid pattern: `^$|/.*`

IsolationMode

The [containerization](#) mode for this connector. The default is `GreengrassContainer`, which means that the connector runs in an isolated runtime environment inside the AWS IoT Greengrass container.

 Note

The default containerization setting for the group does not apply to connectors.

Display name in the AWS IoT console: **Container isolation mode**

Required: false

Type: string

Valid values: GreengrassContainer or NoContainer

Valid pattern: ^NoContainer\$|^GreengrassContainer\$

Version 1 - 3

SplunkEndpoint

The endpoint of your Splunk instance. This value must contain the protocol, hostname, and port.

Display name in the AWS IoT console: **Splunk endpoint**

Required: true

Type: string

Valid pattern: ^(http:\\\\|https:\\\\)?[a-z0-9]+([-.]{1}[a-z0-9]+)*.[a-z]{2,5}(:[0-9]{1,5})?(\\.|*)?\$

MemorySize

The amount of memory (in KB) to allocate to the connector.

Display name in the AWS IoT console: **Memory size**

Required: true

Type: string

Valid pattern: `^[0-9]+$`

`SplunkQueueSize`

The maximum number of items to save in memory before the items are submitted or discarded. When this limit is met, the oldest items in the queue are replaced with newer items. This limit typically applies when there's no connection to the internet.

Display name in the AWS IoT console: **Maximum items to retain**

Required: `true`

Type: `string`

Valid pattern: `^[0-9]+$`

`SplunkFlushIntervalSeconds`

The interval (in seconds) for publishing received data to Splunk HEC. The maximum value is 900. To configure the connector to publish items as they are received (without batching), specify 0.

Display name in the AWS IoT console: **Splunk publish interval**

Required: `true`

Type: `string`

Valid pattern: `[0-9] | [1-9]\d | [1-9]\d\d | 900`

`SplunkTokenSecretArn`

The secret in AWS Secrets Manager that stores the Splunk token. This must be a text type secret.

Display name in the AWS IoT console: **ARN of Splunk auth token secret**

Required: `true`

Type: `string`

Valid pattern: `arn:aws:secretsmanager:[a-z]{2}-[a-z]+-\d{1}:\d{12}?:secret:[a-zA-Z0-9-_\d]{32}`

SplunkTokenSecretArn-ResourceId

The secret resource in the Greengrass group that references the Splunk secret.

Display name in the AWS IoT console: **Splunk auth token resource**

Required: true

Type: string

Valid pattern: .+

SplunkCustomCALocation

The file path of the custom certificate authority (CA) for Splunk (for example, /etc/ssl/certs/splunk.crt).

Display name in the AWS IoT console: **Splunk custom certificate authority location**

Required: false

Type: string

Valid pattern: ^\$|/.*

Create Connector Example (AWS CLI)

The following CLI command creates a `ConnectorDefinition` with an initial version that contains the Splunk Integration connector.

```
aws greengrass create-connector-definition --name MyGreengrassConnectors --initial-  
version '{  
  "Connectors": [  
    {  
      "Id": "MySplunkIntegrationConnector",  
      "ConnectorArn": "arn:aws:greengrass:region::/connectors/SplunkIntegration/  
versions/4",  
      "Parameters": {  
        "SplunkEndpoint": "https://myinstance.cloud.splunk.com:8088",  
        "MemorySize": 200000,  
        "SplunkQueueSize": 10000,  
        "SplunkFlushIntervalSeconds": 5,  
        "SplunkTokenSecretArn": "arn:aws:secretsmanager:region:account-  
id:secret:greengrass-secret-hash",
```

```

        "SplunkTokenSecretArn-ResourceId": "MySplunkResource",
        "IsolationMode" : "GreengrassContainer"
    }
}
]
}'

```

Note

The Lambda function in this connector has a [long-lived](#) lifecycle.

In the AWS IoT Greengrass console, you can add a connector from the group's **Connectors** page. For more information, see [the section called “Get started with connectors \(console\)”](#).

Input data

This connector accepts logging and event data on an MQTT topic and publishes the received data as is to the Splunk API. Input messages must be in JSON format.

Topic filter in subscription

splunk/logs/put

Message properties

request

The event data to send to the Splunk API. Events must meet the specifications of the [services/collector](#) API.

Required: true

Type: object. Only the event property is required.

id

An arbitrary ID for the request. This property is used to map an input request to an output status.

Required: false

Type: string

Limits

All limits that are imposed by the Splunk API apply when using this connector. For more information, see [services/collector](#).

Example input

```
{
  "request": {
    "event": "some event",
    "fields": {
      "severity": "INFO",
      "category": [
        "value1",
        "value2"
      ]
    }
  },
  "id": "request123"
}
```

Output data

This connector publishes output data on two topics:

- Status information on the `splunk/logs/put/status` topic.
- Errors on the `splunk/logs/put/error` topic.

Topic filter: `splunk/logs/put/status`

Use this topic to listen for the status of the requests. Each time that the connector sends a batch of received data to the Splunk API, it publishes a list of the IDs of the requests that succeeded and failed.

Example output

```
{
  "response": {
    "succeeded": [
      "request123",
      ...
    ],

```

```
        "failed": [  
            "request789",  
            ...  
        ]  
    }  
}
```

Topic filter: splunk/logs/put/error

Use this topic to listen for errors from the connector. The `error_message` property that describes the error or timeout encountered while processing the request.

Example output

```
{  
  "response": {  
    "error": "UnauthorizedException",  
    "error_message": "invalid splunk token",  
    "status": "fail"  
  }  
}
```

Note

If the connector detects a retryable error (for example, connection errors), it retries the publish in the next batch.

Usage Example

Use the following high-level steps to set up an example Python 3.7 Lambda function that you can use to try out the connector.

Note

- If you use other Python runtimes, you can create a symlink from `Python3.x` to `Python 3.7`.
- The [Get started with connectors \(console\)](#) and [Get started with connectors \(CLI\)](#) topics contain detailed steps that show you how to configure and deploy an example Twilio Notifications connector.

1. Make sure you meet the [requirements](#) for the connector.
2. Create and publish a Lambda function that sends input data to the connector.

Save the [example code](#) as a PY file. Download and unzip the [AWS IoT Greengrass Core SDK for Python](#). Then, create a zip package that contains the PY file and the greengrasssdk folder at the root level. This zip package is the deployment package that you upload to AWS Lambda.

After you create the Python 3.7 Lambda function, publish a function version and create an alias.

3. Configure your Greengrass group.
 - a. Add the Lambda function by its alias (recommended). Configure the Lambda lifecycle as long-lived (or "Pinned": `true` in the CLI).
 - b. Add the required secret resource and grant read access to the Lambda function.
 - c. Add the connector and configure its [parameters](#).
 - d. Add subscriptions that allow the connector to receive [input data](#) and send [output data](#) on supported topic filters.
 - Set the Lambda function as the source, the connector as the target, and use a supported input topic filter.
 - Set the connector as the source, AWS IoT Core as the target, and use a supported output topic filter. You use this subscription to view status messages in the AWS IoT console.
4. Deploy the group.
5. In the AWS IoT console, on the **Test** page, subscribe to the output data topic to view status messages from the connector. The example Lambda function is long-lived and starts sending messages immediately after the group is deployed.

When you're finished testing, you can set the Lambda lifecycle to on-demand (or "Pinned": `false` in the CLI) and deploy the group. This stops the function from sending messages.

Example

The following example Lambda function sends an input message to the connector.

```
import greengrasssdk
import time
import json
```



```

iot_client = greengrasssdk.client('iot-data')
send_topic = 'splunk/logs/put'

def create_request_with_all_fields():
    return {
        "request": {
            "event": "Access log test message."
        },
        "id" : "req_123"
    }

def publish_basic_message():
    messageToPublish = create_request_with_all_fields()
    print("Message To Publish: ", messageToPublish)
    iot_client.publish(topic=send_topic,
        payload=json.dumps(messageToPublish))

publish_basic_message()

def lambda_handler(event, context):
    return

```

Licenses

This connector is released under the [Greengrass Core Software License Agreement](#).

Changelog

The following table describes the changes in each version of the connector.

Version	Changes
4	Added the <code>IsolationMode</code> parameter to configure the containerization mode for the connector.
3	Upgraded the Lambda runtime to Python 3.7, which changes the runtime requirement.
2	Fix to reduce excessive logging.

Version	Changes
1	Initial release.

A Greengrass group can contain only one version of the connector at a time. For information about upgrading a connector version, see [the section called “Upgrading connector versions”](#).

See also

- [Integrate with services and protocols using connectors](#)
- [the section called “Get started with connectors \(console\)”](#)
- [the section called “Get started with connectors \(CLI\)”](#)

Twilio Notifications connector

Warning

This connector has moved into the *extended life phase*, and AWS IoT Greengrass won't release updates that provide features, enhancements to existing features, security patches, or bug fixes. For more information, see [AWS IoT Greengrass Version 1 maintenance policy](#).

The Twilio Notifications [connector](#) makes automated phone calls or sends text messages through Twilio. You can use this connector to send notifications in response to events in the Greengrass group. For phone calls, the connector can forward a voice message to the recipient.

This connector receives Twilio message information on an MQTT topic, and then triggers a Twilio notification.

Note

For a tutorial that shows how to use the Twilio Notifications connector, see [the section called “Get started with connectors \(console\)”](#) or [the section called “Get started with connectors \(CLI\)”](#).

This connector has the following versions.

Version	ARN
5	arn:aws:greengrass: <i>region</i> ::/connectors/TwilioNotifications/versions/5
4	arn:aws:greengrass: <i>region</i> ::/connectors/TwilioNotifications/versions/4
3	arn:aws:greengrass: <i>region</i> ::/connectors/TwilioNotifications/versions/3
2	arn:aws:greengrass: <i>region</i> ::/connectors/TwilioNotifications/versions/2
1	arn:aws:greengrass: <i>region</i> ::/connectors/TwilioNotifications/versions/1

For information about version changes, see the [Changelog](#).

Requirements

This connector has the following requirements:

Version 4 - 5

- AWS IoT Greengrass Core software v1.9.3 or later. AWS IoT Greengrass must be configured to support local secrets, as described in [Secrets Requirements](#).

Note

This requirement includes allowing access to your Secrets Manager secrets. If you're using the default Greengrass service role, Greengrass has permission to get the values of secrets with names that start with *greengrass-*.

- [Python](#) version 3.7 or 3.8 installed on the core device and added to the PATH environment variable.

 Note

To use Python 3.8, run the following command to create a symbolic link from the the default Python 3.7 installation folder to the installed Python 3.8 binaries.

```
sudo ln -s path-to-python-3.8/python3.8 /usr/bin/python3.7
```

This configures your device to meet the Python requirement for AWS IoT Greengrass.

- A Twilio account SID, auth token, and Twilio-enabled phone number. After you create a Twilio project, these values are available on the project dashboard.

 Note

You can use a Twilio trial account. If you're using a trial account, you must add non-Twilio recipient phone numbers to a list of verified phone numbers. For more information, see [How to Work with your Free Twilio Trial Account](#).

- A text type secret in AWS Secrets Manager that stores the Twilio auth token. For more information, see [Creating a basic secret](#) in the *AWS Secrets Manager User Guide*.

 Note

To create the secret in the Secrets Manager console, enter your token on the **Plaintext** tab. Don't include quotation marks or other formatting. In the API, specify the token as the value for the `SecretString` property.

- A secret resource in the Greengrass group that references the Secrets Manager secret. For more information, see [Deploy secrets to the core](#).

Versions 1 - 3

- AWS IoT Greengrass Core software v1.7 or later. AWS IoT Greengrass must be configured to support local secrets, as described in [Secrets Requirements](#).

Note

This requirement includes allowing access to your Secrets Manager secrets. If you're using the default Greengrass service role, Greengrass has permission to get the values of secrets with names that start with *greengrass-*.

- [Python](#) version 2.7 installed on the core device and added to the PATH environment variable.
- A Twilio account SID, auth token, and Twilio-enabled phone number. After you create a Twilio project, these values are available on the project dashboard.

Note

You can use a Twilio trial account. If you're using a trial account, you must add non-Twilio recipient phone numbers to a list of verified phone numbers. For more information, see [How to Work with your Free Twilio Trial Account](#).

- A text type secret in AWS Secrets Manager that stores the Twilio auth token. For more information, see [Creating a basic secret](#) in the *AWS Secrets Manager User Guide*.

Note

To create the secret in the Secrets Manager console, enter your token on the **Plaintext** tab. Don't include quotation marks or other formatting. In the API, specify the token as the value for the `SecretString` property.

- A secret resource in the Greengrass group that references the Secrets Manager secret. For more information, see [Deploy secrets to the core](#).

Connector Parameters

This connector provides the following parameters.

Version 5

`TWILIO_ACCOUNT_SID`

The Twilio account SID that's used to invoke the Twilio API.

Display name in the AWS IoT console: **Twilio account SID**

Required: true

Type: string

Valid pattern: .+

TwilioAuthTokenSecretArn

The ARN of the Secrets Manager secret that stores the Twilio auth token.

 Note

This is used to access the value of the local secret on the core.

Display name in the AWS IoT console: **ARN of Twilio auth token secret**

Required: true

Type: string

Valid pattern: `arn:aws:secretsmanager:[a-z0-9\-\+]:[0-9]{12}:secret:([a-zA-Z0-9\-\+\/]*[a-zA-Z0-9/_+=,.\@-\+]-[a-zA-Z0-9]+)`

TwilioAuthTokenSecretArn-ResourceId

The ID of the secret resource in the Greengrass group that references the secret for the Twilio auth token.

Display name in the AWS IoT console: **Twilio auth token resource**

Required: true

Type: string

Valid pattern: .+

DefaultFromPhoneNumber

The default Twilio-enabled phone number that Twilio uses to send messages. Twilio uses this number to initiate the text or call.

- If you don't configure a default phone number, you must specify a phone number in the `from_number` property in the input message body.
- If you do configure a default phone number, you can optionally override the default by specifying the `from_number` property in the input message body.

Display name in the AWS IoT console: **Default from phone number**

Required: false

Type: string

Valid pattern: `^\$|\+[0-9]+`

IsolationMode

The [containerization](#) mode for this connector. The default is `GreengrassContainer`, which means that the connector runs in an isolated runtime environment inside the AWS IoT Greengrass container.

 Note

The default containerization setting for the group does not apply to connectors.

Display name in the AWS IoT console: **Container isolation mode**

Required: false

Type: string

Valid values: `GreengrassContainer` or `NoContainer`

Valid pattern: `^NoContainer$|^GreengrassContainer$`

Version 1 - 4

TWILIO_ACCOUNT_SID

The Twilio account SID that's used to invoke the Twilio API.

Display name in the AWS IoT console: **Twilio account SID**

Required: true

Type: string

Valid pattern: .+

TwilioAuthTokenSecretArn

The ARN of the Secrets Manager secret that stores the Twilio auth token.

 Note

This is used to access the value of the local secret on the core.

Display name in the AWS IoT console: **ARN of Twilio auth token secret**

Required: true

Type: string

Valid pattern: `arn:aws:secretsmanager:[a-z0-9\-\_]+:[0-9]{12}:secret:([a-zA-Z0-9\_\-\_+\/)*[a-zA-Z0-9/_+=,.\@-\_]+-[a-zA-Z0-9]+`

TwilioAuthTokenSecretArn-ResourceId

The ID of the secret resource in the Greengrass group that references the secret for the Twilio auth token.

Display name in the AWS IoT console: **Twilio auth token resource**

Required: true

Type: string

Valid pattern: .+

DefaultFromPhoneNumber

The default Twilio-enabled phone number that Twilio uses to send messages. Twilio uses this number to initiate the text or call.

- If you don't configure a default phone number, you must specify a phone number in the `from_number` property in the input message body.
- If you do configure a default phone number, you can optionally override the default by specifying the `from_number` property in the input message body.

Display name in the AWS IoT console: **Default from phone number**

Required: false

Type: string

Valid pattern: `^$|\+[0-9]+`

Create Connector Example (AWS CLI)

The following example CLI command creates a `ConnectorDefinition` with an initial version that contains the Twilio Notifications connector.

```
aws greengrass create-connector-definition --name MyGreengrassConnectors --initial-  
version '{  
  "Connectors": [  
    {  
      "Id": "MyTwilioNotificationsConnector",  
      "ConnectorArn": "arn:aws:greengrass:region::/connectors/  
TwilioNotifications/versions/5",  
      "Parameters": {  
        "TWILIO_ACCOUNT_SID": "abcd12345xyz",  
        "TwilioAuthTokenSecretArn": "arn:aws:secretsmanager:region:account-  
id:secret:greengrass-secret-hash",  
        "TwilioAuthTokenSecretArn-ResourceId": "MyTwilioSecret",  
        "DefaultFromPhoneNumber": "+19999999999",  
        "IsolationMode" : "GreengrassContainer"  
      }  
    }  
  ]  
}'
```

For tutorials that show how add the Twilio Notifications connector to a group, see [the section called “Get started with connectors \(CLI\)”](#) and [the section called “Get started with connectors \(console\)”](#).

Input data

This connector accepts Twilio message information on two MQTT topics. Input messages must be in JSON format.

- Text message information on the `twilio/txt` topic.

- Phone message information on the `twilio/call` topic.

Note

The input message payload can include a text message (`message`) or voice message (`voice_message_location`), but not both.

Topic filter: `twilio/txt`**Message properties**

`request`

Information about the Twilio notification.

Required: `true`

Type: object that includes the following properties:

`recipient`

The message recipient. Only one recipient is supported.

Required: `true`

Type: object that include the following properties:

`name`

The name of the recipient.

Required: `true`

Type: `string`

Valid pattern: `.*`

`phone_number`

The phone number of the recipient.

Required: `true`

Type: string

Valid pattern: \+[1-9]+

message

The text content of the text message. Only text messages are supported on this topic. For voice messages, use `twilio/call`.

Required: true

Type: string

Valid pattern: .+

from_number

The phone number of the sender. Twilio uses this phone number to initiate the message. This property is required if the `DefaultFromPhoneNumber` parameter isn't configured. If `DefaultFromPhoneNumber` is configured, you can use this property to override the default.

Required: false

Type: string

Valid pattern: \+[1-9]+

retries

The number of retries. The default is 0.

Required: false

Type: integer

id

An arbitrary ID for the request. This property is used to map an input request to an output response.

Required: true

Type: string

Valid pattern: . +

Example input

```
{
  "request": {
    "recipient": {
      "name": "Darla",
      "phone_number": "+12345000000",
      "message": "Hello from the edge"
    },
    "from_number": "+19999999999",
    "retries": 3
  },
  "id": "request123"
}
```

Topic filter: `twilio/call`

Message properties

`request`

Information about the Twilio notification.

Required: true

Type: object that includes the following properties:

`recipient`

The message recipient. Only one recipient is supported.

Required: true

Type: object that include the following properties:

`name`

The name of the recipient.

Required: true

Type: string

Valid pattern: .+

phone_number

The phone number of the recipient.

Required: true

Type: string

Valid pattern: \+[1-9]+

voice_message_location

The URL of the audio content for the voice message. This must be in TwiML format. Only voice messages are supported on this topic. For text messages, use `twilio/txt`.

Required: true

Type: string

Valid pattern: .+

from_number

The phone number of the sender. Twilio uses this phone number to initiate the message. This property is required if the `DefaultFromPhoneNumber` parameter isn't configured. If `DefaultFromPhoneNumber` is configured, you can use this property to override the default.

Required: false

Type: string

Valid pattern: \+[1-9]+

retries

The number of retries. The default is 0.

Required: false

Type: integer

id

An arbitrary ID for the request. This property is used to map an input request to an output response.

Required: true

Type: string

Valid pattern: .+

Example input

```
{
  "request": {
    "recipient": {
      "name": "Darla",
      "phone_number": "+12345000000",
      "voice_message_location": "https://some-public-TwiML"
    },
    "from_number": "+19999999999",
    "retries": 3
  },
  "id": "request123"
}
```

Output data

This connector publishes status information as output data on an MQTT topic.

Topic filter in subscription

twilio/message/status

Example output: Success

```
{
  "response": {
    "status": "success",
    "payload": {
      "from_number": "+19999999999",
      "messages": {
```

```

        "message_status": "queued",
        "to_number": "+12345000000",
        "name": "Darla"
    }
},
"id": "request123"
}

```

Example output: Failure

```

{
  "response": {
    "status": "fail",
    "error_message": "Recipient name cannot be None",
    "error": "InvalidParameter",
    "payload": None
  },
  "id": "request123"
}

```

The `payload` property in the output is the response from the Twilio API when the message is sent. If the connector detects that the input data is invalid (for example, it doesn't specify a required input field), the connector returns an error and sets the value to `None`. The following are example payloads:

```

{
  'from_number': '+19999999999',
  'messages': {
    'name': 'Darla',
    'to_number': '+12345000000',
    'message_status': 'undelivered'
  }
}

```

```

{
  'from_number': '+19999999999',
  'messages': {
    'name': 'Darla',
    'to_number': '+12345000000',
    'message_status': 'queued'
  }
}

```

```
}  
}
```

Usage Example

Use the following high-level steps to set up an example Python 3.7 Lambda function that you can use to try out the connector.

Note

The [the section called “Get started with connectors \(console\)”](#) and [the section called “Get started with connectors \(CLI\)”](#) topics contain end-to-end steps that show how to set up, deploy, and test the Twilio Notifications connector.

1. Make sure you meet the [requirements](#) for the connector.
2. Create and publish a Lambda function that sends input data to the connector.

Save the [example code](#) as a PY file. Download and unzip the [AWS IoT Greengrass Core SDK for Python](#). Then, create a zip package that contains the PY file and the greengrasssdk folder at the root level. This zip package is the deployment package that you upload to AWS Lambda.

After you create the Python 3.7 Lambda function, publish a function version and create an alias.

3. Configure your Greengrass group.
 - a. Add the Lambda function by its alias (recommended). Configure the Lambda lifecycle as long-lived (or "Pinned": `true` in the CLI).
 - b. Add the required secret resource and grant read access to the Lambda function.
 - c. Add the connector and configure its [parameters](#).
 - d. Add subscriptions that allow the connector to receive [input data](#) and send [output data](#) on supported topic filters.
 - Set the Lambda function as the source, the connector as the target, and use a supported input topic filter.
 - Set the connector as the source, AWS IoT Core as the target, and use a supported output topic filter. You use this subscription to view status messages in the AWS IoT console.

4. Deploy the group.
5. In the AWS IoT console, on the **Test** page, subscribe to the output data topic to view status messages from the connector. The example Lambda function is long-lived and starts sending messages immediately after the group is deployed.

When you're finished testing, you can set the Lambda lifecycle to on-demand (or "Pinned": `false` in the CLI) and deploy the group. This stops the function from sending messages.

Example

The following example Lambda function sends an input message to the connector. This example triggers a text message.

```
import greengrasssdk
import json

iot_client = greengrasssdk.client('iot-data')
TXT_INPUT_TOPIC = 'twilio/txt'
CALL_INPUT_TOPIC = 'twilio/call'

def publish_basic_message():

    txt = {
        "request": {
            "recipient" : {
                "name": "Darla",
                "phone_number": "+12345000000",
                "message": 'Hello from the edge'
            },
            "from_number" : "+19999999999"
        },
        "id" : "request123"
    }

    print("Message To Publish: ", txt)

    client.publish(topic=TXT_INPUT_TOPIC,
                  payload=json.dumps(txt))

publish_basic_message()

def lambda_handler(event, context):
```

```
return
```

Licenses

The Twilio Notifications connector includes the following third-party software/licensing:

- [twilio-python](#)/MIT

This connector is released under the [Greengrass Core Software License Agreement](#).

Changelog

The following table describes the changes in each version of the connector.

Version	Changes
5	Added the <code>IsolationMode</code> parameter to configure the containerization mode for the connector.
4	Upgraded the Lambda runtime to Python 3.7, which changes the runtime requirement.
3	Fix to reduce excessive logging.
2	Minor bug fixes and improvements.
1	Initial release.

A Greengrass group can contain only one version of the connector at a time. For information about upgrading a connector version, see [the section called “Upgrading connector versions”](#).

See also

- [Integrate with services and protocols using connectors](#)
- [the section called “Get started with connectors \(console\)”](#)
- [the section called “Get started with connectors \(CLI\)”](#)
- [Twilio API Reference](#)


Getting started with Greengrass connectors (console)

This feature is available for AWS IoT Greengrass Core v1.7 and later.

This tutorial shows how to use the AWS Management Console to work with connectors.

Use connectors to accelerate your development life cycle. Connectors are prebuilt, reusable modules that can make it easier to interact with services, protocols, and resources. They can help you deploy business logic to Greengrass devices more quickly. For more information, see [Integrate with services and protocols using connectors](#).

In this tutorial, you configure and deploy the [Twilio Notifications](#) connector. The connector receives Twilio message information as input data, and then triggers a Twilio text message. The data flow is shown in following diagram.



After you configure the connector, you create a Lambda function and a subscription.

- The function evaluates simulated data from a temperature sensor. It conditionally publishes the Twilio message information to an MQTT topic. This is the topic that the connector subscribes to.
- The subscription allows the function to publish to the topic and the connector to receive data from the topic.

The Twilio Notifications connector requires a Twilio auth token to interact with the Twilio API. The token is a text type secret created in AWS Secrets Manager and referenced from a group resource. This enables AWS IoT Greengrass to create a local copy of the secret on the Greengrass core, where

it is encrypted and made available to the connector. For more information, see [Deploy secrets to the core](#).

The tutorial contains the following high-level steps:

1. [Create a Secrets Manager secret](#)
2. [Add a secret resource to a group](#)
3. [Add a connector to the group](#)
4. [Create a Lambda function deployment package](#)
5. [Create a Lambda function](#)
6. [Add a function to the group](#)
7. [Add subscriptions to the group](#)
8. [Deploy the group](#)
9. [the section called "Test the solution"](#)

The tutorial should take about 20 minutes to complete.

Prerequisites

To complete this tutorial, you need:

- A Greengrass group and a Greengrass core (v1.9.3 or later). To learn how to create a Greengrass group and core, see [Getting started with AWS IoT Greengrass](#). The Getting Started tutorial also includes steps for installing the AWS IoT Greengrass Core software.
- Python 3.7 installed on the AWS IoT Greengrass core device.
- AWS IoT Greengrass must be configured to support local secrets, as described in [Secrets Requirements](#).

Note

This requirement includes allowing access to your Secrets Manager secrets. If you're using the default Greengrass service role, Greengrass has permission to get the values of secrets with names that start with *greengrass-*.

- A Twilio account SID, auth token, and Twilio-enabled phone number. After you create a Twilio project, these values are available on the project dashboard.

Note

You can use a Twilio trial account. If you're using a trial account, you must add non-Twilio recipient phone numbers to a list of verified phone numbers. For more information, see [How to Work with your Free Twilio Trial Account](#).

Step 1: Create a Secrets Manager secret

In this step, you use the AWS Secrets Manager console to create a text type secret for your Twilio auth token.

1. Sign in to the [AWS Secrets Manager console](#).

Note

For more information about this process, see [Step 1: Create and store your secret in AWS Secrets Manager](#) in the *AWS Secrets Manager User Guide*.

2. Choose **Store a new secret**.
3. Under **Choose secret type**, choose **Other type of secret**.
4. Under **Specify the key/value pairs to be stored for this secret**, on the **Plaintext** tab, enter your Twilio auth token. Remove all of the JSON formatting and enter only the token value.
5. Keep **aws/secretsmanager** selected for the encryption key, and then choose **Next**.

Note

You aren't charged by AWS KMS if you use the default AWS managed key that Secrets Manager creates in your account.

6. For **Secret name**, enter **greengrass-TwilioAuthToken**, and then choose **Next**.

Note

By default, the Greengrass service role allows AWS IoT Greengrass to get the value of secrets with names that start with *greengrass-*. For more information, see [secrets requirements](#).

7. This tutorial doesn't require rotation, so choose **disable automatic rotation**, and then choose **Next**.
8. On the **Review** page, review your settings, and then choose **Store**.

Next, you create a secret resource in your Greengrass group that references the secret.

Step 2: Add a secret resource to a Greengrass group

In this step, you add a *secret resource* to the Greengrass group. This resource is a reference to the secret that you created in the previous step.

1. In the AWS IoT console navigation pane, under **Manage**, expand **Greengrass devices**, and then choose **Groups (V1)**.
2. Choose the group that you want to add the secret resource to.
3. On the group configuration page, choose the **Resources** tab, and then scroll down to the **Secrets** section. The **Secrets** section displays the secret resources that belong to the group. You can add, edit, and remove secret resources from this section.

Note

Alternatively, the console allows you to create a secret and secret resource when you configure a connector or Lambda function. You can do this from the connector's **Configure parameters** page or the Lambda function's **Resources** page.

4. Choose **Add** under the **Secrets** section.
5. On the **Add a secret resource** page, enter **MyTwilioAuthToken** for the **Resource name**.
6. For the **Secret**, choose **greengrass-TwilioAuthToken**.
7. In the **Select labels (Optional)** section, the **AWSCURRENT** staging label represents the latest version of the secret. This label is always included in a secret resource.

Note

This tutorial requires the `AWSCURRENT` label only. You can optionally include labels that are required by your Lambda function or connector.

8. Choose **Add resource**.

Step 3: Add a connector to the Greengrass group

In this step, you configure parameters for the [Twilio Notifications connector](#) and add it to the group.

1. On the group configuration page, choose **Connectors**, and then choose **Add a connector**.
2. On the **Add connector** page, choose **Twilio Notifications**.
3. Choose the version.
4. In the **Configuration** section:
 - For **Twilio auth token resource**, enter the resource that you created in the previous step.

Note

When you enter the resource, the **ARN of Twilio auth token secret** property is populated for you.

- For **Default from phone number**, enter your Twilio-enabled phone number.
 - For **Twilio account SID**, enter your Twilio account SID.
5. Choose **Add resource**.

Step 4: Create a Lambda function deployment package

To create a Lambda function, you must first create a Lambda function *deployment package* that contains the function code and dependencies. Greengrass Lambda functions require the [AWS IoT Greengrass Core SDK](#) for tasks such as communicating with MQTT messages in the core environment and accessing local secrets. This tutorial creates a Python function, so you use the Python version of the SDK in the deployment package.

1. From the [AWS IoT Greengrass Core SDK](#) downloads page, download the AWS IoT Greengrass Core SDK for Python to your computer.
2. Unzip the downloaded package to get the SDK. The SDK is the greengrasssdk folder.
3. Save the following Python code function in a local file named `temp_monitor.py`.

```
import greengrasssdk
import json
import random

client = greengrasssdk.client('iot-data')

# publish to the Twilio Notifications connector through the twilio/txt topic
def function_handler(event, context):
    temp = event['temperature']

    # check the temperature
    # if greater than 30C, send a notification
    if temp > 30:
        data = build_request(event)
        client.publish(topic='twilio/txt', payload=json.dumps(data))
        print('published:' + str(data))

    print('temperature:' + str(temp))
    return

# build the Twilio request from the input data
def build_request(event):
    to_name = event['to_name']
    to_number = event['to_number']
    temp_report = 'temperature:' + str(event['temperature'])

    return {
        "request": {
            "recipient": {
                "name": to_name,
                "phone_number": to_number,
                "message": temp_report
            }
        },
        "id": "request_" + str(random.randint(1,101))
    }
```


4. Zip the following items into a file named `temp_monitor_python.zip`. When creating the ZIP file, include only the code and dependencies, not the containing folder.
 - **temp_monitor.py**. App logic.
 - **greengrasssdk**. Required library for Python Greengrass Lambda functions that publish MQTT messages.


This is your Lambda function deployment package.

Now, create a Lambda function that uses the deployment package.

Step 5: Create a Lambda function in the AWS Lambda console

In this step, you use the AWS Lambda console to create a Lambda function and configure it to use your deployment package. Then, you publish a function version and create an alias.

1. First, create the Lambda function.
 - a. In the AWS Management Console, choose **Services**, and open the AWS Lambda console.
 - b. Choose **Create function** and then choose **Author from scratch**.
 - c. In the **Basic information** section, use the following values:
 - For **Function name**, enter **TempMonitor**.
 - For **Runtime**, choose **Python 3.7**.
 - For **Permissions**, keep the default setting. This creates an execution role that grants basic Lambda permissions. This role isn't used by AWS IoT Greengrass.
 - d. At the bottom of the page, choose **Create function**.
2. Next, register the handler and upload your Lambda function deployment package.
 - a. On the **Code** tab, under **Code source**, choose **Upload from**. From the dropdown, choose **.zip file**.



- b. Choose **Upload**, and then choose your `temp_monitor_python.zip` deployment package. Then, choose **Save**.
- c. On the **Code** tab for the function, under **Runtime settings**, choose **Edit**, and then enter the following values.
 - For **Runtime**, choose **Python 3.7**.
 - For **Handler**, enter `temp_monitor.function_handler`
- d. Choose **Save**.

Note

The **Test** button on the AWS Lambda console doesn't work with this function. The AWS IoT Greengrass Core SDK doesn't contain modules that are required to run your Greengrass Lambda functions independently in the AWS Lambda console. These modules (for example, `greengrass_common`) are supplied to the functions after they are deployed to your Greengrass core.

3. Now, publish the first version of your Lambda function and create an [alias for the version](#).

Note

Greengrass groups can reference a Lambda function by alias (recommended) or by version. Using an alias makes it easier to manage code updates because you don't have to change your subscription table or group definition when the function code is updated. Instead, you just point the alias to the new function version.

- a. From the **Actions** menu, choose **Publish new version**.
- b. For **Version description**, enter **First version**, and then choose **Publish**.
- c. On the **TempMonitor: 1** configuration page, from the **Actions** menu, choose **Create alias**.

- d. On the **Create a new alias** page, use the following values:
 - For **Name**, enter **GG_TempMonitor**.
 - For **Version**, choose **1**.

 Note

AWS IoT Greengrass doesn't support Lambda aliases for **\$LATEST** versions.

- e. Choose **Create**.

Now you're ready to add the Lambda function to your Greengrass group.

Step 6: Add a Lambda function to the Greengrass group

In this step, you add the Lambda function to the group and then configure its lifecycle and environment variables. For more information, see [the section called "Controlling Greengrass Lambda function execution"](#).

1. On the group configuration page, choose the **Lambda functions** tab.
2. Under **My Lambda functions**, choose **Add**.
3. On the **Add Lambda function** page, choose **TempMonitor** for your Lambda function.
4. For **Lambda function version**, choose **Alias: GG_TempMonitor**.
5. Choose **Add Lambda function**.

Step 7: Add subscriptions to the Greengrass group

In this step, you add a subscription that enables the Lambda function to send input data to the connector. The connector defines the MQTT topics that it subscribes to, so this subscription uses one of the topics. This is the same topic that the example function publishes to.

For this tutorial, you also create subscriptions that allow the function to receive simulated temperature readings from AWS IoT and allow AWS IoT to receive status information from the connector.

1. On the group configuration page, choose the **Subscriptions** tab, and then choose **Add Subscription**.

2. On the **Create a subscription** page, configure the source and target, as follows:
 - a. For **Source type**, choose **Lambda function**, and then choose **TempMonitor**.
 - b. For **Target type**, choose **Connector**, and then choose **Twilio Notifications**.
3. For the **Topic filter**, choose **twilio/txt**.
4. Choose **Create subscription**.
5. Repeat steps 1 - 4 to create a subscription that allows AWS IoT to publish messages to the function.
 - a. For **Source type**, choose **Service**, and then choose **IoT Cloud**.
 - b. For **Select a target**, choose **Lambda function**, and then choose **TempMonitor**.
 - c. For **Topic filter**, enter **temperature/input**.
6. Repeat steps 1 - 4 to create a subscription that allows the connector to publish messages to AWS IoT.
 - a. For **Source type**, choose **Connector**, and then choose **Twilio Notifications**.
 - b. For **Target type**, choose **Service**, and then choose **IoT Cloud**.
 - c. For **Topic filter**, **twilio/message/status** is entered for you. This is the predefined topic that the connector publishes to.

Step 8: Deploy the Greengrass group

Deploy the group to the core device.

1. Make sure that the AWS IoT Greengrass core is running. Run the following commands in your Raspberry Pi terminal, as needed:
 - a. To check whether the daemon is running:

```
ps aux | grep -E 'greengrass.*daemon'
```

If the output contains a root entry for `/greengrass/ggc/packages/ggc-version/bin/daemon`, then the daemon is running.

Note

The version in the path depends on the AWS IoT Greengrass Core software version that's installed on your core device.

- b. To start the daemon:

```
cd /greengrass/ggc/core/  
sudo ./greengrassd start
```

2. On the group configuration page, choose **Deploy**.
3.
 - a. In the **Lambda functions** tab, under the **System Lambda functions** section, select **IP detector** and choose **Edit**.
 - b. In the **Edit IP detector settings** dialog box, select **Automatically detect and override MQTT broker endpoints**.
 - c. Choose **Save**.

This enables devices to automatically acquire connectivity information for the core, such as IP address, DNS, and port number. Automatic detection is recommended, but AWS IoT Greengrass also supports manually specified endpoints. You're only prompted for the discovery method the first time that the group is deployed.

Note

If prompted, grant permission to create the [Greengrass service role](#) and associate it with your AWS account in the current AWS Region. This role allows AWS IoT Greengrass to access your resources in AWS services.

The **Deployments** page shows the deployment timestamp, version ID, and status. When completed, the status displayed for the deployment should be **Completed**.

For troubleshooting help, see [Troubleshooting](#).

Note

A Greengrass group can contain only one version of the connector at a time. For information about upgrading a connector version, see [the section called “Upgrading connector versions”](#).

Test the solution

1. On the AWS IoT console home page, choose **Test**.
2. For **Subscribe to topic**, use the following values, and then choose **Subscribe**. The Twilio Notifications connector publishes status information to this topic.

Property	Value
Subscription topic	twilio/message/status
MQTT payload display	Display payloads as strings

3. For **Publish to topic**, use the following values, and then choose **Publish** to invoke the function.

Property	Value
Topic	temperature/input
Message	Replace <i>recipient-name</i> with a name and <i>recipient-phone-number</i> with the phone number of the text message recipient. Example: +12345000000 <pre>{ "to_name": " <i>recipient-name</i> ", "to_number": " <i>recipient-phone-number</i> ", "temperature": 31 }</pre>

Property	Value
	If you're using a trial account, you must add non-Twilio recipient phone numbers to a list of verified phone numbers. For more information, see Verify your Personal Phone Number .

If successful, the recipient receives the text message and the console displays the success status from the [output data](#).

Now, change the temperature in the input message to **29** and publish. Because this is less than 30, the TempMonitor function doesn't trigger a Twilio message.

See also

- [Integrate with services and protocols using connectors](#)
- [the section called "AWS-provided Greengrass connectors"](#)


Getting started with Greengrass connectors (CLI)

This feature is available for AWS IoT Greengrass Core v1.7 and later.

This tutorial shows how to use the AWS CLI to work with connectors.

Use connectors to accelerate your development life cycle. Connectors are prebuilt, reusable modules that can make it easier to interact with services, protocols, and resources. They can help you deploy business logic to Greengrass devices more quickly. For more information, see [Integrate with services and protocols using connectors](#).

In this tutorial, you configure and deploy the [Twilio Notifications](#) connector. The connector receives Twilio message information as input data, and then triggers a Twilio text message. The data flow is shown in following diagram.



After you configure the connector, you create a Lambda function and a subscription.

- The function evaluates simulated data from a temperature sensor. It conditionally publishes the Twilio message information to an MQTT topic. This is the topic that the connector subscribes to.
- The subscription allows the function to publish to the topic and the connector to receive data from the topic.

The Twilio Notifications connector requires a Twilio auth token to interact with the Twilio API. The token is a text type secret created in AWS Secrets Manager and referenced from a group resource. This enables AWS IoT Greengrass to create a local copy of the secret on the Greengrass core, where it is encrypted and made available to the connector. For more information, see [Deploy secrets to the core](#).

The tutorial contains the following high-level steps:

1. [Create a Secrets Manager secret](#)
2. [Create a resource definition and version](#)
3. [Create a connector definition and version](#)
4. [Create a Lambda function deployment package](#)
5. [Create a Lambda function](#)
6. [Create a function definition and version](#)
7. [Create a subscription definition and version](#)
8. [Create a group version](#)

9. [Create a deployment](#)

10. [the section called "Test the solution"](#)

The tutorial should take about 30 minutes to complete.

Using the AWS IoT Greengrass API

It's helpful to understand the following patterns when you work with Greengrass groups and group components (for example, the connectors, functions, and resources in the group).

- At the top of the hierarchy, a component has a *definition* object that is a container for *version* objects. In turn, a version is a container for the connectors, functions, or other component types.
- When you deploy to the Greengrass core, you deploy a specific group version. A group version can contain one version of each type of component. A core is required, but the others are included as needed.
- Versions are immutable, so you must create new versions when you want to make changes.

Tip

If you receive an error when you run an AWS CLI command, add the `--debug` parameter and then rerun the command to get more information about the error.

The AWS IoT Greengrass API lets you create multiple definitions for a component type. For example, you can create a `FunctionDefinition` object every time that you create a `FunctionDefinitionVersion`, or you can add new versions to an existing definition. This flexibility allows you to customize your version management system.

Prerequisites

To complete this tutorial, you need:

- A Greengrass group and a Greengrass core (v1.9.3 or later). To learn how to create a Greengrass group and core, see [Getting started with AWS IoT Greengrass](#). The Getting Started tutorial also includes steps for installing the AWS IoT Greengrass Core software.
- Python 3.7 installed on the AWS IoT Greengrass core device.

- AWS IoT Greengrass must be configured to support local secrets, as described in [Secrets Requirements](#).

Note

This requirement includes allowing access to your Secrets Manager secrets. If you're using the default Greengrass service role, Greengrass has permission to get the values of secrets with names that start with *greengrass-*.

- A Twilio account SID, auth token, and Twilio-enabled phone number. After you create a Twilio project, these values are available on the project dashboard.

Note

You can use a Twilio trial account. If you're using a trial account, you must add non-Twilio recipient phone numbers to a list of verified phone numbers. For more information, see [How to Work with your Free Twilio Trial Account](#).

- AWS CLI installed and configured on your computer. For more information, see [Installing the AWS Command Line Interface](#) and [Configuring the AWS CLI](#) in the *AWS Command Line Interface User Guide*.

The examples in this tutorial are written for Linux and other Unix-based systems. If you're using Windows, see [Specifying parameter values for the AWS Command Line Interface](#) to learn about differences in syntax.

If the command contains a JSON string, the tutorial provides an example that has the JSON on a single line. On some systems, it might be easier to edit and run commands using this format.

Step 1: Create a Secrets Manager secret

In this step, you use the AWS Secrets Manager API to create a secret for your Twilio auth token.

1. First, create the secret.
 - Replace *twilio-auth-token* with your Twilio auth token.

```
aws secretsmanager create-secret --name greengrass-TwilioAuthToken --secret-string twilio-auth-token
```

Note

By default, the Greengrass service role allows AWS IoT Greengrass to get the value of secrets with names that start with *greengrass-*. For more information, see [secrets requirements](#).

2. Copy the ARN of the secret from the output. You use this to create the secret resource and to configure the Twilio Notifications connector.

Step 2: Create a resource definition and version

In this step, you use the AWS IoT Greengrass API to create a secret resource for your Secrets Manager secret.

1. Create a resource definition that includes an initial version.
 - Replace *secret-arn* with the ARN of the secret that you copied in the previous step.

JSON Expanded

```
aws greengrass create-resource-definition --name MyGreengrassResources --initial-version '{
  "Resources": [
    {
      "Id": "TwilioAuthToken",
      "Name": "MyTwilioAuthToken",
      "ResourceDataContainer": {
        "SecretsManagerSecretResourceData": {
          "ARN": "secret-arn"
        }
      }
    }
  ]
}
```

```
}'
```

JSON Single-line

```
aws greengrass create-resource-definition \
--name MyGreengrassResources \
--initial-version '{"Resources": [{"Id": "TwilioAuthToken",
  "Name": "MyTwilioAuthToken", "ResourceDataContainer":
  {"SecretsManagerSecretResourceData": {"ARN": "secret-arn"}}}]}'
```

2. Copy the LatestVersionArn of the resource definition from the output. You use this value to add the resource definition version to the group version that you deploy to the core.

Step 3: Create a connector definition and version

In this step, you configure parameters for the Twilio Notifications connector.

1. Create a connector definition with an initial version.
 - Replace *account-sid* with your Twilio account SID.
 - Replace *secret-arn* with the ARN of your Secrets Manager secret. The connector uses this to get the value of the local secret.
 - Replace *phone-number* with your Twilio-enabled phone number. Twilio uses this to initiate the text message. This can be overridden in the input message payload. Use the following format: +19999999999.

JSON Expanded

```
aws greengrass create-connector-definition --name MyGreengrassConnectors --
initial-version '{
  "Connectors": [
    {
      "Id": "MyTwilioNotificationsConnector",
      "ConnectorArn": "arn:aws:greengrass:region::/connectors/
TwilioNotifications/versions/4",
      "Parameters": {
        "TWILIO_ACCOUNT_SID": "account-sid",
```

```

        "TwilioAuthTokenSecretArn": "secret-arn",
        "TwilioAuthTokenSecretArn-ResourceId": "TwilioAuthToken",
        "DefaultFromPhoneNumber": "phone-number"
    }
}
]
}'

```

JSON Single-line

```

aws greengrass create-connector-definition \
--name MyGreengrassConnectors \
--initial-version '{"Connectors": [{"Id": "MyTwilioNotificationsConnector",
  "ConnectorArn": "arn:aws:greengrass:region::/connectors/TwilioNotifications/
versions/4", "Parameters": {"TWILIO_ACCOUNT_SID": "account-sid",
  "TwilioAuthTokenSecretArn": "secret-arn", "TwilioAuthTokenSecretArn-
ResourceId": "TwilioAuthToken", "DefaultFromPhoneNumber": "phone-number"}}]}'

```

Note

TwilioAuthToken is the ID that you used in the previous step to create the secret resource.

2. Copy the LatestVersionArn of the connector definition from the output. You use this value to add the connector definition version to the group version that you deploy to the core.

Step 4: Create a Lambda function deployment package

To create a Lambda function, you must first create a Lambda function *deployment package* that contains the function code and dependencies. Greengrass Lambda functions require the [AWS IoT Greengrass Core SDK](#) for tasks such as communicating with MQTT messages in the core environment and accessing local secrets. This tutorial creates a Python function, so you use the Python version of the SDK in the deployment package.

1. From the [AWS IoT Greengrass Core SDK](#) downloads page, download the AWS IoT Greengrass Core SDK for Python to your computer.
2. Unzip the downloaded package to get the SDK. The SDK is the greengrasssdk folder.
3. Save the following Python code function in a local file named `temp_monitor.py`.

```
import greengrasssdk
import json
import random

client = greengrasssdk.client('iot-data')

# publish to the Twilio Notifications connector through the twilio/txt topic
def function_handler(event, context):
    temp = event['temperature']

    # check the temperature
    # if greater than 30C, send a notification
    if temp > 30:
        data = build_request(event)
        client.publish(topic='twilio/txt', payload=json.dumps(data))
        print('published:' + str(data))

    print('temperature:' + str(temp))
    return

# build the Twilio request from the input data
def build_request(event):
    to_name = event['to_name']
    to_number = event['to_number']
    temp_report = 'temperature:' + str(event['temperature'])

    return {
        "request": {
            "recipient": {
                "name": to_name,
                "phone_number": to_number,
                "message": temp_report
            }
        },
        "id": "request_" + str(random.randint(1,101))
    }
```

4. Zip the following items into a file named `temp_monitor_python.zip`. When creating the ZIP file, include only the code and dependencies, not the containing folder.
 - **temp_monitor.py**. App logic.

- **greengrasssdk**. Required library for Python Greengrass Lambda functions that publish MQTT messages.

This is your Lambda function deployment package.

Step 5: Create a Lambda function

Now, create a Lambda function that uses the deployment package.

1. Create an IAM role so you can pass in the role ARN when you create the function.

JSON Expanded

```
aws iam create-role --role-name Lambda_empty --assume-role-policy '{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Principal": {
        "Service": "lambda.amazonaws.com"
      },
      "Action": "sts:AssumeRole"
    }
  ]
}'
```

JSON Single-line

```
aws iam create-role --role-name Lambda_empty --assume-role-policy '{"Version":
"2012-10-17", "Statement": [{"Effect": "Allow", "Principal": {"Service":
"lambda.amazonaws.com"},"Action": "sts:AssumeRole"}]}'
```

Note

AWS IoT Greengrass doesn't use this role because permissions for your Greengrass Lambda functions are specified in the Greengrass group role. For this tutorial, you create an empty role.

2. Copy the Arn from the output.
3. Use the AWS Lambda API to create the TempMonitor function. The following command assumes that the zip file is in the current directory.
 - Replace *role-arn* with the Arn that you copied.

```
aws lambda create-function \  
--function-name TempMonitor \  
--zip-file fileb://temp_monitor_python.zip \  
--role role-arn \  
--handler temp_monitor.function_handler \  
--runtime python3.7
```

4. Publish a version of the function.

```
aws lambda publish-version --function-name TempMonitor --description 'First  
version'
```

5. Create an alias for the published version.

Greengrass groups can reference a Lambda function by alias (recommended) or by version. Using an alias makes it easier to manage code updates because you don't have to change your subscription table or group definition when the function code is updated. Instead, you just point the alias to the new function version.

 Note

AWS IoT Greengrass doesn't support Lambda aliases for **\$LATEST** versions.

```
aws lambda create-alias --function-name TempMonitor --name GG_TempMonitor --  
function-version 1
```

6. Copy the AliasArn from the output. You use this value when you configure the function for AWS IoT Greengrass and when you create a subscription.

Now you're ready to configure the function for AWS IoT Greengrass.

Step 6: Create a function definition and version

To use a Lambda function on an AWS IoT Greengrass core, you create a function definition version that references the Lambda function by alias and defines the group-level configuration. For more information, see [the section called “Controlling Greengrass Lambda function execution”](#).

1. Create a function definition that includes an initial version.
 - Replace *alias-arn* with the AliasArn that you copied when you created the alias.

JSON Expanded

```
aws greengrass create-function-definition --name MyGreengrassFunctions --
initial-version '{
  "Functions": [
    {
      "Id": "TempMonitorFunction",
      "FunctionArn": "alias-arn",
      "FunctionConfiguration": {
        "Executable": "temp_monitor.function_handler",
        "MemorySize": 16000,
        "Timeout": 5
      }
    }
  ]
}'
```

JSON Single-line

```
aws greengrass create-function-definition \
--name MyGreengrassFunctions \
--initial-version '{"Functions": [{"Id": "TempMonitorFunction",
"FunctionArn": "alias-arn", "FunctionConfiguration": {"Executable":
"temp_monitor.function_handler", "MemorySize": 16000,"Timeout": 5}}]}'
```

2. Copy the LatestVersionArn from the output. You use this value to add the function definition version to the group version that you deploy to the core.
3. Copy the Id from the output. You use this value later when you update the function.

Step 7: Create a subscription definition and version

In this step, you add a subscription that enables the Lambda function to send input data to the connector. The connector defines the MQTT topics that it subscribes to, so this subscription uses one of the topics. This is the same topic that the example function publishes to.

For this tutorial, you also create subscriptions that allow the function to receive simulated temperature readings from AWS IoT and allow AWS IoT to receive status information from the connector.

1. Create a subscription definition that contains an initial version that includes the subscriptions.
 - Replace *alias-arn* with the AliasArn that you copied when you created the alias for the function. Use this ARN for both subscriptions that use it.

JSON Expanded

```
aws greengrass create-subscription-definition --initial-version '{
  "Subscriptions": [
    {
      "Id": "TriggerNotification",
      "Source": "alias-arn",
      "Subject": "twilio/txt",
      "Target": "arn:aws:greengrass:region::/connectors/
TwilioNotifications/versions/4"
    },
    {
      "Id": "TemperatureInput",
      "Source": "cloud",
      "Subject": "temperature/input",
      "Target": "alias-arn"
    },
    {
      "Id": "OutputStatus",
      "Source": "arn:aws:greengrass:region::/connectors/
TwilioNotifications/versions/4",
      "Subject": "twilio/message/status",
      "Target": "cloud"
    }
  ]
}
```

```
}'
```

JSON Single-line

```
aws greengrass create-subscription-definition \  
--initial-version '{"Subscriptions": [{"Id": "TriggerNotification", "Source":  
  "alias-arn", "Subject": "twilio/txt", "Target": "arn:aws:greengrass:region:/  
connectors/TwilioNotifications/versions/4"}, {"Id": "TemperatureInput",  
  "Source": "cloud", "Subject": "temperature/input", "Target": "alias-arn"},  
{"Id": "OutputStatus", "Source": "arn:aws:greengrass:region:/connectors/  
TwilioNotifications/versions/4", "Subject": "twilio/message/status", "Target":  
  "cloud"}]}'
```

2. Copy the LatestVersionArn from the output. You use this value to add the subscription definition version to the group version that you deploy to the core.

Step 8: Create a group version

Now, you're ready to create a group version that contains all of the items that you want to deploy. You do this by creating a group version that references the target version of each component type.

First, get the group ID and the ARN of the core definition version. These values are required to create the group version.

1. Get the ID of the group and latest group version:
 - a. Get the IDs of the target Greengrass group and group version. This procedure assumes that this is the latest group and group version. The following query returns the most recently created group.

```
aws greengrass list-groups --query "reverse(sort_by(Groups,  
  &CreationTimestamp))[0]"
```

Or, you can query by name. Group names are not required to be unique, so multiple groups might be returned.

```
aws greengrass list-groups --query "Groups[?Name=='MyGroup']"
```

 Note

You can also find these values in the AWS IoT console. The group ID is displayed on the group's **Settings** page. Group version IDs are displayed on the group's **Deployments** tab.

- b. Copy the Id of the target group from the output. You use this to get the core definition version and when you deploy the group.
 - c. Copy the LatestVersion from the output, which is the ID of the last version added to the group. You use this to get the core definition version.
2. Get the ARN of the core definition version:
 - a. Get the group version. For this step, we assume that the latest group version includes a core definition version.
 - Replace *group-id* with the Id that you copied for the group.
 - Replace *group-version-id* with the LatestVersion that you copied for the group.

```
aws greengrass get-group-version \  
--group-id group-id \  
--group-version-id group-version-id
```

- b. Copy the CoreDefinitionVersionArn from the output.
3. Create a group version.
 - Replace *group-id* with the Id that you copied for the group.
 - Replace *core-definition-version-arn* with the CoreDefinitionVersionArn that you copied for the core definition version.
 - Replace *resource-definition-version-arn* with the LatestVersionArn that you copied for the resource definition.
 - Replace *connector-definition-version-arn* with the LatestVersionArn that you copied for the connector definition.
 - Replace *function-definition-version-arn* with the LatestVersionArn that you copied for the function definition.

- Replace *subscription-definition-version-arn* with the LatestVersionArn that you copied for the subscription definition.

```
aws greengrass create-group-version \  
--group-id group-id \  
--core-definition-version-arn core-definition-version-arn \  
--resource-definition-version-arn resource-definition-version-arn \  
--connector-definition-version-arn connector-definition-version-arn \  
--function-definition-version-arn function-definition-version-arn \  
--subscription-definition-version-arn subscription-definition-version-arn
```

4. Copy the value of Version from the output. This is the ID of the group version. You use this value to deploy the group version.

Step 9: Create a deployment

Deploy the group to the core device.

1. In a core device terminal, make sure that the AWS IoT Greengrass daemon is running.
 - a. To check whether the daemon is running:

```
ps aux | grep -E 'greengrass.*daemon'
```

If the output contains a root entry for `/greengrass/ggc/packages/1.11.6/bin/daemon`, then the daemon is running.

- b. To start the daemon:

```
cd /greengrass/ggc/core/  
sudo ./greengrassd start
```

2. Create a deployment.
 - Replace *group-id* with the Id that you copied for the group.
 - Replace *group-version-id* with the Version that you copied for the new group version.

```
aws greengrass create-deployment \  
--deployment-type NewDeployment \  

```

```
--group-id group-id \  
--group-version-id group-version-id
```

3. Copy the DeploymentId from the output.
4. Get the deployment status.
 - Replace *group-id* with the Id that you copied for the group.
 - Replace *deployment-id* with the DeploymentId that you copied for the deployment.

```
aws greengrass get-deployment-status \  
--group-id group-id \  
--deployment-id deployment-id
```

If the status is Success, the deployment was successful. For troubleshooting help, see [Troubleshooting](#).

Test the solution

1. On the AWS IoT console home page, choose **Test**.
2. For **Subscribe to topic**, use the following values, and then choose **Subscribe**. The Twilio Notifications connector publishes status information to this topic.

Property	Value
Subscription topic	twilio/message/status
MQTT payload display	Display payloads as strings

3. For **Publish to topic**, use the following values, and then choose **Publish** to invoke the function.

Property	Value
Topic	temperature/input
Message	Replace <i>recipient-name</i> with a name and <i>recipient-phone-number</i> with

Property	Value
	the phone number of the text message recipient. Example: +12345000000 <pre>{ "to_name": " <i>recipient-name</i> ", "to_number": " <i>recipient-phone-number</i> ", "temperature": 31 }</pre> If you're using a trial account, you must add non-Twilio recipient phone numbers to a list of verified phone numbers. For more information, see Verify your Personal Phone Number.

If successful, the recipient receives the text message and the console displays the success status from the [output data](#).

Now, change the temperature in the input message to **29** and publish. Because this is less than 30, the TempMonitor function doesn't trigger a Twilio message.

See also

- [Integrate with services and protocols using connectors](#)
- [the section called "AWS-provided Greengrass connectors"](#)
- [the section called "Get started with connectors \(console\)"](#)
- [AWS Secrets Manager commands](#) in the *AWS CLI Command Reference*
- [AWS Identity and Access Management \(IAM\) commands](#) in the *AWS CLI Command Reference*
- [AWS Lambda commands](#) in the *AWS CLI Command Reference*
- [AWS IoT Greengrass commands](#) in the *AWS CLI Command Reference*

Greengrass Discovery RESTful API

All client devices that communicate with an AWS IoT Greengrass core must be a member of a Greengrass group. Each group must have a Greengrass core. The Discovery API enables devices to retrieve information required to connect to a Greengrass core that is in the same Greengrass group as the client device. When a client device first comes online, it can connect to the AWS IoT Greengrass service and use the Discovery API to find:

- The group to which it belongs. A client device can be a member of up to 10 groups.
- The IP address and port for the Greengrass core in the group.
- The group CA certificate, which can be used to authenticate the Greengrass core device.

Note

Client devices can also use the AWS IoT Device SDKs to discover connectivity information for a Greengrass core. For more information, see [AWS IoT Device SDK](#).

To use this API, send HTTP requests to the Discovery API endpoint. For example:

```
https://greengrass-ats.iot.region.amazonaws.com:port/greengrass/discover/thing/thing-name
```

For a list of supported Amazon Web Services Regions and endpoints for the AWS IoT Greengrass Discovery API, see [AWS IoT Greengrass endpoints and quotas](#) in the *AWS General Reference*. This is a data plane only API. The endpoints for group management and AWS IoT Core operations are different from the Discovery API endpoints.

Request

The request contains the standard HTTP headers and is sent to the Greengrass Discovery endpoint, as shown in the following examples.

The port number depends on whether the core is configured to send HTTPS traffic over port 8443 or port 443. For more information, see [the section called "Connect on port 443 or through a network proxy"](#).

Port 8443

```
HTTP GET https://greengrass-ats.iot.region.amazonaws.com:8443/greengrass/discover/thing/thing-name
```

Port 443

```
HTTP GET https://greengrass-ats.iot.region.amazonaws.com:443/greengrass/discover/thing/thing-name
```

Clients that connect on port 443 must implement the [Application Layer Protocol Negotiation \(ALPN\)](#) TLS extension and pass `x-amzn-http-ca` as the `ProtocolName` in the `ProtocolNameList`. For more information, see [Protocols](#) in the *AWS IoT Developer Guide*.

Note

These examples use the Amazon Trust Services (ATS) endpoint, which is used with ATS root CA certificates (recommended). Endpoints must match the root CA certificate type. For more information, see [the section called “Service endpoints must match the certificate type”](#).

Response

Upon success, the response includes the standard HTTP headers plus the following code and body:

```
HTTP 200
BODY: response document
```

For more information, see [Example discover response documents](#).

Discovery authorization

Retrieving the connectivity information requires a policy that allows the caller to perform the `greengrass:Discover` action. TLS mutual authentication with a client certificate is the only accepted form of authentication. The following is an example policy that allows a caller to perform this action:

```
{
```

```

"Version": "2012-10-17",
"Statement": [{
  "Effect": "Allow",
  "Action": "greengrass:Discover",
  "Resource": ["arn:aws:iot:us-west-2:123456789012:thing/MyThingName"]
}]
}

```

Example discover response documents

The following document shows the response for a client device that is a member of a group with one Greengrass core, one endpoint, and one group CA certificate:

```

{
  "GGGroups": [
    {
      "GGGroupId": "gg-group-01-id",
      "Cores": [
        {
          "thingArn": "core-01-thing-arn",
          "Connectivity": [
            {
              "id": "core-01-connection-id",
              "hostAddress": "core-01-address",
              "portNumber": core-01-port,
              "metadata": "core-01-description"
            }
          ]
        }
      ],
      "CAs": [
        "-----BEGIN CERTIFICATE-----cert-contents-----END CERTIFICATE-----"
      ]
    }
  ]
}

```

The following document shows the response for a client device that is a member of two groups with one Greengrass core, multiple endpoints, and multiple group CA certificates:

```

{
  "GGGroups": [

```

```

{
  "GGGroupId": "gg-group-01-id",
  "Cores": [
    {
      "thingArn": "core-01-thing-arn",
      "Connectivity": [
        {
          "id": "core-01-connection-id",
          "hostAddress": "core-01-address",
          "portNumber": core-01-port,
          "metadata": "core-01-connection-1-description"
        },
        {
          "id": "core-01-connection-id-2",
          "hostAddress": "core-01-address-2",
          "portNumber": core-01-port-2,
          "metadata": "core-01-connection-2-description"
        }
      ]
    }
  ],
  "CAs": [
    "-----BEGIN CERTIFICATE-----cert-contents-----END CERTIFICATE-----",
    "-----BEGIN CERTIFICATE-----cert-contents-----END CERTIFICATE-----",
    "-----BEGIN CERTIFICATE-----cert-contents-----END CERTIFICATE-----"
  ]
},
{
  "GGGroupId": "gg-group-02-id",
  "Cores": [
    {
      "thingArn": "core-02-thing-arn",
      "Connectivity" : [
        {
          "id": "core-02-connection-id",
          "hostAddress": "core-02-address",
          "portNumber": core-02-port,
          "metadata": "core-02-connection-1-description"
        }
      ],
      "CAs": [
        "-----BEGIN CERTIFICATE-----cert-contents-----END CERTIFICATE-----",
        "-----BEGIN CERTIFICATE-----cert-contents-----END CERTIFICATE-----",
        "-----BEGIN CERTIFICATE-----cert-contents-----END CERTIFICATE-----"
      ]
    }
  ]
}

```

```

    ]
  }
}
}

```

Note

A Greengrass group must define exactly one Greengrass core. Any response from the AWS IoT Greengrass service that contains a list of Greengrass cores contains only one Greengrass core.

If you have cURL installed, you can test the discovery request. For example:

```

$ curl --cert 1a23bc4d56.cert.pem --key 1a23bc4d56.private.key https://greengrass-ats.iot.us-west-2.amazonaws.com:8443/greengrass/discover/thing/MyDevice
{"GGGroups":[{"GGGroupId":"1234a5b6-78cd-901e-2fgh-3i45j6k1789","Cores":[{"thingArn":"arn:aws:iot:us-west-2:123456789012:thing/MyFirstGroup_Core","Connectivity":[{"Id":"AUTOIP_192.168.1.4_1","HostAddress":"192.168.1.5","PortNumber":8883,"Metadata":""}]}],"CAs":["-----BEGIN CERTIFICATE-----\ncert-contents\n-----END CERTIFICATE-----\n"]}]}

```

Security in AWS IoT Greengrass

Cloud security at AWS is the highest priority. As an AWS customer, you benefit from a data center and network architecture that is built to meet the requirements of the most security-sensitive organizations.

Security is a shared responsibility between AWS and you. The [shared responsibility model](#) describes this as security *of* the cloud and security *in* the cloud:

- **Security of the cloud** – AWS is responsible for protecting the infrastructure that runs AWS services in the AWS Cloud. AWS also provides you with services that you can use securely. Third-party auditors regularly test and verify the effectiveness of our security as part of the [AWS Compliance Programs](#). To learn about the compliance programs that apply to AWS IoT Greengrass, see [AWS Services in Scope by Compliance Program](#).
- **Security in the cloud** – Your responsibility is determined by the AWS service that you use. You are also responsible for other factors, including the sensitivity of your data, your company's requirements, and applicable laws and regulations.

When you use AWS IoT Greengrass, you are also responsible for securing your devices, local network connection, and private keys.

This documentation helps you understand how to apply the shared responsibility model when using AWS IoT Greengrass. The following topics show you how to configure AWS IoT Greengrass to meet your security and compliance objectives. You also learn how to use other AWS services that help you to monitor and secure your AWS IoT Greengrass resources.

Topics


- [Overview of AWS IoT Greengrass security](#)
- [Data protection in AWS IoT Greengrass](#)
- [Device authentication and authorization for AWS IoT Greengrass](#)
- [Identity and access management for AWS IoT Greengrass](#)
- [Compliance validation for AWS IoT Greengrass](#)
- [Resilience in AWS IoT Greengrass](#)
- [Infrastructure security in AWS IoT Greengrass](#)

- [Configuration and vulnerability analysis in AWS IoT Greengrass](#)
- [AWS IoT Greengrass and interface VPC endpoints \(AWS PrivateLink\)](#)
- [Security best practices for AWS IoT Greengrass](#)

Overview of AWS IoT Greengrass security

AWS IoT Greengrass uses X.509 certificates, AWS IoT policies, and IAM policies and roles to secure the applications that run on devices in your local Greengrass environment.

The following diagram shows the components of the AWS IoT Greengrass security model:



A - Greengrass service role

A customer-created IAM role assumed by AWS IoT Greengrass when accessing to your AWS resources from AWS IoT Core, AWS Lambda, and other AWS services. For more information, see [the section called “Greengrass service role”](#).

B - Core device certificate

An X.509 certificate used to authenticate a Greengrass core with AWS IoT Core and AWS IoT Greengrass. For more information, see [the section called “Device authentication and authorization”](#).

C - Device certificate

An X.509 certificate used to authenticate a client device, which is also known as a connected device, with AWS IoT Core and AWS IoT Greengrass. For more information, see [the section called “Device authentication and authorization”](#).

D - Group role

A customer-created IAM role assumed by AWS IoT Greengrass when calling AWS services from a Greengrass core.

You use this role to specify access permissions that your user-defined Lambda functions and connectors need to access AWS services, such as DynamoDB. You also use it to allow AWS IoT Greengrass to export stream manager streams to AWS services and write to CloudWatch Logs. For more information, see [the section called “Greengrass group role”](#).

Note

AWS IoT Greengrass doesn't use the Lambda execution role that's specified in AWS Lambda for the cloud version of a Lambda function.

E - MQTT server certificate

The certificate used for Transport Layer Security (TLS) mutual authentication between a Greengrass core device and client devices in the Greengrass group. The certificate is signed by the group CA certificate, which is stored in the AWS Cloud.

Device connection workflow

This section describes how client devices connect to the AWS IoT Greengrass service and Greengrass core devices. Client devices are registered AWS IoT Core devices that are in the same Greengrass group as the core device.

- A Greengrass core device uses its device certificate, private key, and the AWS IoT Core root CA certificate to connect to the AWS IoT Greengrass service. On the core device, the crypto object in the [configuration file](#) specifies the file path for these items.
- The Greengrass core device downloads group membership information from the AWS IoT Greengrass service.

- When a deployment is made to the Greengrass core device, the Device Certificate Manager (DCM) handles local server certificate management for the Greengrass core device.
- A client device connects to the AWS IoT Greengrass service using its device certificate, private key, and the AWS IoT Core root CA certificate. After making the connection, the client device uses the Greengrass Discovery Service to find the IP address of its Greengrass core device. The client device also downloads the group CA certificate, which is used for TLS mutual authentication with the Greengrass core device.
- A client device attempts to connect to the Greengrass core device, passing its device certificate and client ID. If the client ID matches the thing name of the client device and the certificate is valid (part of the Greengrass group), the connection is made. Otherwise, the connection is terminated.

The AWS IoT policy for client devices must grant the `greengrass:Discover` permission to allow client devices to discover connectivity information for the core. For more information about the policy statement, see [the section called “Discovery authorization”](#).

Configuring AWS IoT Greengrass security

To configure your Greengrass application's security

1. Create an AWS IoT Core thing for your Greengrass core device.
2. Generate a key pair and device certificate for your Greengrass core device.
3. Create and attach an [AWS IoT policy](#) to the device certificate. The certificate and policy allow the Greengrass core device access to AWS IoT Core and AWS IoT Greengrass services. For more information, see [Minimal AWS IoT policy for the core device](#).

Note

The use of [thing policy variables](#) (`iot:Connection.Thing.*`) in the AWS IoT policy for a core device is not supported. The core uses the same device certificate to make [multiple connections](#) to AWS IoT Core but the client ID in a connection might not be an exact match of the core thing name.

4. Create a [Greengrass service role](#). This IAM role authorizes AWS IoT Greengrass to access resources from other AWS services on your behalf. This allows AWS IoT Greengrass to perform essential tasks, such as retrieving AWS Lambda functions and managing device shadows.

You can use the same service role across AWS Regions, but it must be associated with your AWS account in every AWS Region where you use AWS IoT Greengrass.

5. (Optional) Create a [Greengrass group role](#). This IAM role grants permission to Lambda functions and connectors running on a Greengrass core to call AWS services. For example, the [Kinesis Firehose connector](#) requires permission to write records to an Amazon Data Firehose delivery stream.

You can attach only one role to a Greengrass group.

6. Create an AWS IoT Core thing for each device that connects to your Greengrass core.

Note

You can also use existing AWS IoT Core things and certificates.

7. Create device certificates, key pairs, and AWS IoT policies for each device that connects to your Greengrass core.

AWS IoT Greengrass core security principals

The Greengrass core uses the following security principals: AWS IoT client, local MQTT server, and local secrets manager. The configuration for these principals is stored in the `crypto` object in the `config.json` configuration file. For more information, see [the section called "AWS IoT Greengrass core configuration file"](#).

This configuration includes the path to the private key used by the principal component for authentication and encryption. AWS IoT Greengrass supports two modes of private key storage: hardware-based or file system-based (default). For more information about storing keys on hardware security modules, see [the section called "Hardware security integration"](#).

AWS IoT Client

The AWS IoT client (IoT client) manages communication over the internet between the Greengrass core and AWS IoT Core. AWS IoT Greengrass uses X.509 certificates with public and private keys for mutual authentication when establishing TLS connections for this communication. For more information, see [X.509 certificates and AWS IoT Core](#) in the *AWS IoT Core Developer Guide*.

The IoT client supports RSA and EC certificates and keys. The certificate and private key path are specified for the `IoTCertificate` principal in `config.json`.

MQTT Server

The local MQTT server manages communication over the local network between the Greengrass core and client devices in the group. AWS IoT Greengrass uses X.509 certificates with public and private keys for mutual authentication when establishing TLS connections for this communication.

By default, AWS IoT Greengrass generates an RSA private key for you. To configure the core to use a different private key, you must provide the key path for the `MQTTServerCertificate` principal in `config.json`. You are responsible for rotating a customer-provided key.

Private key support

	RSA key	EC key
Key type	Supported	Supported
Key parameters	Minimum 2048-bit length	NIST P-256 or NIST P-384 curve
Disk format	PKCS#1, PKCS#8	SECG1, PKCS#8
Minimum GGC version	- Use default RSA key: 1.0 - Specify an RSA key: 1.7	Specify an EC key: 1.9

The configuration of the private key determines related processes. For the list of cipher suites that the Greengrass core supports as a server, see [the section called “TLS cipher suites support”](#).

If no private key is specified (default)

- AWS IoT Greengrass rotates the key based on your rotation settings.
- The core generates an RSA key, which is used to generate the certificate.
- The MQTT server certificate has an RSA public key and an SHA-256 RSA signature.

If an RSA private key is specified (requires GGC v1.7 or later)

- You are responsible for rotating the key.
- The core uses the specified key to generate the certificate.
- The RSA key must have a minimum length of 2048 bits.

- The MQTT server certificate has an RSA public key and an SHA-256 RSA signature.

If an EC private key is specified (requires GGC v1.9 or later)

- You are responsible for rotating the key.
- The core uses the specified key to generate the certificate.
- The EC private key must use an NIST P-256 or NIST P-384 curve.
- The MQTT server certificate has an EC public key and an SHA-256 RSA signature.

The MQTT server certificate presented by the core has an SHA-256 RSA signature, regardless of the key type. For this reason, clients must support SHA-256 RSA certificate validation to establish a secure connection with the core.

Secrets Manager

The local secrets manager securely manages local copies of secrets that you create in AWS Secrets Manager. It uses a private key to secure the data key that's used to encrypt the secrets. For more information, see [Deploy secrets to the core](#).

By default, the IoT client private key is used, but you can specify a different private key for the SecretsManager principal in `config.json`. Only the RSA key type is supported. For more information, see [the section called "Specify the private key for secret encryption"](#).

Note

Currently, AWS IoT Greengrass supports only the [PKCS#1 v1.5](#) padding mechanism for encryption and decryption of local secrets when using hardware-based private keys. If you're following vendor-provided instructions to manually generate hardware-based private keys, make sure to choose PKCS#1 v1.5. AWS IoT Greengrass doesn't support Optimal Asymmetric Encryption Padding (OAEP).

Private key support

	RSA key	EC key
Key type	Supported	Not supported
Key parameters	Minimum 2048-bit length	Not applicable
Disk format	PKCS#1, PKCS#8	Not applicable

	RSA key	EC key
Minimum GGC version	1.7	Not applicable

Managed subscriptions in the MQTT messaging workflow

AWS IoT Greengrass uses a subscription table to define how MQTT messages can be exchanged between client devices, functions, and connectors in a Greengrass group, and with AWS IoT Core or the local shadow service. Each subscription specifies a source, target, and MQTT topic (or subject) over which messages are sent or received. AWS IoT Greengrass allows messages to be sent from a source to a target only if a corresponding subscription is defined.

A subscription defines the message flow in one direction only, from the source to the target. To support two-way message exchange, you must create two subscriptions, one for each direction.

TLS cipher suites support

AWS IoT Greengrass uses the AWS IoT Core transport security model to encrypt communication with the cloud by using [TLS cipher suites](#). In addition, AWS IoT Greengrass data is encrypted when at rest (in the cloud). For more information about AWS IoT Core transport security and supported cipher suites, see [Transport security](#) in the *AWS IoT Core Developer Guide*.

Supported Cipher Suites for Local Network Communication

As opposed to AWS IoT Core, the AWS IoT Greengrass core supports the following *local network* TLS cipher suites for certificate-signing algorithms. All of these cipher suites are supported when private keys are stored on the file system. A subset are supported when the core is configured to use hardware security modules (HSM). For more information, see [the section called “Security principals”](#) and [the section called “Hardware security integration”](#). The table also includes the minimum version of AWS IoT Greengrass Core software required for support.

	Cipher	HSM support	Minimum GGC version
TLSv1.2	TLS_ECDHE _RSA_WITH _AES_128_CBC_SHA	Supported	1.0

Cipher	HSM support	Minimum GGC version
TLS_ECDHE _RSA_WITH _AES_256_CBC_SHA	Supported	1.0
TLS_ECDHE _RSA_WITH _AES_256_ GCM_SHA384	Supported	1.0
TLS_RSA_W ITH_AES_1 28_CBC_SHA	Not supported	1.0
TLS_RSA_W ITH_AES_1 28_GCM_SHA256	Not supported	1.0
TLS_RSA_W ITH_AES_2 56_CBC_SHA	Not supported	1.0
TLS_RSA_W ITH_AES_2 56_GCM_SHA384	Not supported	1.0
TLS_ECDHE _ECDSA_WI TH_AES_12 8_GCM_SHA256	Supported	1.9
TLS_ECDHE _ECDSA_WI TH_AES_25 6_GCM_SHA384	Supported	1.9

	Cipher	HSM support	Minimum GGC version
TLSv1.1	TLS_ECDHE _RSA_WITH _AES_128_CBC_SHA	Supported	1.0
	TLS_ECDHE _RSA_WITH _AES_256_CBC_SHA	Supported	1.0
	TLS_RSA_W ITH_AES_1 28_CBC_SHA	Not supported	1.0
	TLS_RSA_W ITH_AES_2 56_CBC_SHA	Not supported	1.0
TLSv1.0	TLS_ECDHE _RSA_WITH _AES_128_CBC_SHA	Supported	1.0
	TLS_ECDHE _RSA_WITH _AES_256_CBC_SHA	Supported	1.0
	TLS_RSA_W ITH_AES_1 28_CBC_SHA	Not supported	1.0
	TLS_RSA_W ITH_AES_2 56_CBC_SHA	Not supported	1.0

Data protection in AWS IoT Greengrass

The AWS [shared responsibility model](#) applies to data protection in AWS IoT Greengrass. As described in this model, AWS is responsible for protecting the global infrastructure that runs all of the AWS Cloud. You are responsible for maintaining control over your content that is hosted on this infrastructure. You are also responsible for the security configuration and management tasks for the AWS services that you use. For more information about data privacy, see the [Data Privacy FAQ](#). For information about data protection in Europe, see the [AWS Shared Responsibility Model and GDPR](#) blog post on the *AWS Security Blog*.

For data protection purposes, we recommend that you protect AWS account credentials and set up individual users with AWS IAM Identity Center or AWS Identity and Access Management (IAM). That way, each user is given only the permissions necessary to fulfill their job duties. We also recommend that you secure your data in the following ways:

- Use multi-factor authentication (MFA) with each account.
- Use SSL/TLS to communicate with AWS resources. We require TLS 1.2 and recommend TLS 1.3.
- Set up API and user activity logging with AWS CloudTrail. For information about using CloudTrail trails to capture AWS activities, see [Working with CloudTrail trails](#) in the *AWS CloudTrail User Guide*.
- Use AWS encryption solutions, along with all default security controls within AWS services.
- Use advanced managed security services such as Amazon Macie, which assists in discovering and securing sensitive data that is stored in Amazon S3.
- If you require FIPS 140-3 validated cryptographic modules when accessing AWS through a command line interface or an API, use a FIPS endpoint. For more information about the available FIPS endpoints, see [Federal Information Processing Standard \(FIPS\) 140-3](#).

We strongly recommend that you never put confidential or sensitive information, such as your customers' email addresses, into tags or free-form text fields such as a **Name** field. This includes when you work with AWS IoT Greengrass or other AWS services using the console, API, AWS CLI, or AWS SDKs. Any data that you enter into tags or free-form text fields used for names may be used for billing or diagnostic logs. If you provide a URL to an external server, we strongly recommend that you do not include credentials information in the URL to validate your request to that server.

For more information about protecting sensitive information in AWS IoT Greengrass, see [the section called “Don't log sensitive information”](#).

For more information about data protection, see the [AWS Shared Responsibility Model and GDPR](#) blog post on the *AWS Security Blog*.

Topics

- [Data encryption](#)
- [Hardware security integration](#)

Data encryption

AWS IoT Greengrass uses encryption to protect data while in-transit (over the internet or local network) and at rest (stored in the AWS Cloud).

Devices in a AWS IoT Greengrass environment often collect data that's sent to AWS services for further processing. For more information about data encryption on other AWS services, see the security documentation for that service.

Topics

- [Encryption in transit](#)
- [Encryption at rest](#)
- [Key management for the Greengrass core device](#)

Encryption in transit

AWS IoT Greengrass has three modes of communication where data is in transit:

- [the section called "Data in transit over the internet"](#). Communication between a Greengrass core and AWS IoT Greengrass over the internet is encrypted.
- [the section called "Data in transit over the local network"](#). Communication between a Greengrass core and client devices over a local network is encrypted.
- [the section called "Data on the core device"](#). Communication between components on the Greengrass core device is not encrypted.

Data in transit over the internet

AWS IoT Greengrass uses Transport Layer Security (TLS) to encrypt all communication over the internet. All data sent to the AWS Cloud is sent over a TLS connection using MQTT or HTTPS

protocols, so it is secure by default. AWS IoT Greengrass uses the AWS IoT transport security model. For more information, see [Transport security](#) in the *AWS IoT Core Developer Guide*.

Data in transit over the local network

AWS IoT Greengrass uses TLS to encrypt all communication over the local network between the Greengrass core and client devices. For more information, see [Supported Cipher Suites for Local Network Communication](#).

It is your responsibility to protect the local network and private keys.

For Greengrass core devices, it's your responsibility to:

- Keep the kernel updated with the latest security patches.
- Keep system libraries updated with the latest security patches.
- Protect private keys. For more information, see [the section called "Key management"](#).

For client devices, it's your responsibility to:

- Keep the TLS stack up to date.
- Protect private keys.

Data on the core device

AWS IoT Greengrass doesn't encrypt data exchanged locally on the Greengrass core device because the data doesn't leave the device. This includes communication between user-defined Lambda functions, connectors, the AWS IoT Greengrass Core SDK, and system components, such as stream manager.

Encryption at rest

AWS IoT Greengrass stores your data:

- [the section called "Data at rest in the AWS Cloud"](#). This data is encrypted.
- [the section called "Data at rest on the Greengrass core"](#). This data is not encrypted (except local copies of your secrets).

Data at rest in the AWS Cloud

AWS IoT Greengrass encrypts customer data stored in the AWS Cloud. This data is protected using AWS KMS keys that are managed by AWS IoT Greengrass.

Data at rest on the Greengrass core

AWS IoT Greengrass relies on Unix file permissions and full-disk encryption (if enabled) to protect data at rest on the core. It is your responsibility to secure the file system and device.

However, AWS IoT Greengrass does encrypt local copies of your secrets retrieved from AWS Secrets Manager. For more information, see [the section called “Secrets encryption”](#).

Key management for the Greengrass core device

It's the responsibility of the customer to guarantee secure storage of cryptographic (public and private) keys on the Greengrass core device. AWS IoT Greengrass uses public and private keys for the following scenarios:

- The IoT client key is used with the IoT certificate to authenticate the Transport Layer Security (TLS) handshake when a Greengrass core connects to AWS IoT Core. For more information, see [the section called “Device authentication and authorization”](#).

Note

The key and certificate are also referred to as the core private key and the core device certificate.

- The MQTT server key is used the MQTT server certificate to authenticate TLS connections between core and client devices. For more information, see [the section called “Device authentication and authorization”](#).
- The local secrets manager also uses the IoT client key to protect the data key used to encrypt local secrets, but you can provide your own private key. For more information, see [the section called “Secrets encryption”](#).

A Greengrass core supports private key storage using file system permissions, [hardware security modules](#), or both. If you use file system-based private keys, you are responsible for their secure storage on the core device.

On a Greengrass core, the location of your private keys are specified in the `crypto` section of the `config.json` file. If you configure the core to use a customer-provided key for the MQTT server certificate, it is your responsibility to rotate the key. For more information, see [the section called “Security principals”](#).

For client devices, it's your responsibility to keep the TLS stack up to date and protect private keys. Private keys are used with device certificates to authenticate TLS connections with the AWS IoT Greengrass service.


Hardware security integration

This feature is available for AWS IoT Greengrass Core v1.7 and later.

AWS IoT Greengrass supports the use of hardware security modules (HSM) through the [PKCS#11 interface](#) for secure storage and offloading of private keys. This prevents keys from being exposed or duplicated in software. Private keys can be securely stored on hardware modules, such as HSMs, Trusted Platform Modules (TPM), or other cryptographic elements.

Search for devices that are qualified for this feature in the [AWS Partner Device Catalog](#).

The following diagram shows the hardware security architecture for an AWS IoT Greengrass core.



On a standard installation, AWS IoT Greengrass uses two private keys. One key is used by the AWS IoT client (IoT client) component during the Transport Layer Security (TLS) handshake when a Greengrass core connects to AWS IoT Core. (This key is also referred to as the core private key.) The other key is used by the local MQTT server, which enables Greengrass devices to communicate with the Greengrass core. If you want to use hardware security for both components, you can use a shared private key or separate private keys. For more information, see [the section called "Provisioning practices"](#).

Note

On a standard installation, the local secrets manager also uses the IoT client key for its encryption process, but you can use your own private key. It must be an RSA key with a minimum length of 2048 bits. For more information, see [the section called “Specify the private key for secret encryption”](#).

Requirements

Before you can configure hardware security for a Greengrass core, you must have the following:

- A hardware security module (HSM) that supports your target private key configuration for the IoT client, local MQTT server, and local secrets manager components. The configuration can include one, two, or three hardware-based private keys, depending on whether you configure the components to share keys. For more information about private key support, see [the section called “Security principals”](#).
- For RSA keys: An RSA-2048 key size (or larger) and [PKCS#1 v1.5](#) signature scheme.
- For EC keys: An NIST P-256 or NIST P-384 curve.

Note

Search for devices that are qualified for this feature in the [AWS Partner Device Catalog](#).

- A PKCS#11 provider library that is loadable at runtime (using libdl) and provides [PKCS#11](#) functions.
- The hardware module must be resolvable by slot label, as defined in the PKCS#11 specification.
- The private key must be generated and loaded on the HSM by using the vendor-provided provisioning tools.
- The private key must be resolvable by object label.
- The core device certificate. This is an IoT client certificate that corresponds to the private key.
- If you're using the Greengrass OTA update agent, the [OpenSSL libp11 PKCS#11](#) wrapper library must be installed. For more information, see [the section called “Configure OTA updates”](#).

In addition, make sure that the following conditions are met:

- The IoT client certificates that are associated with the private key are registered in AWS IoT and activated. You can verify this in the AWS IoT console under **Manage**, expand **All devices**, choose **Things** and choose the **Certificates** tab for the core thing.
- The AWS IoT Greengrass Core software v1.7 or later is installed on the core device, as described in [Module 2](#) of the Getting Started tutorial. Version 1.9 or later is required to use an EC key for the MQTT server.
- The certificates are attached to the Greengrass core. You can verify this from the **Manage** page for the core thing in the AWS IoT console.

Note

Currently, AWS IoT Greengrass doesn't support loading the CA certificate or IoT client certificate directly from the HSM. The certificates must be loaded as plain-text files on the file system in a location that can be read by Greengrass.

Hardware security configuration for an AWS IoT Greengrass core

Hardware security is configured in the Greengrass configuration file. This is the [config.json](#) file that's located in the `/greengrass-root/config` directory.

Note

To walk through the process of setting up an HSM configuration using a pure software implementation, see [the section called "Module 7: Simulating hardware security integration"](#).

Important

The simulated configuration in the example doesn't provide any security benefits. It's intended to allow you to learn about the PKCS#11 specification and do initial testing of your software if you plan to use a hardware-based HSM in the future.

To configure hardware security in AWS IoT Greengrass, you edit the `crypto` object in `config.json`.

When using hardware security, the crypto object is used to specify paths to certificates, private keys, and assets for the PKCS#11 provider library on the core, as shown in the following example.

```
"crypto": {
  "PKCS11" : {
    "OpenSSLEngine" : "/path-to-p11-openssl-engine",
    "P11Provider" : "/path-to-pkcs11-provider-so",
    "slotLabel" : "crypto-token-name",
    "slotUserPin" : "crypto-token-user-pin"
  },
  "principals" : {
    "IoTCertificate" : {
      "privateKeyPath" : "pkcs11:object=core-private-key-label;type=private",
      "certificatePath" : "file:///path-to-core-device-certificate"
    },
    "MQTTServerCertificate" : {
      "privateKeyPath" : "pkcs11:object=server-private-key-label;type=private"
    },
    "SecretsManager" : {
      "privateKeyPath": "pkcs11:object=core-private-key-label;type=private"
    }
  },
  "caPath" : "file:///path-to-root-ca"
```

The crypto object contains the following properties:

Field	Description	Notes
caPath	The absolute path to the AWS IoT root CA.	Must be a file URI of the form: file:///absolute/path/to/file .

 Note

Make sure that your [endpoints correspond to your certificate type](#).

PKCS11

Field	Description	Notes
OpenSSLEngine	Optional. The absolute path to the OpenSSL engine .so file to enable PKCS#11 support on OpenSSL.	Must be a path to a file on the file system. This property is required if you're using the Greengrass OTA update agent with hardware security. For more information, see the section called "Configure OTA updates" .
P11Provider	The absolute path to the PKCS#11 implementation's libdl-loadable library.	Must be a path to a file on the file system.
slotLabel	The slot label that's used to identify the hardware module.	Must conform to PKCS#11 label specifications.
slotUserPin	The user PIN that's used to authenticate the Greengrass core to the module.	Must have sufficient permissions to perform C_Sign with the configured private keys.
principals		
IoTCertificate	The certificate and private key that the core uses to make requests to AWS IoT.	
IoTCertificate .privateKeyPath	The path to the core private key.	For file system storage, must be a file URI of the form: file:///absolute/path/to/file . For HSM storage, must be an RFC 7512 PKCS#11 path that specifies the object label.

Field	Description	Notes
IoTCertificate .certificatePath	The absolute path to the core device certificate.	Must be a file URI of the form: <code>file:///absolute/path/to/file</code> .
MQTTServerCertificate	Optional. The private key that the core uses in combination with the certificate to act as an MQTT server or gateway.	
MQTTServerCertificate .privateKeyPath	The path to the local MQTT server private key.	Use this value to specify your own private key for the local MQTT server. For file system storage, must be a file URI of the form: <code>file:///absolute/path/to/file</code>. For HSM storage, must be an RFC 7512 PKCS#11 path that specifies the object label. If this property is omitted, AWS IoT Greengrass rotates the key based your rotation settings. If specified, the customer is responsible for rotating the key.
SecretsManager	The private key that secures the data key used for encryption. For more information, see Deploy secrets to the core .	

Field	Description	Notes
SecretsManager .privateKeyPath	The path to the local secrets manager private key.	Only an RSA key is supported. For file system storage, must be a file URI of the form: <code>file:///absolute/path/to/file</code>. For HSM storage, must be an RFC 7512 PKCS#11 path that specifies the object label. The private key must be generated using the PKCS#1 v1.5 padding mechanism.

Field	Description	Notes
caPath	The absolute path to the AWS IoT root CA.	Must be a file URI of the form: <code>file:///absolute/path/to/file</code>. Note Make sure that your endpoints correspond to your certificate type.
PKCS11		
OpenSSLEngine	Optional. The absolute path to the OpenSSL engine .so file to enable PKCS#11 support on OpenSSL.	Must be a path to a file on the file system. This property is required if you're using the Greengrass OTA update agent with

Field	Description	Notes
		hardware security. For more information, see the section called "Configure OTA updates" .
P11Provider	The absolute path to the PKCS#11 implementation's libdl-loadable library.	Must be a path to a file on the file system.
slotLabel	The slot label that's used to identify the hardware module.	Must conform to PKCS#11 label specifications.
slotUserPin	The user PIN that's used to authenticate the Greengrass core to the module.	Must have sufficient permissions to perform C_Sign with the configured private keys.
principals		
IoTCertificate	The certificate and private key that the core uses to make requests to AWS IoT.	
IoTCertificate .privateKeyPath	The path to the core private key.	For file system storage, must be a file URI of the form: file:///absolute/path/to/file . For HSM storage, must be an RFC 7512 PKCS#11 path that specifies the object label.
IoTCertificate .certificatePath	The absolute path to the core device certificate.	Must be a file URI of the form: file:///absolute/path/to/file .
MQTTServerCertificate	Optional. The private key that the core uses in combination with the certificate to act as an MQTT server or gateway.	

Field	Description	Notes
MQTTServerCertificate.privateKeyPath	The path to the local MQTT server private key.	Use this value to specify your own private key for the local MQTT server. For file system storage, must be a file URI of the form: <code>file:///absolute/path/to/file</code>. For HSM storage, must be an RFC 7512 PKCS#11 path that specifies the object label. If this property is omitted, AWS IoT Greengrass rotates the key based your rotation settings. If specified, the customer is responsible for rotating the key.
SecretsManager	The private key that secures the data key used for encryption. For more information, see Deploy secrets to the core .	
SecretsManager.privateKeyPath	The path to the local secrets manager private key.	Only an RSA key is supported. For file system storage, must be a file URI of the form: <code>file:///absolute/path/to/file</code>. For HSM storage, must be an RFC 7512 PKCS#11 path that specifies the object label. The private key must be generated using the PKCS#1 v1.5 padding mechanism.

Field	Description	Notes
caPath	The absolute path to the AWS IoT root CA.	Must be a file URI of the form: <code>file:///absolute/path/to/file</code> .
		 Note Make sure that your endpoints correspond to your certificate type.
PKCS11		
OpenSSLEngine	Optional. The absolute path to the OpenSSL engine .so file to enable PKCS#11 support on OpenSSL.	Must be a path to a file on the file system. This property is required if you're using the Greengrass OTA update agent with hardware security. For more information, see the section called "Configure OTA updates".
P11Provider	The absolute path to the PKCS#11 implementation's libdl-loadable library.	Must be a path to a file on the file system.
slotLabel	The slot label that's used to identify the hardware module.	Must conform to PKCS#11 label specifications.
slotUserPin	The user PIN that's used to authenticate the Greengrass core to the module.	Must have sufficient permissions to perform C_Sign with the configured private keys.

Field	Description	Notes
principals		
IoTCertificate	The certificate and private key that the core uses to make requests to AWS IoT.	
IoTCertificate .privateKeyPath	The path to the core private key.	For file system storage, must be a file URI of the form: file:///absolute/path/to/file . For HSM storage, must be an RFC 7512 PKCS#11 path that specifies the object label.
IoTCertificate .certificatePath	The absolute path to the core device certificate.	Must be a file URI of the form: file:///absolute/path/to/file .
MQTTServerCertificate	Optional. The private key that the core uses in combination with the certificate to act as an MQTT server or gateway.	

Field	Description	Notes
MQTTServerCertificate.privateKeyPath	The path to the local MQTT server private key.	Use this value to specify your own private key for the local MQTT server. For file system storage, must be a file URI of the form: <code>file:///absolute/path/to/file</code> . For HSM storage, must be an RFC 7512 PKCS#11 path that specifies the object label. If this property is omitted, AWS IoT Greengrass rotates the key based your rotation settings. If specified, the customer is responsible for rotating the key.
SecretsManager	The private key that secures the data key used for encryption. For more information, see Deploy secrets to the core .	
SecretsManager.privateKeyPath	The path to the local secrets manager private key.	Only an RSA key is supported. For file system storage, must be a file URI of the form: <code>file:///absolute/path/to/file</code> . For HSM storage, must be an RFC 7512 PKCS#11 path that specifies the object label. The private key must be generated using the PKCS#1 v1.5 padding mechanism.

Provisioning practices for AWS IoT Greengrass hardware security

The following are security and performance-related provisioning practices.

Security

- Generate private keys directly on the HSM by using the internal hardware random-number generator.

Note

If you configure private keys to use with this feature (by following the instructions provided by the hardware vendor), be aware that AWS IoT Greengrass currently supports only the PKCS1 v1.5 padding mechanism for encryption and decryption of [local secrets](#). AWS IoT Greengrass doesn't support Optimal Asymmetric Encryption Padding (OAEP).

- Configure private keys to prohibit export.
- Use the provisioning tool that's provided by the hardware vendor to generate a certificate signing request (CSR) using the hardware-protected private key, and then use the AWS IoT console to generate a client certificate.

Note

The practice of rotating keys doesn't apply when private keys are generated on an HSM.


Performance

The following diagram shows the IoT client component and local MQTT server on the AWS IoT Greengrass core. If you want to use an HSM configuration for both components, you can use the same private key or separate private keys. If you use separate keys, they must be stored in the same slot.

Note

AWS IoT Greengrass doesn't impose any limits on the number of keys that you store on the HSM, so you can store private keys for the IoT client, MQTT server, and secrets

manager components. However, some HSM vendors might impose limits on the number of keys you can store in a slot.



In general, the IoT client key is not used very frequently because the AWS IoT Greengrass Core software maintains long-lived connections to the cloud. However, the MQTT server key is used every time that a Greengrass device connects to the core. These interactions directly affect performance.

When the MQTT server key is stored on the HSM, the rate at which devices can connect depends on the number of RSA signature operations per second that the HSM can perform. For example, if the HSM takes 300 milliseconds to perform an RSASSA-PKCS1-v1.5 signature on an RSA-2048 private key, then only three devices can connect to the Greengrass core per second. After the connections are made, the HSM is no longer used and the standard [quotas for AWS IoT Greengrass](#) apply.

To mitigate performance bottlenecks, you can store the private key for the MQTT server on the file system instead of on the HSM. With this configuration, the MQTT server behaves as if hardware security isn't enabled.

AWS IoT Greengrass supports multiple key-storage configurations for the IoT client and MQTT server components, so you can optimize for your security and performance requirements. The following table includes example configurations.

Configuration	IoT key	MQTT key	Performance
HSM Shared Key	HSM: Key A	HSM: Key A	Limited by the HSM or CPU
HSM Separate Keys	HSM: Key A	HSM: Key B	Limited by the HSM or CPU
HSM for IoT only	HSM: Key A	File System: Key B	Limited by the CPU
Legacy	File System: Key A	File System: Key B	Limited by the CPU

To configure the Greengrass core to use file system-based keys for the MQTT server, omit the `principals.MQTTServerCertificate` section from `config.json` (or specify a file-based path to the key if you're not using the default key generated by AWS IoT Greengrass). The resulting crypto object looks like this:

```
"crypto": {
  "PKCS11": {
    "OpenSSLEngine": "...",
    "P11Provider": "...",
    "slotLabel": "...",
    "slotUserPin": "..."
  },
  "principals": {
    "IoTCertificate": {
      "privateKeyPath": "...",
      "certificatePath": "..."
    },
    "SecretsManager": {
      "privateKeyPath": "..."
    }
  },
  "caPath" : "..."
}
```

Supported cipher suites for hardware security integration

AWS IoT Greengrass supports a set of cipher suites when the core is configured for hardware security. This is a subset of the cipher suites that are supported when the core is configured to use file-based security. For more information, see [the section called “TLS cipher suites support”](#).

Note

When connecting to the Greengrass core from Greengrass devices over the local network, be sure to use one of the supported cipher suites to make the TLS connection.

Configure support for over-the-air updates

To enable over-the-air (OTA) updates of the AWS IoT Greengrass Core software when using hardware security, you must install the OpenSC libp11 [PKCS#11 wrapper library](#) and edit the Greengrass configuration file. For more information about OTA updates, see [OTA updates of AWS IoT Greengrass Core software](#).

1. Stop the Greengrass daemon.

```
cd /greengrass-root/ggc/core/  
sudo ./greengrassd stop
```

Note

greengrass-root represents the path where the AWS IoT Greengrass Core software is installed on your device. Typically, this is the /greengrass directory.

2. Install the OpenSSL engine. OpenSSL 1.0 or 1.1 are supported.

```
sudo apt-get install libengine-pkcs11-openssl
```

3. Find the path to the OpenSSL engine (libpkcs11.so) on your system:

- a. Get the list of installed packages for the library.

```
sudo dpkg -l libengine-pkcs11-openssl
```

The `libpkcs11.so` file is located in the `engines` directory.

- b. Copy the full path to the file (for example, `/usr/lib/ssl/engines/libpkcs11.so`).
4. Open the Greengrass configuration file. This is the [config.json](#) file in the `/greengrass-root/config` directory.
5. For the `OpenSSL Engine` property, enter the path to the `libpkcs11.so` file.

```
{
  "crypto": {
    "caPath" : "file:///path-to-root-ca",
    "PKCS11" : {
      "OpenSSL Engine" : "/path-to-p11-openssl-engine",
      "P11Provider" : "/path-to-pkcs11-provider-so",
      "slotLabel" : "crypto-token-name",
      "slotUserPin" : "crypto-token-user-pin"
    },
    ...
  }
  ...
}
```

Note

If the `OpenSSL Engine` property doesn't exist in the `PKCS11` object, then add it.

6. Start the Greengrass daemon.

```
cd /greengrass-root/ggc/core/
sudo ./greengrassd start
```

Backward compatibility with earlier versions of the AWS IoT Greengrass core software

The AWS IoT Greengrass Core software with hardware security support is fully backward compatible with `config.json` files that are generated for v1.6 and earlier. If the `crypto` object is not present in the `config.json` configuration file, then AWS IoT Greengrass uses the file-based `coreThing.certPath`, `coreThing.keyPath`, and `coreThing.caPath` properties. This

backward compatibility applies to Greengrass OTA updates, which do not overwrite a file-based configuration that's specified in `config.json`.

Hardware without PKCS#11 support

The PKCS#11 library is typically provided by the hardware vendor or is open source. For example, with standards-compliant hardware (such as TPM1.2), it might be possible to use existing open source software. However, if your hardware doesn't have a corresponding PKCS#11 library implementation, or if you want to write a custom PKCS#11 provider, you should contact your AWS Enterprise Support representative with integration-related questions.

See also

- *PKCS #11 Cryptographic Token Interface Usage Guide Version 2.40*. Edited by John Leiseboer and Robert Griffin. 16 November 2014. OASIS Committee Note 02. <http://docs.oasis-open.org/pkcs11/pkcs11-ug/v2.40/cn02/pkcs11-ug-v2.40-cn02.html>. Latest version: <http://docs.oasis-open.org/pkcs11/pkcs11-ug/v2.40/pkcs11-ug-v2.40.html>.
- [RFC 7512](#)
- [PKCS #1: RSA Encryption Version 1.5](#)

Device authentication and authorization for AWS IoT Greengrass

Devices in AWS IoT Greengrass environments use X.509 certificates for authentication and AWS IoT policies for authorization. Certificates and policies allow devices to securely connect with each other, AWS IoT Core, and AWS IoT Greengrass.

X.509 certificates are digital certificates that use the X.509 public key infrastructure standard to associate a public key with the identity contained in a certificate. X.509 certificates are issued by a trusted entity called a certificate authority (CA). The CA maintains one or more special certificates called CA certificates that it uses to issue X.509 certificates. Only the certificate authority has access to CA certificates.

AWS IoT policies define the set of operations allowed for AWS IoT devices. Specifically, they allow and deny access to AWS IoT Core and AWS IoT Greengrass data plane operations, such as publishing MQTT messages and retrieving device shadows.

All devices require an entry in the AWS IoT Core registry and an activated X.509 certificate with an attached AWS IoT policy. Devices fall into two categories:

- **Greengrass cores.** Greengrass core devices use certificates and AWS IoT policies to connect to AWS IoT Core. The certificates and policies also allow AWS IoT Greengrass to deploy configuration information, Lambda functions, connectors, and managed subscriptions to core devices.
- **Client devices.** Client devices (also called *connected devices*, *Greengrass devices*, or *devices*) are devices that connect to a Greengrass core over MQTT. They use certificates and policies to connect to AWS IoT Core and the AWS IoT Greengrass service. This allows client devices to use the AWS IoT Greengrass Discovery Service to find and connect to a core device. A client device uses the same certificate to connect to the AWS IoT Core device gateway and core device. Client devices also use discovery information for mutual authentication with the core device. For more information, see [the section called “Device connection workflow”](#) and [the section called “Manage device authentication with the Greengrass core”](#).

X.509 certificates

Communication between core and client devices and between devices and AWS IoT Core or AWS IoT Greengrass must be authenticated. This mutual authentication is based on registered X.509 device certificates and cryptographic keys.

In an AWS IoT Greengrass environment, devices use certificates with public and private keys for the following Transport Layer Security (TLS) connections:

- The AWS IoT client component on the Greengrass core connecting to AWS IoT Core and AWS IoT Greengrass over the internet.
- Client devices connecting to AWS IoT Greengrass to get core discovery information over the internet.
- The MQTT server component on the Greengrass core connecting to client devices in the group over the local network.

The AWS IoT Greengrass core device stores certificates in two locations:

- Core device certificate in `/greengrass-root/certs`. Typically, the core device certificate is named `hash.cert.pem` (for example, `86c84488a5.cert.pem`). This certificate is used by the

AWS IoT client for mutual authentication when the core connects to the AWS IoT Core and AWS IoT Greengrass services.

- MQTT server certificate in `/greengrass-root/ggc/var/state/server`. The MQTT server certificate is named `server.crt`. This certificate is used for mutual authentication between the local MQTT server (on the Greengrass core) and Greengrass devices.

Note

`greengrass-root` represents the path where the AWS IoT Greengrass Core software is installed on your device. Typically, this is the `/greengrass` directory.

For more information, see [the section called “Security principals”](#).

Certificate authority (CA) certificates

Core devices and client devices download a root CA certificate used for authentication with AWS IoT Core and AWS IoT Greengrass services. We recommend that you use an Amazon Trust Services (ATS) root CA certificate, such as [Amazon Root CA 1](#). For more information, see [CA certificates for server authentication](#) in the *AWS IoT Core Developer Guide*.

Note

Your root CA certificate type must match your endpoint. Use an ATS root CA certificate with an ATS endpoint (preferred) or a VeriSign root CA certificate with a legacy endpoint. Only some Amazon Web Services Regions support legacy endpoints. For more information, see [the section called “Service endpoints must match the certificate type”](#).

Client devices also download the Greengrass group CA certificate. This is used to validate the MQTT server certificate on the Greengrass core during mutual authentication. For more information, see [the section called “Device connection workflow”](#). The default expiration of the MQTT server certificate is seven days.

Certificate rotation on the local MQTT server

Client devices use the local MQTT server certificate for mutual authentication with the Greengrass core device. By default, this certificate expires in seven days. This limited period is based on security

best practices. The MQTT server certificate is signed by the group CA certificate, which is stored in the cloud.

For certificate rotation to occur, your Greengrass core device must be online and able to access the AWS IoT Greengrass service directly on a regular basis. When the certificate expires, the core device attempts to connect to the AWS IoT Greengrass service to obtain a new certificate. If the connection is successful, the core device downloads a new MQTT server certificate and restarts the local MQTT service. At this point, all client devices that are connected to the core are disconnected. If the core device is offline at the time of expiry, it does not receive the replacement certificate. Any new attempts to connect to the core device are rejected. Existing connections are not affected. Client devices cannot connect to the core device until the connection to the AWS IoT Greengrass service is restored and a new MQTT server certificate can be downloaded.

You can set the expiration to any value between 7 and 30 days, depending on your needs. More frequent rotation requires more frequent cloud connection. Less frequent rotation can pose security concerns. If you want to set the certificate expiration to a value higher than 30 days, contact Support.

In the AWS IoT console, you can manage the certificate on the group's **Settings** page. In the AWS IoT Greengrass API, you can use the [UpdateGroupCertificateConfiguration](#) action.

When the MQTT server certificate expires, any attempt to validate the certificate fails. Client devices must be able to detect the failure and terminate the connection.

AWS IoT policies for data plane operations

Use AWS IoT policies to authorize access to the AWS IoT Core and AWS IoT Greengrass data plane. The AWS IoT Core data plane consists of operations for devices, users, and applications, such as connecting to AWS IoT Core and subscribing to topics. The AWS IoT Greengrass data plane consists of operations for Greengrass devices, such as retrieving deployments and updating connectivity information.

An AWS IoT policy is a JSON document that's similar to an [IAM policy](#). It contains one or more policy statements that specify the following properties:

- **Effect.** The access mode, which can be Allow or Deny.
- **Action.** The list of actions that are allowed or denied by the policy.
- **Resource.** The list of resources on which the action is allowed or denied.

AWS IoT policies support `*` as a wildcard character, and treat MQTT wildcard characters (`+` and `#`) as literal strings. For more information about the `*` wildcard, see [Using wildcard in resource ARNs](#) in the *AWS Identity and Access Management User Guide*.

For more information, see [AWS IoT policies](#) and [AWS IoT policy actions](#) in the *AWS IoT Core Developer Guide*.

Note

AWS IoT Core enables you to attach AWS IoT policies to thing groups to define permissions for groups of devices. Thing group policies don't allow access to AWS IoT Greengrass data plane operations. To allow a thing access to an AWS IoT Greengrass data plane operation, add the permission to an AWS IoT policy that you attach to the thing's certificate.

AWS IoT Greengrass policy actions

Greengrass Core Actions

AWS IoT Greengrass defines the following policy actions that Greengrass core devices can use in AWS IoT policies:

`greengrass:AssumeRoleForGroup`

Permission for a Greengrass core device to retrieve credentials using the Token Exchange Service (TES) system Lambda function. The permissions that are tied to the retrieved credentials are based on the policy that's attached to the configured group role.

This permission is checked when a Greengrass core device attempts to retrieve credentials (assuming the credentials are not cached locally).

`greengrass:CreateCertificate`

Permission for a Greengrass core device to create its own server certificate.

This permission is checked when a Greengrass core device creates a certificate. Greengrass core devices attempt to create a server certificate upon first run, when the core's connectivity information changes, and on designated rotation periods.

`greengrass:GetConnectivityInfo`

Permission for a Greengrass core device to retrieve its own connectivity information.

This permission is checked when a Greengrass core device attempts to retrieve its connectivity information from AWS IoT Core.

`greengrass:GetDeployment`

Permission for a Greengrass core device to retrieve deployments.

This permission is checked when a Greengrass core device attempts to retrieve deployments and deployment statuses from the cloud.

`greengrass:GetDeploymentArtifacts`

Permission for a Greengrass core device to retrieve deployment artifacts such as group information or Lambda functions.

This permission is checked when a Greengrass core device receives a deployment and then attempts to retrieve deployment artifacts.

`greengrass:UpdateConnectivityInfo`

Permission for a Greengrass core device to update its own connectivity information with IP or hostname information.

This permission is checked when a Greengrass core device attempts to update its connectivity information in the cloud.

`greengrass:UpdateCoreDeploymentStatus`

Permission for a Greengrass core device to update the status of a deployment.

This permission is checked when a Greengrass core device receives a deployment and then attempts to update the deployment status.

Greengrass Device Actions

AWS IoT Greengrass defines the following policy action that client devices can use in AWS IoT policies:

`greengrass:Discover`

Permission for a client device to use the [Discovery API](#) to retrieve its group's core connectivity information and group certificate authority.

This permission is checked when a client device calls the Discovery API with TLS mutual authentication.

Minimal AWS IoT policy for the AWS IoT Greengrass core device

The following example policy includes the minimum set of actions required to support basic Greengrass functionality for your core device.

- The policy lists the MQTT topics and topic filters that the core device can publish messages to, subscribe to, and receive messages on, including topics used for shadow state. To support message exchange between AWS IoT Core, Lambda functions, connectors, and client devices in the Greengrass group, specify the topics and topic filters that you want to allow. For more information, see [Publish/Subscribe policy examples](#) in the *AWS IoT Core Developer Guide*.
- The policy includes a section that allows AWS IoT Core to get, update, and delete the core device's shadow. To allow shadow sync for client devices in the Greengrass group, specify the target Amazon Resource Names (ARNs) in the Resource list (for example, `arn:aws:iot:region:account-id:thing/device-name`).
- The use of [thing policy variables](#) (`iot:Connection.Thing.*`) in the AWS IoT policy for a core device is not supported. The core uses the same device certificate to make [multiple connections](#) to AWS IoT Core but the client ID in a connection might not be an exact match of the core thing name.
- For the `greengrass:UpdateCoreDeploymentStatus` permission, the final segment in the Resource ARN is the URL-encoded ARN of the core device.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "iot:Connect"
      ],
      "Resource": [
        "arn:aws:iot:region:account-id:client/core-name-*"
      ]
    },
    {
```

```

        "Effect": "Allow",
        "Action": [
            "iot:Publish",
            "iot:Receive"
        ],
        "Resource": [
            "arn:aws:iot:region:account-id:topic/$aws/things/core-name-*"
        ]
    },
    {
        "Effect": "Allow",
        "Action": [
            "iot:Subscribe"
        ],
        "Resource": [
            "arn:aws:iot:region:account-id:topicfilter/$aws/things/core-name-*"
        ]
    },
    {
        "Effect": "Allow",
        "Action": [
            "iot:GetThingShadow",
            "iot:UpdateThingShadow",
            "iot:DeleteThingShadow"
        ],
        "Resource": [
            "arn:aws:iot:region:account-id:thing/core-name-*"
        ]
    },
    {
        "Effect": "Allow",
        "Action": [
            "greengrass:AssumeRoleForGroup",
            "greengrass:CreateCertificate"
        ],
        "Resource": [
            "*"
        ]
    },
    {
        "Effect": "Allow",
        "Action": [
            "greengrass:GetDeployment"
        ],

```

```

        "Resource": [
            "arn:aws:greengrass:region:account-id:/greengrass/groups/group-id/
deployments/*"
        ],
    },
    {
        "Effect": "Allow",
        "Action": [
            "greengrass:GetDeploymentArtifacts"
        ],
        "Resource": [
            "arn:aws:greengrass:region:account-id:/greengrass/groups/group-id/
deployments/*"
        ],
    },
    {
        "Effect": "Allow",
        "Action": [
            "greengrass:UpdateCoreDeploymentStatus"
        ],
        "Resource": [
            "arn:aws:greengrass:region:account-id:/greengrass/groups/group-id/
deployments/*/cores/arn%3Aaws%3Aiot%3Aregion%3Aaccount-id%3Athing%2Fcore-name"
        ],
    },
    {
        "Effect": "Allow",
        "Action": [
            "greengrass:GetConnectivityInfo",
            "greengrass:UpdateConnectivityInfo"
        ],
        "Resource": [
            "arn:aws:iot:region:account-id:thing/core-name-"
        ],
    }
]
}

```

Note

AWS IoT policies for client devices typically require similar permissions for `iot:Connect`, `iot:Publish`, `iot:Receive`, and `iot:Subscribe` actions.

To allow a client device to automatically detect connectivity information for the cores in the Greengrass groups that the device belongs to, the AWS IoT policy for a client device must include the `greengrass:Discover` action. In the Resource section, specify the ARN of the client device, not the ARN of the Greengrass core device. For example:

```
{
  "Effect": "Allow",
  "Action": [
    "greengrass:Discover"
  ],
  "Resource": [
    "arn:aws:iot:region:account-id:thing/device-name"
  ]
}
```

The AWS IoT policy for client devices doesn't typically require permissions for `iot:GetThingShadow`, `iot:UpdateThingShadow`, or `iot:DeleteThingShadow` actions, because the Greengrass core handles shadow sync operations for client devices. In this case, make sure that the Resource section for shadow actions in the core's AWS IoT policy includes the ARNs of the client devices.

In the AWS IoT console, you can view and edit the policy that's attached to your core's certificate.

1. In the navigation pane, under **Manage**, expand **All devices**, and then choose **Things**.
2. Choose your core.
3. On your core's configuration page, choose the **Certificates** tab.
4. In the **Certificates** tab, choose your certificate.
5. On the certificate's configuration page, choose **Policies**, and then choose the policy.

If you want to edit the policy, choose **Edit active version**.

6. Review the policy and add, remove, or edit permissions as needed.
7. To set a new policy version as the active version, under **Policy version status**, select **Set the edited version as the active version for this policy**.
8. Choose **Save as new version**.

Identity and access management for AWS IoT Greengrass

AWS Identity and Access Management (IAM) is an AWS service that helps an administrator securely control access to AWS resources. IAM administrators control who can be *authenticated* (signed in) and *authorized* (have permissions) to use AWS IoT Greengrass resources. IAM is an AWS service that you can use with no additional charge.

Note

This topic describes IAM concepts and features. For information about IAM features supported by AWS IoT Greengrass, see [the section called “How AWS IoT Greengrass works with IAM”](#).

Audience

How you use AWS Identity and Access Management (IAM) differs, depending on the work that you do in AWS IoT Greengrass.

Service user – If you use the AWS IoT Greengrass service to do your job, then your administrator provides you with the credentials and permissions that you need. As you use more AWS IoT Greengrass features to do your work, you might need additional permissions. Understanding how access is managed can help you request the right permissions from your administrator. If you cannot access a feature in AWS IoT Greengrass, see [Troubleshooting identity and access issues for AWS IoT Greengrass](#).

Service administrator – If you're in charge of AWS IoT Greengrass resources at your company, you probably have full access to AWS IoT Greengrass. It's your job to determine which AWS IoT Greengrass features and resources your service users should access. You must then submit requests to your IAM administrator to change the permissions of your service users. Review the information on this page to understand the basic concepts of IAM. To learn more about how your company can use IAM with AWS IoT Greengrass, see [How AWS IoT Greengrass works with IAM](#).

IAM administrator – If you're an IAM administrator, you might want to learn details about how you can write policies to manage access to AWS IoT Greengrass. To view example AWS IoT Greengrass identity-based policies that you can use in IAM, see [Identity-based policy examples for AWS IoT Greengrass](#).

Authenticating with identities

Authentication is how you sign in to AWS using your identity credentials. You must be *authenticated* (signed in to AWS) as the AWS account root user, as an IAM user, or by assuming an IAM role.

You can sign in to AWS as a federated identity by using credentials provided through an identity source. AWS IAM Identity Center (IAM Identity Center) users, your company's single sign-on authentication, and your Google or Facebook credentials are examples of federated identities. When you sign in as a federated identity, your administrator previously set up identity federation using IAM roles. When you access AWS by using federation, you are indirectly assuming a role.

Depending on the type of user you are, you can sign in to the AWS Management Console or the AWS access portal. For more information about signing in to AWS, see [How to sign in to your AWS account](#) in the *AWS Sign-In User Guide*.

If you access AWS programmatically, AWS provides a software development kit (SDK) and a command line interface (CLI) to cryptographically sign your requests by using your credentials. If you don't use AWS tools, you must sign requests yourself. For more information about using the recommended method to sign requests yourself, see [AWS Signature Version 4 for API requests](#) in the *IAM User Guide*.

Regardless of the authentication method that you use, you might be required to provide additional security information. For example, AWS recommends that you use multi-factor authentication (MFA) to increase the security of your account. To learn more, see [Multi-factor authentication](#) in the *AWS IAM Identity Center User Guide* and [AWS Multi-factor authentication in IAM](#) in the *IAM User Guide*.

AWS account root user

When you create an AWS account, you begin with one sign-in identity that has complete access to all AWS services and resources in the account. This identity is called the AWS account *root user* and is accessed by signing in with the email address and password that you used to create the account. We strongly recommend that you don't use the root user for your everyday tasks. Safeguard your root user credentials and use them to perform the tasks that only the root user can perform. For the complete list of tasks that require you to sign in as the root user, see [Tasks that require root user credentials](#) in the *IAM User Guide*.

IAM users and groups

An [IAM user](#) is an identity within your AWS account that has specific permissions for a single person or application. Where possible, we recommend relying on temporary credentials instead of creating IAM users who have long-term credentials such as passwords and access keys. However, if you have specific use cases that require long-term credentials with IAM users, we recommend that you rotate access keys. For more information, see [Rotate access keys regularly for use cases that require long-term credentials](#) in the *IAM User Guide*.

An [IAM group](#) is an identity that specifies a collection of IAM users. You can't sign in as a group. You can use groups to specify permissions for multiple users at a time. Groups make permissions easier to manage for large sets of users. For example, you could have a group named *IAMAdmins* and give that group permissions to administer IAM resources.

Users are different from roles. A user is uniquely associated with one person or application, but a role is intended to be assumable by anyone who needs it. Users have permanent long-term credentials, but roles provide temporary credentials. To learn more, see [Use cases for IAM users](#) in the *IAM User Guide*.

IAM roles

An [IAM role](#) is an identity within your AWS account that has specific permissions. It is similar to an IAM user, but is not associated with a specific person. To temporarily assume an IAM role in the AWS Management Console, you can [switch from a user to an IAM role \(console\)](#). You can assume a role by calling an AWS CLI or AWS API operation or by using a custom URL. For more information about methods for using roles, see [Methods to assume a role](#) in the *IAM User Guide*.

IAM roles with temporary credentials are useful in the following situations:

- **Federated user access** – To assign permissions to a federated identity, you create a role and define permissions for the role. When a federated identity authenticates, the identity is associated with the role and is granted the permissions that are defined by the role. For information about roles for federation, see [Create a role for a third-party identity provider \(federation\)](#) in the *IAM User Guide*. If you use IAM Identity Center, you configure a permission set. To control what your identities can access after they authenticate, IAM Identity Center correlates the permission set to a role in IAM. For information about permissions sets, see [Permission sets](#) in the *AWS IAM Identity Center User Guide*.
- **Temporary IAM user permissions** – An IAM user or role can assume an IAM role to temporarily take on different permissions for a specific task.

- **Cross-account access** – You can use an IAM role to allow someone (a trusted principal) in a different account to access resources in your account. Roles are the primary way to grant cross-account access. However, with some AWS services, you can attach a policy directly to a resource (instead of using a role as a proxy). To learn the difference between roles and resource-based policies for cross-account access, see [Cross account resource access in IAM](#) in the *IAM User Guide*.
- **Cross-service access** – Some AWS services use features in other AWS services. For example, when you make a call in a service, it's common for that service to run applications in Amazon EC2 or store objects in Amazon S3. A service might do this using the calling principal's permissions, using a service role, or using a service-linked role.
- **Forward access sessions (FAS)** – When you use an IAM user or role to perform actions in AWS, you are considered a principal. When you use some services, you might perform an action that then initiates another action in a different service. FAS uses the permissions of the principal calling an AWS service, combined with the requesting AWS service to make requests to downstream services. FAS requests are only made when a service receives a request that requires interactions with other AWS services or resources to complete. In this case, you must have permissions to perform both actions. For policy details when making FAS requests, see [Forward access sessions](#).
- **Service role** – A service role is an [IAM role](#) that a service assumes to perform actions on your behalf. An IAM administrator can create, modify, and delete a service role from within IAM. For more information, see [Create a role to delegate permissions to an AWS service](#) in the *IAM User Guide*.
- **Service-linked role** – A service-linked role is a type of service role that is linked to an AWS service. The service can assume the role to perform an action on your behalf. Service-linked roles appear in your AWS account and are owned by the service. An IAM administrator can view, but not edit the permissions for service-linked roles.
- **Applications running on Amazon EC2** – You can use an IAM role to manage temporary credentials for applications that are running on an EC2 instance and making AWS CLI or AWS API requests. This is preferable to storing access keys within the EC2 instance. To assign an AWS role to an EC2 instance and make it available to all of its applications, you create an instance profile that is attached to the instance. An instance profile contains the role and enables programs that are running on the EC2 instance to get temporary credentials. For more information, see [Use an IAM role to grant permissions to applications running on Amazon EC2 instances](#) in the *IAM User Guide*.

Managing access using policies

You control access in AWS by creating policies and attaching them to AWS identities or resources. A policy is an object in AWS that, when associated with an identity or resource, defines their permissions. AWS evaluates these policies when a principal (user, root user, or role session) makes a request. Permissions in the policies determine whether the request is allowed or denied. Most policies are stored in AWS as JSON documents. For more information about the structure and contents of JSON policy documents, see [Overview of JSON policies](#) in the *IAM User Guide*.

Administrators can use AWS JSON policies to specify who has access to what. That is, which **principal** can perform **actions** on what **resources**, and under what **conditions**.

By default, users and roles have no permissions. To grant users permission to perform actions on the resources that they need, an IAM administrator can create IAM policies. The administrator can then add the IAM policies to roles, and users can assume the roles.

IAM policies define permissions for an action regardless of the method that you use to perform the operation. For example, suppose that you have a policy that allows the `iam:GetRole` action. A user with that policy can get role information from the AWS Management Console, the AWS CLI, or the AWS API.

Identity-based policies

Identity-based policies are JSON permissions policy documents that you can attach to an identity, such as an IAM user, group of users, or role. These policies control what actions users and roles can perform, on which resources, and under what conditions. To learn how to create an identity-based policy, see [Define custom IAM permissions with customer managed policies](#) in the *IAM User Guide*.

Identity-based policies can be further categorized as *inline policies* or *managed policies*. Inline policies are embedded directly into a single user, group, or role. Managed policies are standalone policies that you can attach to multiple users, groups, and roles in your AWS account. Managed policies include AWS managed policies and customer managed policies. To learn how to choose between a managed policy or an inline policy, see [Choose between managed policies and inline policies](#) in the *IAM User Guide*.

Resource-based policies

Resource-based policies are JSON policy documents that you attach to a resource. Examples of resource-based policies are IAM *role trust policies* and Amazon S3 *bucket policies*. In services that

support resource-based policies, service administrators can use them to control access to a specific resource. For the resource where the policy is attached, the policy defines what actions a specified principal can perform on that resource and under what conditions. You must [specify a principal](#) in a resource-based policy. Principals can include accounts, users, roles, federated users, or AWS services.

Resource-based policies are inline policies that are located in that service. You can't use AWS managed policies from IAM in a resource-based policy.

Access control lists (ACLs)

Access control lists (ACLs) control which principals (account members, users, or roles) have permissions to access a resource. ACLs are similar to resource-based policies, although they do not use the JSON policy document format.

Amazon S3, AWS WAF, and Amazon VPC are examples of services that support ACLs. To learn more about ACLs, see [Access control list \(ACL\) overview](#) in the *Amazon Simple Storage Service Developer Guide*.

Other policy types

AWS supports additional, less-common policy types. These policy types can set the maximum permissions granted to you by the more common policy types.

- **Permissions boundaries** – A permissions boundary is an advanced feature in which you set the maximum permissions that an identity-based policy can grant to an IAM entity (IAM user or role). You can set a permissions boundary for an entity. The resulting permissions are the intersection of an entity's identity-based policies and its permissions boundaries. Resource-based policies that specify the user or role in the `Principal` field are not limited by the permissions boundary. An explicit deny in any of these policies overrides the allow. For more information about permissions boundaries, see [Permissions boundaries for IAM entities](#) in the *IAM User Guide*.
- **Service control policies (SCPs)** – SCPs are JSON policies that specify the maximum permissions for an organization or organizational unit (OU) in AWS Organizations. AWS Organizations is a service for grouping and centrally managing multiple AWS accounts that your business owns. If you enable all features in an organization, then you can apply service control policies (SCPs) to any or all of your accounts. The SCP limits permissions for entities in member accounts, including each AWS account root user. For more information about Organizations and SCPs, see [Service control policies](#) in the *AWS Organizations User Guide*.

- **Resource control policies (RCPs)** – RCPs are JSON policies that you can use to set the maximum available permissions for resources in your accounts without updating the IAM policies attached to each resource that you own. The RCP limits permissions for resources in member accounts and can impact the effective permissions for identities, including the AWS account root user, regardless of whether they belong to your organization. For more information about Organizations and RCPs, including a list of AWS services that support RCPs, see [Resource control policies \(RCPs\)](#) in the *AWS Organizations User Guide*.
- **Session policies** – Session policies are advanced policies that you pass as a parameter when you programmatically create a temporary session for a role or federated user. The resulting session's permissions are the intersection of the user or role's identity-based policies and the session policies. Permissions can also come from a resource-based policy. An explicit deny in any of these policies overrides the allow. For more information, see [Session policies](#) in the *IAM User Guide*.

Multiple policy types

When multiple types of policies apply to a request, the resulting permissions are more complicated to understand. To learn how AWS determines whether to allow a request when multiple policy types are involved, see [Policy evaluation logic](#) in the *IAM User Guide*.

See also

- [the section called “How AWS IoT Greengrass works with IAM”](#)
- [the section called “Identity-based policy examples”](#)
- [the section called “Troubleshooting identity and access issues”](#)

How AWS IoT Greengrass works with IAM

Before you use IAM to manage access to AWS IoT Greengrass, you should understand the IAM features that you can use with AWS IoT Greengrass.

IAM feature	Supported by Greengrass?
Identity-based policies with resource-level permissions	Yes
Resource-based policies	No

IAM feature	Supported by Greengrass?
Access control lists (ACLs)	No
Tags-based authorization	No
Temporary credentials	Yes
Service-linked roles	No
Service roles	Yes

For a high-level view of how other AWS services work with IAM, see [AWS services that work with IAM](#) in the *IAM User Guide*.

Identity-based policies for AWS IoT Greengrass

With IAM identity-based policies, you can specify allowed or denied actions and resources and the conditions under which actions are allowed or denied. AWS IoT Greengrass supports specific actions, resources, and condition keys. To learn about all of the elements that you use in a policy, see [IAM JSON policy elements reference](#) in the *IAM User Guide*.

Actions

Administrators can use AWS JSON policies to specify who has access to what. That is, which **principal** can perform **actions** on what **resources**, and under what **conditions**.

The Action element of a JSON policy describes the actions that you can use to allow or deny access in a policy. Policy actions usually have the same name as the associated AWS API operation. There are some exceptions, such as *permission-only actions* that don't have a matching API operation. There are also some operations that require multiple actions in a policy. These additional actions are called *dependent actions*.

Include actions in a policy to grant permissions to perform the associated operation.

Policy actions for AWS IoT Greengrass use the `greengrass:` prefix before the action. For example, to allow someone to use the `ListGroups` API operation to list the groups in their AWS account, you include the `greengrass:ListGroups` action in their policy. Policy statements must include either an Action or NotAction element. AWS IoT Greengrass defines its own set of actions that describe tasks that you can perform with this service.

To specify multiple actions in a single statement, list them between brackets ([]) and separate them with commas, as follows:

```
"Action": [  
    "greengrass:action1",  
    "greengrass:action2",  
    "greengrass:action3"  
]
```

You can use wildcards (*) to specify multiple actions. For example, to specify all actions that begin with the word `List`, include the following action:

```
"Action": "greengrass:List*"
```

Note

We recommend that you avoid the use of wildcards to specify all available actions for a service. As a best practice, you should grant least privilege and narrowly scope permissions in a policy. For more information, see [the section called “Grant minimum possible permissions”](#).

For the complete list of AWS IoT Greengrass actions, see [Actions Defined by AWS IoT Greengrass](#) in the *IAM User Guide*.

Resources

Administrators can use AWS JSON policies to specify who has access to what. That is, which **principal** can perform **actions** on what **resources**, and under what **conditions**.

The Resource JSON policy element specifies the object or objects to which the action applies. Statements must include either a Resource or a NotResource element. As a best practice, specify a resource using its [Amazon Resource Name \(ARN\)](#). You can do this for actions that support a specific resource type, known as *resource-level permissions*.

For actions that don't support resource-level permissions, such as listing operations, use a wildcard (*) to indicate that the statement applies to all resources.

```
"Resource": "*"
```

The following table contains the AWS IoT Greengrass resource ARNs that can be used in the Resource element of a policy statement. For a mapping of supported resource-level permissions for AWS IoT Greengrass actions, see [Actions Defined by AWS IoT Greengrass](#) in the *IAM User Guide*.

Resource	ARN
Group	arn:\${Partition}:greengrass:\${Region}:\${Account}:/greengrass/groups/\${GroupId}
GroupVersion	arn:\${Partition}:greengrass:\${Region}:\${Account}:/greengrass/groups/\${GroupId}/versions/\${VersionId}
CertificateAuthority	arn:\${Partition}:greengrass:\${Region}:\${Account}:/greengrass/groups/\${GroupId}/certificateauthorities/\${CertificateAuthorityId}
Deployment	arn:\${Partition}:greengrass:\${Region}:\${Account}:/greengrass/groups/\${GroupId}/deployments/\${DeploymentId}
BulkDeployment	arn:\${Partition}:greengrass:\${Region}:\${Account}:/greengrass/bulk/deployments/\${BulkDeploymentId}
Connector Definition	arn:\${Partition}:greengrass:\${Region}:\${Account}:/greengrass/definition/connectors/\${ConnectorDefinitionId}
Connector DefinitionVersion	arn:\${Partition}:greengrass:\${Region}:\${Account}:/greengrass/definition/connectors/\${ConnectorDefinitionId}/versions/\${VersionId}
CoreDefinition	arn:\${Partition}:greengrass:\${Region}:\${Account}:/greengrass/definition/cores/\${CoreDefinitionId}
CoreDefinitionVersion	arn:\${Partition}:greengrass:\${Region}:\${Account}:/greengrass/definition/cores/\${CoreDefinitionId}/versions/\${VersionId}

Resource	ARN
<u>DeviceDefinition</u>	arn:\${Partition}:greengrass:\${Region}:\${Account}:/greengrass/definition/devices/\${DeviceDefinitionId}
<u>DeviceDefinitionVersion</u>	arn:\${Partition}:greengrass:\${Region}:\${Account}:/greengrass/definition/devices/\${DeviceDefinitionId}/versions/\${VersionId}
<u>FunctionDefinition</u>	arn:\${Partition}:greengrass:\${Region}:\${Account}:/greengrass/definition/functions/\${FunctionDefinitionId}
<u>FunctionDefinitionVersion</u>	arn:\${Partition}:greengrass:\${Region}:\${Account}:/greengrass/definition/functions/\${FunctionDefinitionId}/versions/\${VersionId}
<u>LoggerDefinition</u>	arn:\${Partition}:greengrass:\${Region}:\${Account}:/greengrass/definition/loggers/\${LoggerDefinitionId}
<u>LoggerDefinitionVersion</u>	arn:\${Partition}:greengrass:\${Region}:\${Account}:/greengrass/definition/loggers/\${LoggerDefinitionId}/versions/\${VersionId}
<u>ResourceDefinition</u>	arn:\${Partition}:greengrass:\${Region}:\${Account}:/greengrass/definition/resources/\${ResourceDefinitionId}
<u>ResourceDefinitionVersion</u>	arn:\${Partition}:greengrass:\${Region}:\${Account}:/greengrass/definition/resources/\${ResourceDefinitionId}/versions/\${VersionId}
<u>SubscriptionDefinition</u>	arn:\${Partition}:greengrass:\${Region}:\${Account}:/greengrass/definition/subscriptions/\${SubscriptionDefinitionId}

Resource	ARN
SubscriptionDefinitionVersion	arn:\${Partition}:greengrass:\${Region}:\${Account}:/ greengrass/definition/subscriptions/\${SubscriptionDefinitionId}/versions/\${VersionId}
ConnectivityInfo	arn:\${Partition}:greengrass:\${Region}:\${Account}:/ greengrass/things/\${ThingName}/connectivityInfo

The following example Resource element specifies the ARN of a group in the US West (Oregon) Region in the AWS account 123456789012:

```
"Resource": "arn:aws:greengrass:us-west-2:123456789012:/greengrass/groups/a1b2c3d4-5678-90ab-cdef-EXAMPLE11111"
```

Or, to specify all groups that belong to an AWS account in a specific AWS Region, use the wildcard in place of the group ID:

```
"Resource": "arn:aws:greengrass:us-west-2:123456789012:/greengrass/groups/*"
```

Some AWS IoT Greengrass actions (for example, some list operations), cannot be performed on a specific resource. In those cases, you must use the wildcard alone.

```
"Resource": ""
```

To specify multiple resource ARNs in a statement, list them between brackets ([]) and separate them with commas, as follows:

```
"Resource": [
  "resource-arn1",
  "resource-arn2",
  "resource-arn3"
]
```

For more information about ARN formats, see [Amazon Resource Names \(ARNs\) and AWS service namespaces](#) in the *Amazon Web Services General Reference*.

Condition keys

Administrators can use AWS JSON policies to specify who has access to what. That is, which **principal** can perform **actions** on what **resources**, and under what **conditions**.

The Condition element (or Condition *block*) lets you specify conditions in which a statement is in effect. The Condition element is optional. You can create conditional expressions that use [condition operators](#), such as equals or less than, to match the condition in the policy with values in the request.

If you specify multiple Condition elements in a statement, or multiple keys in a single Condition element, AWS evaluates them using a logical AND operation. If you specify multiple values for a single condition key, AWS evaluates the condition using a logical OR operation. All of the conditions must be met before the statement's permissions are granted.

You can also use placeholder variables when you specify conditions. For example, you can grant an IAM user permission to access a resource only if it is tagged with their IAM user name. For more information, see [IAM policy elements: variables and tags](#) in the *IAM User Guide*.

AWS supports global condition keys and service-specific condition keys. To see all AWS global condition keys, see [AWS global condition context keys](#) in the *IAM User Guide*.

AWS IoT Greengrass supports the following global condition keys.

Key	Description
aws:CurrentTime	Filters access by checking date/time conditions for the current date and time.
aws:EpochTime	Filters access by checking date/time conditions for the current date and time in epoch or Unix time.
aws:MultiFactorAuthAge	Filters access by checking how long ago (in seconds) the security credentials validated by multi-factor authentication (MFA) in the request were issued using MFA.
aws:MultiFactorAuthPresent	Filters access by checking whether multi-factor authentication (MFA) was used to validate the temporary security credentials that made the current request.

Key	Description
<code>aws:RequestTag/\${TagKey}</code>	Filters create requests based on the allowed set of values for each of the mandatory tags.
<code>aws:ResourceTag/\${TagKey}</code>	Filters actions based on the tag value associated with the resource.
<code>aws:SecureTransport</code>	Filters access by checking whether the request was sent using SSL.
<code>aws:TagKeys</code>	Filters create requests based on the presence of mandatory tags in the request.
<code>aws:UserAgent</code>	Filters access by the requester's client application.

For more information, see [AWS global condition context keys](#) in the *IAM User Guide*.

Examples

To view examples of AWS IoT Greengrass identity-based policies, see [the section called “Identity-based policy examples”](#).

Resource-based policies for AWS IoT Greengrass

AWS IoT Greengrass does not support [resource-based policies](#).

Access control lists (ACLs)

AWS IoT Greengrass does not support [ACLs](#).

Authorization based on AWS IoT Greengrass tags

AWS IoT Greengrass does not support tags-based authorization.

IAM roles for AWS IoT Greengrass

An [IAM role](#) is an entity within your AWS account that has specific permissions.

Using temporary credentials with AWS IoT Greengrass

Temporary credentials are used to sign in with federation, assume an IAM role, or to assume a cross-account role. You obtain temporary security credentials by calling AWS STS API operations such as [AssumeRole](#) or [GetFederationToken](#).

On the Greengrass core, temporary credentials for the [group role](#) are made available to user-defined Lambda functions and connectors. If your Lambda functions use the AWS SDK, you don't need to add logic to obtain the credentials because the AWS SDK does this for you.

Service-linked roles

AWS IoT Greengrass does not support [service-linked roles](#).

Service roles

This feature allows a service to assume a [service role](#) on your behalf. This role allows the service to access resources in other services to complete an action on your behalf. Service roles appear in your IAM account and are owned by the account. This means that an IAM administrator can change the permissions for this role. However, doing so might break the functionality of the service.

AWS IoT Greengrass uses a service role to access some of your AWS resources on your behalf. For more information, see [the section called “Greengrass service role”](#).

Choosing an IAM role in the AWS IoT Greengrass console

In the AWS IoT Greengrass console, you might need to choose a Greengrass service role or a Greengrass group role from a list of IAM roles in your account.

- The Greengrass service role allows AWS IoT Greengrass to access your AWS resources in other services on your behalf. Typically, you don't need to choose the service role because the console can create and configure it for you. For more information, see [the section called “Greengrass service role”](#).
- The Greengrass group role is used to allow Greengrass Lambda functions and connectors in the group to access your AWS resources. It can also give AWS IoT Greengrass permissions to export streams to AWS services and write CloudWatch logs. For more information, see [the section called “Greengrass group role”](#).

Greengrass service role

The Greengrass service role is an AWS Identity and Access Management (IAM) service role that authorizes AWS IoT Greengrass to access resources from AWS services on your behalf. This makes it possible for AWS IoT Greengrass to perform essential tasks, such as retrieving your AWS Lambda functions and managing AWS IoT shadows.

To allow AWS IoT Greengrass to access your resources, the Greengrass service role must be associated with your AWS account and specify AWS IoT Greengrass as a trusted entity. The role must include the [AWSGreengrassResourceAccessRolePolicy](#) managed policy or a custom policy that defines equivalent permissions for the AWS IoT Greengrass features that you use. This policy is maintained by AWS and defines the set of permissions that AWS IoT Greengrass uses to access your AWS resources.

You can reuse the same Greengrass service role across AWS Regions, but you must associate it with your account in every AWS Region where you use AWS IoT Greengrass. Group deployment fails if the service role doesn't exist in the current AWS account and Region.

The following sections describe how to create and manage the Greengrass service role in the AWS Management Console or AWS CLI.

- [Manage the service role \(console\)](#)
- [Manage the service role \(CLI\)](#)

Note

In addition to the service role that authorizes service-level access, you can assign a *group role* to an AWS IoT Greengrass group. The group role is a separate IAM role that controls how Greengrass Lambda functions and connectors in the group can access AWS services.

Managing the Greengrass service role (console)

The AWS IoT console makes it easy to manage your Greengrass service role. For example, when you create or deploy a Greengrass group, the console checks whether your AWS account is attached to a Greengrass service role in the AWS Region that's currently selected in the console. If not, the console can create and configure a service role for you. For more information, see [the section called "Create the Greengrass service role"](#).

You can use the AWS IoT console for the following role management tasks:

- [Find your Greengrass service role](#)
- [Create the Greengrass service role](#)
- [Change the Greengrass service role](#)
- [Detach the Greengrass service role](#)

 Note

The user who is signed in to the console must have permissions to view, create, or change the service role.

Find your Greengrass service role (console)

Use the following steps to find the service role that AWS IoT Greengrass is using in the current AWS Region.

1. From the [AWS IoT console](#) navigation pane, choose **Settings**.
2. Scroll to the **Greengrass service role** section to see your service role and its policies.

If you don't see a service role, you can let the console create or configure one for you. For more information, see [Create the Greengrass service role](#).

Create the Greengrass service role (console)

The console can create and configure a default Greengrass service role for you. This role has the following properties.

Property	Value
Name	Greengrass_ServiceRole
Trusted entity	AWS service: greengrass

Property	Value
Policy	AWSGreengrassResourceAccessRolePolicy

Note

If [Greengrass device setup](#) creates the service role, the role name is `GreengrassServiceRole_random-string`.

When you create or deploy a Greengrass group from the AWS IoT console, the console checks whether a Greengrass service role is associated with your AWS account in the AWS Region that's currently selected in the console. If not, the console prompts you to allow AWS IoT Greengrass to read and write to AWS services on your behalf.

If you grant permission, the console checks whether a role named `Greengrass_ServiceRole` exists in your AWS account.

- If the role exists, the console attaches the service role to your AWS account in the current AWS Region.
- If the role doesn't exist, the console creates a default Greengrass service role and attaches it to your AWS account in the current AWS Region.

Note

If you want to create a service role with custom role policies, use the IAM console to create or modify the role. For more information, see [Creating a role to delegate permissions to an AWS service](#) or [Modifying a role](#) in the *IAM User Guide*. Make sure that the role grants permissions that are equivalent to the `AWSGreengrassResourceAccessRolePolicy` managed policy for the features and resources that you use. We recommend that you also include the `aws:SourceArn` and `aws:SourceAccount` global condition context keys in your trust policy to help prevent the *confused deputy* security problem. The condition context keys restrict access to allow only those requests that come from the specified account and Greengrass workspace. For more information about the confused deputy problem, see [Cross-service confused deputy prevention](#).

If you create a service role, return to the AWS IoT console and attach the role to the group. You can do this under **Greengrass service role** on the group's **Settings** page.

Change the Greengrass service role (console)

Use the following procedure to choose a different Greengrass service role to attach to your AWS account in the AWS Region currently selected in the console.

1. From the [AWS IoT console](#) navigation pane, choose **Settings**.
2. Under **Greengrass service role**, choose **Change role**.

The **Update Greengrass service role** dialog box opens and shows the IAM roles in your AWS account that define AWS IoT Greengrass as a trusted entity.

3. Choose the Greengrass service role to attach.
4. Choose **Attach role**.

Note

To allow the console to create a default Greengrass service role for you, choose **Create role for me** instead of choosing a role from the list. The **Create role for me** link does not appear if a role named `Greengrass_ServiceRole` is in your AWS account.

Detach the Greengrass service role (console)

Use the following procedure to detach the Greengrass service role from your AWS account in the AWS Region currently selected in the console. This revokes permissions for AWS IoT Greengrass to access AWS services in the current AWS Region.

Important

Detaching the service role might interrupt active operations.

1. From the [AWS IoT console](#) navigation pane, choose **Settings**.
2. Under **Greengrass service role**, choose **Detach role**.
3. In the confirmation dialog box, choose **Detach**.

Note

If you no longer need the role, you can delete it in the IAM console. For more information, see [Deleting roles or instance profiles](#) in the *IAM User Guide*.

Other roles might allow AWS IoT Greengrass to access your resources. To find all roles that allow AWS IoT Greengrass to assume permissions on your behalf, in the IAM console, on the **Roles** page, look for roles that include **AWS service: greengrass** in the **Trusted entities** column.

Managing the Greengrass service role (CLI)

In the following procedures, we assume that the AWS CLI is installed and configured to use your AWS account ID. For more information, see [Installing the AWS command line interface](#) and [Configuring the AWS CLI](#) in the *AWS Command Line Interface User Guide*.

You can use the AWS CLI for the following role management tasks:

- [Get your Greengrass service role](#)
- [Create the Greengrass service role](#)
- [Remove the Greengrass service role](#)

Get the Greengrass service role (CLI)

Use the following procedure to find out if a Greengrass service role is associated with your AWS account in an AWS Region.

- Get the service role. Replace *region* with your AWS Region (for example, us-west-2).

```
aws Greengrass get-service-role-for-account --region region
```

If a Greengrass service role is already associated with your account, the following role metadata is returned.

```
{
  "AssociatedAt": "timestamp",
  "RoleArn": "arn:aws:iam::account-id:role/path/role-name"
}
```

If no role metadata is returned, then you must create the service role (if it doesn't exist) and associate it with your account in the AWS Region.

Create the Greengrass service role (CLI)

Use the following steps to create a role and associate it with your AWS account.

To create the service role using IAM

1. Create the role with a trust policy that allows AWS IoT Greengrass to assume the role. This example creates a role named `Greengrass_ServiceRole`, but you can use a different name. We recommend that you also include the `aws:SourceArn` and `aws:SourceAccount` global condition context keys in your trust policy to help prevent the *confused deputy* security problem. The condition context keys restrict access to allow only those requests that come from the specified account and Greengrass workspace. For more information about the confused deputy problem, see [Cross-service confused deputy prevention](#).

Linux, macOS, or Unix

```
aws iam create-role --role-name Greengrass_ServiceRole --assume-role-policy-
document '{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Principal": {
        "Service": "greengrass.amazonaws.com"
      },
      "Action": "sts:AssumeRole",
      "Condition": {
```

```

    "StringEquals": {
      "aws:SourceAccount": "account-id"
    },
    "ArnLike": {
      "aws:SourceArn": "arn:aws:greengrass:region:account-id:*"
    }
  }
}
]
}'

```

Windows command prompt

```

aws iam create-role --role-name Greengrass_ServiceRole --assume-role-
policy-document "{\"Version\":\"2012-10-17\",\"Statement\": [{\"Effect
\": \"Allow\", \"Principal\": {\"Service\": \"greengrass.amazonaws.com\"},
\"Action\": \"sts:AssumeRole\", \"Condition\": {\"ArnLike\": {\"aws:SourceArn
\": \"arn:aws:greengrass:region:account-id:*\"}, \"StringEquals\":
{\"aws:SourceAccount\": \"account-id\"}}}]}"

```

2. Copy the role ARN from the role metadata in the output. You use the ARN to associate the role with your account.
3. Attach the AWSGreengrassResourceAccessRolePolicy policy to the role.

```

aws iam attach-role-policy --role-name Greengrass_ServiceRole --policy-arn
arn:aws:iam::aws:policy/service-role/AWSGreengrassResourceAccessRolePolicy

```

To associate the service role with your AWS account

- Associate the role with your account. Replace *role-arn* with the service role ARN and *region* with your AWS Region (for example, us-west-2).

```

aws greengrass associate-service-role-to-account --role-arn role-arn --
region region

```

If successful, the following response is returned.

```

{
  "AssociatedAt": "timestamp"
}

```

```
}
```

Remove the Greengrass service role (CLI)

Use the following steps to disassociate the Greengrass service role from your AWS account.

- Disassociate the service role from your account. Replace *region* with your AWS Region (for example, `us-west-2`).

```
aws greengrass disassociate-service-role-from-account --region region
```

If successful, the following response is returned.

```
{
  "DisassociatedAt": "timestamp"
}
```

Note

You should delete the service role if you're not using it in any AWS Region. First use [delete-role-policy](#) to detach the `AWSGreengrassResourceAccessRolePolicy` managed policy from the role, and then use [delete-role](#) to delete the role. For more information, see [Deleting roles or instance profiles](#) in the *IAM User Guide*.

See also

- [Creating a role to delegate permissions to an AWS service](#) in the *IAM User Guide*
- [Modifying a role](#) in the *IAM User Guide*
- [Deleting roles or instance profiles](#) in the *IAM User Guide*
- AWS IoT Greengrass commands in the *AWS CLI Command Reference*
 - [associate-service-role-to-account](#)
 - [disassociate-service-role-from-account](#)
 - [get-service-role-for-account](#)

- IAM commands in the *AWS CLI Command Reference*
 - [attach-role-policy](#)
 - [create-role](#)
 - [delete-role](#)
 - [delete-role-policy](#)

Greengrass group role

The Greengrass group role is an IAM role that authorizes code running on a Greengrass core to access your AWS resources. You create the role and manage permissions in AWS Identity and Access Management (IAM) and attach the role to your Greengrass group. A Greengrass group has one group role. To add or change permissions, you can attach a different role or change the IAM policies that are attached to the role.

The role must define AWS IoT Greengrass as a trusted entity. Depending on your business case, the group role might contain IAM policies that define:

- Permissions for user-defined [Lambda functions](#) to access AWS services.
- Permissions for [connectors](#) to access AWS services.
- Permissions for [stream manager](#) to export streams to AWS IoT Analytics and Kinesis Data Streams.
- Permissions to allow [CloudWatch logging](#).

The following sections describe how to attach or detach a Greengrass group role in the AWS Management Console or AWS CLI.

- [Manage the group role \(console\)](#)
- [Manage the group role \(CLI\)](#)

Note

In addition to the group role that authorizes access from the Greengrass core, you can assign a [Greengrass service role](#) that allows AWS IoT Greengrass to access AWS resources on your behalf.

Managing the Greengrass group role (console)

You can use the AWS IoT console for the following role management tasks:

- [Find your Greengrass group role](#)
- [Add or change the Greengrass group role](#)
- [Remove the Greengrass group role](#)

Note

The user who is signed in to the console must have permissions to manage the role.

Find your Greengrass group role (console)

Follow these steps to find the role that is attached to a Greengrass group.

1. In the AWS IoT console navigation pane, under **Manage**, expand **Greengrass devices**, and then choose **Groups (V1)**.
2. Choose the target group.
3. On the group configuration page, choose **View settings**.

If a role is attached to the group, it appears under **Group role**.

Add or change the Greengrass group role (console)

Follow these steps to choose an IAM role from your AWS account to add to a Greengrass group.

A group role has the following requirements:

- AWS IoT Greengrass defined as a trusted entity.
- The permission policies attached to the role must grant the permissions to your AWS resources that are required by the Lambda functions and connectors in the group, and by Greengrass system components.

Note

We recommend that you also include the `aws:SourceArn` and `aws:SourceAccount` global condition context keys in your trust policy to help prevent the *confused deputy* security problem. The condition context keys restrict access to allow only those requests that come from the specified account and Greengrass workspace. For more information about the confused deputy problem, see [Cross-service confused deputy prevention](#).

Use the IAM console to create and configure the role and its permissions. For steps that create an example role that allows access to an Amazon DynamoDB table, see [the section called “Configure the group role”](#). For general steps, see [Creating a role for an AWS service \(console\)](#) in the *IAM User Guide*.

After the role is configured, use the AWS IoT console to add the role to the group.

Note

This procedure is required only to choose a role for the group. It's not required after changing the permissions of the currently selected group role.

1. In the AWS IoT console navigation pane, under **Manage**, expand **Greengrass devices**, and then choose **Groups (V1)**.
2. Choose the target group.
3. On the group configuration page, choose **View settings**.
4. Under **Group role**, choose to add or change the role:
 - To add the role, choose **Associate role** and then select your role from your list of roles. These are the roles in your AWS account that define AWS IoT Greengrass as a trusted entity.
 - To choose a different role, choose **Edit role** and then select your role from your list of roles.
5. Choose **Save**.

Remove the Greengrass group role (console)

Follow these steps to detach the role from a Greengrass group.

1. In the AWS IoT console navigation pane, under **Manage**, expand **Greengrass devices**, and then choose **Groups (V1)**.
2. Choose the target group.
3. On the group configuration page, choose **View settings**.
4. Under **Group role**, choose **Disassociate role**.
5. In the confirmation dialog box, choose **Disassociate role**. This step removes the role from the group but doesn't delete the role. If you want to delete the role, use the IAM console.

Managing the Greengrass group role (CLI)

You can use the AWS CLI for the following role management tasks:

- [Get your Greengrass group role](#)
- [Create the Greengrass group role](#)
- [Remove the Greengrass group role](#)

Get the Greengrass group role (CLI)

Follow these steps to find out if a Greengrass group has an associated role.

1. Get the ID of the target group from the list of your groups.

```
aws greengrass list-groups
```

The following is an example `list-groups` response. Each group in the response includes an `Id` property that contains the group ID.

```
{
  "Groups": [
    {
```



```

        "LatestVersionArn": "arn:aws:us-west-2:123456789012:/greengrass/
groups/00dedaaa-ac16-484d-ad77-c3eedEXAMPLE/versions/4cbc3f07-fc5e-48c4-
a50e-7d356EXAMPLE",
        "Name": "MyFirstGroup",
        "LastUpdatedTimestamp": "2019-11-11T05:47:31.435Z",
        "LatestVersion": "4cbc3f07-fc5e-48c4-a50e-7d356EXAMPLE",
        "CreationTimestamp": "2019-11-11T05:47:31.435Z",
        "Id": "00dedaaa-ac16-484d-ad77-c3eedEXAMPLE",
        "Arn": "arn:aws:us-west-2:123456789012:/greengrass/groups/00dedaaa-
ac16-484d-ad77-c3eedEXAMPLE"
    },
    {
        "LatestVersionArn": "arn:aws:us-west-2:123456789012:/greengrass/
groups/036ceaf9-9319-4716-ba2a-237f9EXAMPLE/versions/8fe9e8ec-64d1-4647-
b0b0-01dc8EXAMPLE",
        "Name": "GreenhouseSensors",
        "LastUpdatedTimestamp": "2020-01-07T19:58:36.774Z",
        "LatestVersion": "8fe9e8ec-64d1-4647-b0b0-01dc8EXAMPLE",
        "CreationTimestamp": "2020-01-07T19:58:36.774Z",
        "Id": "036ceaf9-9319-4716-ba2a-237f9EXAMPLE",
        "Arn": "arn:aws:us-west-2:123456789012:/greengrass/
groups/036ceaf9-9319-4716-ba2a-237f9EXAMPLE"
    },
    ...
]
}

```

For more information, including examples that use the query option to filter results, see [the section called “Getting the group ID”](#).

2. Copy the Id of the target group from the output.
3. Get the group role. Replace *group-id* with the ID of the target group.

```
aws greengrass get-associated-role --group-id group-id
```

If a role is associated with your Greengrass group, the following role metadata is returned.

```

{
  "AssociatedAt": "timestamp",
  "RoleArn": "arn:aws:iam::account-id:role/path/role-name"
}

```

If your group doesn't have an associated role, the following error is returned.

An error occurred (404) when calling the GetAssociatedRole operation: You need to attach an IAM role to this deployment group.

Create the Greengrass group role (CLI)

Follow these steps to create a role and associate it with a Greengrass group.

To create the group role using IAM

1. Create the role with a trust policy that allows AWS IoT Greengrass to assume the role. This example creates a role named `MyGreengrassGroupRole`, but you can use a different name. We recommend that you also include the `aws:SourceArn` and `aws:SourceAccount` global condition context keys in your trust policy to help prevent the *confused deputy* security problem. The condition context keys restrict access to allow only those requests that come from the specified account and Greengrass workspace. For more information about the confused deputy problem, see [Cross-service confused deputy prevention](#).

Linux, macOS, or Unix

```
aws iam create-role --role-name MyGreengrassGroupRole --assume-role-policy-document '{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Principal": {
        "Service": "greengrass.amazonaws.com"
      },
      "Action": "sts:AssumeRole",
      "Condition": {
        "StringEquals": {
          "aws:SourceAccount": "account-id"
        },
        "ArnLike": {
          "aws:SourceArn": "arn:aws:greengrass:region:account-id:/greengrass/groups/group-id"
        }
      }
    }
  ]
}
```

```
    }
  }
}
]
```

Windows command prompt

```
aws iam create-role --role-name MyGreengrassGroupRole --assume-role-policy-document '{"Version": "2012-10-17", "Statement": [{"Effect": "Allow", "Principal": {"Service": "greengrass.amazonaws.com"}, "Action": "sts:AssumeRole", "Condition": {"ArnLike": {"aws:SourceArn": "arn:aws:greengrass:region:account-id:/greengrass/groups/group-id"}, "StringEquals": {"aws:SourceAccount": "account-id"}}}]}'
```

2. Copy the role ARN from the role metadata in the output. You use the ARN to associate the role with your group.
3. Attach managed or inline policies to the role to support your business case. For example, if a user-defined Lambda function reads from Amazon S3, you might attach the AmazonS3ReadOnlyAccess managed policy to the role.

```
aws iam attach-role-policy --role-name MyGreengrassGroupRole --policy-arn arn:aws:iam::aws:policy/AmazonS3ReadOnlyAccess
```

If successful, no response is returned.

To associate the role with your Greengrass group

1. Get the ID of the target group from the list of your groups.

```
aws greengrass list-groups
```

The following is an example `list-groups` response. Each group in the response includes an `Id` property that contains the group ID.

```
{
  "Groups": [
```

```

    {
      "LatestVersionArn": "arn:aws:us-west-2:123456789012:/greengrass/
groups/00dedaaa-ac16-484d-ad77-c3eedEXAMPLE/versions/4cbc3f07-fc5e-48c4-
a50e-7d356EXAMPLE",
      "Name": "MyFirstGroup",
      "LastUpdatedTimestamp": "2019-11-11T05:47:31.435Z",
      "LatestVersion": "4cbc3f07-fc5e-48c4-a50e-7d356EXAMPLE",
      "CreationTimestamp": "2019-11-11T05:47:31.435Z",
      "Id": "00dedaaa-ac16-484d-ad77-c3eedEXAMPLE",
      "Arn": "arn:aws:us-west-2:123456789012:/greengrass/groups/00dedaaa-
ac16-484d-ad77-c3eedEXAMPLE"
    },
    {
      "LatestVersionArn": "arn:aws:us-west-2:123456789012:/greengrass/
groups/036ceaf9-9319-4716-ba2a-237f9EXAMPLE/versions/8fe9e8ec-64d1-4647-
b0b0-01dc8EXAMPLE",
      "Name": "GreenhouseSensors",
      "LastUpdatedTimestamp": "2020-01-07T19:58:36.774Z",
      "LatestVersion": "8fe9e8ec-64d1-4647-b0b0-01dc8EXAMPLE",
      "CreationTimestamp": "2020-01-07T19:58:36.774Z",
      "Id": "036ceaf9-9319-4716-ba2a-237f9EXAMPLE",
      "Arn": "arn:aws:us-west-2:123456789012:/greengrass/
groups/036ceaf9-9319-4716-ba2a-237f9EXAMPLE"
    },
    ...
  ]
}

```

For more information, including examples that use the query option to filter results, see [the section called “Getting the group ID”](#).

2. Copy the Id of the target group from the output.
3. Associate the role with your group. Replace *group-id* with the ID of the target group and *role-arn* with the ARN of the group role.

```
aws greengrass associate-role-to-group --group-id group-id --role-arn role-arn
```

If successful, the following response is returned.

```

{
  "AssociatedAt": "timestamp"
}

```

```
}
```

Remove the Greengrass group role (CLI)

Follow these steps to disassociate the group role from your Greengrass group.

1. Get the ID of the target group from the list of your groups.

```
aws greengrass list-groups
```

The following is an example `list-groups` response. Each group in the response includes an `Id` property that contains the group ID.

```
{
  "Groups": [
    {
      "LatestVersionArn": "arn:aws:us-west-2:123456789012:/greengrass/
groups/00dedaaa-ac16-484d-ad77-c3eedEXAMPLE/versions/4cbc3f07-fc5e-48c4-
a50e-7d356EXAMPLE",
      "Name": "MyFirstGroup",
      "LastUpdatedTimestamp": "2019-11-11T05:47:31.435Z",
      "LatestVersion": "4cbc3f07-fc5e-48c4-a50e-7d356EXAMPLE",
      "CreationTimestamp": "2019-11-11T05:47:31.435Z",
      "Id": "00dedaaa-ac16-484d-ad77-c3eedEXAMPLE",
      "Arn": "arn:aws:us-west-2:123456789012:/greengrass/groups/00dedaaa-
ac16-484d-ad77-c3eedEXAMPLE"
    },
    {
      "LatestVersionArn": "arn:aws:us-west-2:123456789012:/greengrass/
groups/036ceaf9-9319-4716-ba2a-237f9EXAMPLE/versions/8fe9e8ec-64d1-4647-
b0b0-01dc8EXAMPLE",
      "Name": "GreenhouseSensors",
      "LastUpdatedTimestamp": "2020-01-07T19:58:36.774Z",
      "LatestVersion": "8fe9e8ec-64d1-4647-b0b0-01dc8EXAMPLE",
      "CreationTimestamp": "2020-01-07T19:58:36.774Z",
      "Id": "036ceaf9-9319-4716-ba2a-237f9EXAMPLE",
      "Arn": "arn:aws:us-west-2:123456789012:/greengrass/
groups/036ceaf9-9319-4716-ba2a-237f9EXAMPLE"
    },
    ...
  ]
}
```

```
]
}
```

For more information, including examples that use the `query` option to filter results, see [the section called “Getting the group ID”](#).

2. Copy the Id of the target group from the output.
3. Disassociate the role from your group. Replace *group-id* with the ID of the target group.

```
aws greengrass disassociate-role-from-group --group-id group-id
```

If successful, the following response is returned.

```
{
  "DisassociatedAt": "timestamp"
}
```

Note

You can delete the group role if you're not using it. First use [delete-role-policy](#) to detach each managed policy from the role, and then use [delete-role](#) to delete the role. For more information, see [Deleting roles or instance profiles](#) in the *IAM User Guide*.

See also

- Related topics in the *IAM User Guide*
 - [Creating a role to delegate permissions to an AWS service](#)
 - [Modifying a role](#)
 - [Adding and removing IAM identity permissions](#)
 - [Deleting roles or instance profiles](#)
- AWS IoT Greengrass commands in the *AWS CLI Command Reference*
 - [list-groups](#)
 - [associate-role-to-group](#)
 - [disassociate-role-from-group](#)
 - [get-associated-role](#)

- IAM commands in the *AWS CLI Command Reference*
 - [attach-role-policy](#)
 - [create-role](#)
 - [delete-role](#)
 - [delete-role-policy](#)

Cross-service confused deputy prevention

The confused deputy problem is a security issue where an entity that doesn't have permission to perform an action can coerce a more-privileged entity to perform the action. In AWS, cross-service impersonation can result in the confused deputy problem. Cross-service impersonation can occur when one service (the *calling service*) calls another service (the *called service*). The calling service can be manipulated to use its permissions to act on another customer's resources in a way it should not otherwise have permission to access. To prevent this, AWS provides tools that help you protect your data for all services with service principals that have been given access to resources in your account.

We recommend using the [aws:SourceArn](#) and [aws:SourceAccount](#) global condition context keys in resource policies to limit the permissions that AWS IoT Greengrass gives another service to the resource. If you use both global condition context keys, the `aws:SourceAccount` value and the account in the `aws:SourceArn` value must use the same account ID when used in the same policy statement.

The value of `aws:SourceArn` must be the Greengrass customer resource that is associated with the `sts:AssumeRole` request.

The most effective way to protect against the confused deputy problem is to use the `aws:SourceArn` global condition context key with the full ARN of the resource. If you don't know the full ARN of the resource or if you are specifying multiple resources, use the `aws:SourceArn` global context condition key with wildcards (*) for the unknown portions of the ARN. For example, `arn:aws:greengrass:region:account-id:*`.

For examples of policies that use the `aws:SourceArn` and `aws:SourceAccount` global condition context keys, see the following topics:

- [Create the Greengrass service role](#)
- [Create the Greengrass group role](#)

- [Create and configure an IAM execution role for bulk deployments](#)

Identity-based policy examples for AWS IoT Greengrass

By default, IAM users and roles don't have permission to create or modify AWS IoT Greengrass resources. They also can't perform tasks using the AWS Management Console, AWS CLI, or AWS API. An IAM administrator must create IAM policies that grant users and roles permission to perform specific API operations on the specified resources they need. The administrator must then attach those policies to the IAM users or groups that require those permissions.

Policy best practices

Identity-based policies determine whether someone can create, access, or delete AWS IoT Greengrass resources in your account. These actions can incur costs for your AWS account. When you create or edit identity-based policies, follow these guidelines and recommendations:

- **Get started with AWS managed policies and move toward least-privilege permissions** – To get started granting permissions to your users and workloads, use the *AWS managed policies* that grant permissions for many common use cases. They are available in your AWS account. We recommend that you reduce permissions further by defining AWS customer managed policies that are specific to your use cases. For more information, see [AWS managed policies](#) or [AWS managed policies for job functions](#) in the *IAM User Guide*.
- **Apply least-privilege permissions** – When you set permissions with IAM policies, grant only the permissions required to perform a task. You do this by defining the actions that can be taken on specific resources under specific conditions, also known as *least-privilege permissions*. For more information about using IAM to apply permissions, see [Policies and permissions in IAM](#) in the *IAM User Guide*.
- **Use conditions in IAM policies to further restrict access** – You can add a condition to your policies to limit access to actions and resources. For example, you can write a policy condition to specify that all requests must be sent using SSL. You can also use conditions to grant access to service actions if they are used through a specific AWS service, such as AWS CloudFormation. For more information, see [IAM JSON policy elements: Condition](#) in the *IAM User Guide*.
- **Use IAM Access Analyzer to validate your IAM policies to ensure secure and functional permissions** – IAM Access Analyzer validates new and existing policies so that the policies adhere to the IAM policy language (JSON) and IAM best practices. IAM Access Analyzer provides more than 100 policy checks and actionable recommendations to help you author secure and

functional policies. For more information, see [Validate policies with IAM Access Analyzer](#) in the *IAM User Guide*.

- **Require multi-factor authentication (MFA)** – If you have a scenario that requires IAM users or a root user in your AWS account, turn on MFA for additional security. To require MFA when API operations are called, add MFA conditions to your policies. For more information, see [Secure API access with MFA](#) in the *IAM User Guide*.

For more information about best practices in IAM, see [Security best practices in IAM](#) in the *IAM User Guide*.

AWS managed policies for AWS IoT Greengrass

AWS IoT Greengrass maintains the following AWS managed policies that you can use to grant permissions to IAM users and roles.

Policy	Description
AWSGreengrassFullAccess	Allows all AWS IoT Greengrass actions for all of your AWS resources. This policy is recommended for AWS IoT Greengrass service administrators or testing purposes.
AWSGreengrassReadOnlyAccess	Allows List and Get AWS IoT Greengrass actions for all of your AWS resources.
AWSGreengrassResourceAccessRolePolicy	Allows access to resources from AWS services including AWS Lambda and AWS IoT Device Shadow. This is the default policy used for the Greengrass service role . This policy is designed to provide general ease of access. You can define a custom policy that is more restrictive.
GreengrassOTAUpdateArtifactAccess	Allows read-only access to over-the-air (OTA) update artifacts for the AWS IoT Greengrass Core software in all AWS Regions.

Policy examples

The following example customer-defined policies grant permissions for common scenarios.

Examples

- [Allow users to view their own permissions](#)

To learn how to create an IAM identity-based policy using these example JSON policy documents, see [Creating policies on the JSON tab](#) in the *IAM User Guide*.

Allow users to view their own permissions

This example shows how you might create a policy that allows IAM users to view the inline and managed policies that are attached to their user identity. This policy includes permissions to complete this action on the console or programmatically using the AWS CLI or AWS API.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "ViewOwnUserInfo",
      "Effect": "Allow",
      "Action": [
        "iam:GetUserPolicy",
        "iam:ListGroupsForUser",
        "iam:ListAttachedUserPolicies",
        "iam:ListUserPolicies",
        "iam:GetUser"
      ],
      "Resource": ["arn:aws:iam::*:user/${aws:username}"]
    },
    {
      "Sid": "NavigateInConsole",
      "Effect": "Allow",
      "Action": [
        "iam:GetGroupPolicy",
        "iam:GetPolicyVersion",
        "iam:GetPolicy",
        "iam:ListAttachedGroupPolicies",
        "iam:ListGroupPolicies",
        "iam:ListPolicyVersions",
        "iam:ListPolicies",

```

```
        "iam:ListUsers"
      ],
      "Resource": "*"
    }
  ]
}
```

Troubleshooting identity and access issues for AWS IoT Greengrass

Use the following information to help you diagnose and fix common issues that you might encounter when working with AWS IoT Greengrass and IAM.

Issues

- [I'm not authorized to perform an action in AWS IoT Greengrass](#)
- [Error: Greengrass is not authorized to assume the Service Role associated with this account, or the error: Failed: TES service role is not associated with this account.](#)
- [Error: Permission denied when attempting to use role arn:aws:iam::<account-id>:role/<role-name> to access s3 url https://<region>-greengrass-updates.s3.<region>.amazonaws.com/core/<architecture>/greengrass-core-<distribution-version>.tar.gz.](#)
- [Device shadow does not sync with the cloud.](#)
- [I'm not authorized to perform iam:PassRole](#)
- [I'm an administrator and want to allow others to access AWS IoT Greengrass](#)
- [I want to allow people outside of my AWS account to access my AWS IoT Greengrass resources](#)

For general troubleshooting help, see [Troubleshooting](#).

I'm not authorized to perform an action in AWS IoT Greengrass

If you receive an error that states you're not authorized to perform an action, you must contact your administrator for assistance. Your administrator is the person who provided you with your user name and password.

The following example error occurs when the mateojackson IAM user tries to view details about a core definition version, but does not have `greengrass:GetCoreDefinitionVersion` permissions.

```
User: arn:aws:iam::123456789012:user/mateojackson is not authorized to perform:
greengrass:GetCoreDefinitionVersion on resource: resource: arn:aws:greengrass:us-
```

```
west-2:123456789012:/greengrass/definition/cores/78cd17f3-bc68-ee18-47bd-5bda5EXAMPLE/versions/368e9ffa-4939-6c75-859c-0bd4cEXAMPLE
```

In this case, Mateo asks his administrator to update his policies to allow him to access the `arn:aws:greengrass:us-west-2:123456789012:/greengrass/definition/cores/78cd17f3-bc68-ee18-47bd-5bda5EXAMPLE/versions/368e9ffa-4939-6c75-859c-0bd4cEXAMPLE` resource using the `greengrass:GetCoreDefinitionVersion` action.

Error: Greengrass is not authorized to assume the Service Role associated with this account, or the error: Failed: TES service role is not associated with this account.

Solution: You might see this error when the deployment fails. Check that a Greengrass service role is associated with your AWS account in the current AWS Region. For more information, see [the section called “Manage the service role \(CLI\)”](#) or [the section called “Manage the service role \(console\)”](#).

Error: Permission denied when attempting to use role `arn:aws:iam::<account-id>:role/<role-name>` to access s3 url `https://<region>-greengrass-updates.s3.<region>.amazonaws.com/core/<architecture>/greengrass-core-<distribution-version>.tar.gz`.

Solution: You might see this error when an over-the-air (OTA) update fails. In the signer role policy, add the target AWS Region as a Resource. This signer role is used to presign the S3 URL for the AWS IoT Greengrass software update. For more information, see [S3 URL signer role](#).

Device shadow does not sync with the cloud.

Solution: Make sure that AWS IoT Greengrass has permissions for `iot:UpdateThingShadow` and `iot:GetThingShadow` actions in the [Greengrass service role](#). If the service role uses the `AWSGreengrassResourceAccessRolePolicy` managed policy, these permissions are included by default.

See [Troubleshooting shadow synchronization timeout issues](#).

The following are general IAM issues that you might encounter when working with AWS IoT Greengrass.

I'm not authorized to perform iam:PassRole

If you receive an error that you're not authorized to perform the `iam:PassRole` action, your policies must be updated to allow you to pass a role to AWS IoT Greengrass.

Some AWS services allow you to pass an existing role to that service instead of creating a new service role or service-linked role. To do this, you must have permissions to pass the role to the service.

The following example error occurs when an IAM user named `marymajor` tries to use the console to perform an action in AWS IoT Greengrass. However, the action requires the service to have permissions that are granted by a service role. Mary does not have permissions to pass the role to the service.

```
User: arn:aws:iam::123456789012:user/marymajor is not authorized to perform:
iam:PassRole
```

In this case, Mary's policies must be updated to allow her to perform the `iam:PassRole` action.

If you need help, contact your AWS administrator. Your administrator is the person who provided you with your sign-in credentials.

I'm an administrator and want to allow others to access AWS IoT Greengrass

To allow others to access AWS IoT Greengrass, you must grant permission to the people or applications that need access. If you are using AWS IAM Identity Center to manage people and applications, you assign permission sets to users or groups to define their level of access. Permission sets automatically create and assign IAM policies to IAM roles that are associated with the person or application. For more information, see [Permission sets](#) in the *AWS IAM Identity Center User Guide*.

If you are not using IAM Identity Center, you must create IAM entities (users or roles) for the people or applications that need access. You must then attach a policy to the entity that grants them the correct permissions in AWS IoT Greengrass. After the permissions are granted, provide the credentials to the user or application developer. They will use those credentials to access AWS. To learn more about creating IAM users, groups, policies, and permissions, see [IAM Identities](#) and [Policies and permissions in IAM](#) in the *IAM User Guide*.

I want to allow people outside of my AWS account to access my AWS IoT Greengrass resources

You can create an IAM role that users in other accounts or people outside of your organization can use to access your AWS resources. You can specify the who is trusted to assume the role. For more information, see [Providing access to an IAM user in another AWS account that you own](#) and [Providing access to Amazon Web Services accounts owned by third parties](#) in the *IAM User Guide*.

AWS IoT Greengrass doesn't support cross-account access based on resource-based policies or access control lists (ACLs).

Compliance validation for AWS IoT Greengrass

To learn whether an AWS service is within the scope of specific compliance programs, see [AWS services in Scope by Compliance Program](#) and choose the compliance program that you are interested in. For general information, see [AWS Compliance Programs](#).

You can download third-party audit reports using AWS Artifact. For more information, see [Downloading Reports in AWS Artifact](#).

Your compliance responsibility when using AWS services is determined by the sensitivity of your data, your company's compliance objectives, and applicable laws and regulations. AWS provides the following resources to help with compliance:

- [Security Compliance & Governance](#) – These solution implementation guides discuss architectural considerations and provide steps for deploying security and compliance features.
- [HIPAA Eligible Services Reference](#) – Lists HIPAA eligible services. Not all AWS services are HIPAA eligible.
- [AWS Compliance Resources](#) – This collection of workbooks and guides might apply to your industry and location.
- [AWS Customer Compliance Guides](#) – Understand the shared responsibility model through the lens of compliance. The guides summarize the best practices for securing AWS services and map the guidance to security controls across multiple frameworks (including National Institute of Standards and Technology (NIST), Payment Card Industry Security Standards Council (PCI), and International Organization for Standardization (ISO)).
- [Evaluating Resources with Rules](#) in the *AWS Config Developer Guide* – The AWS Config service assesses how well your resource configurations comply with internal practices, industry guidelines, and regulations.

- [AWS Security Hub](#) – This AWS service provides a comprehensive view of your security state within AWS. Security Hub uses security controls to evaluate your AWS resources and to check your compliance against security industry standards and best practices. For a list of supported services and controls, see [Security Hub controls reference](#).
- [Amazon GuardDuty](#) – This AWS service detects potential threats to your AWS accounts, workloads, containers, and data by monitoring your environment for suspicious and malicious activities. GuardDuty can help you address various compliance requirements, like PCI DSS, by meeting intrusion detection requirements mandated by certain compliance frameworks.
- [AWS Audit Manager](#) – This AWS service helps you continuously audit your AWS usage to simplify how you manage risk and compliance with regulations and industry standards.

Resilience in AWS IoT Greengrass

The AWS global infrastructure is built around Amazon Web Services Regions and Availability Zones. Each AWS Region provides multiple physically separated and isolated Availability Zones, which are connected with low-latency, high-throughput, and highly redundant networking. With Availability Zones, you can design and operate applications and databases that automatically fail over between zones without interruption. Availability Zones are more highly available, fault tolerant, and scalable than traditional single or multiple data center infrastructures.

For more information about Amazon Web Services Regions and Availability Zones, see [AWS Global Infrastructure](#).

In addition to the AWS global infrastructure, AWS IoT Greengrass offers several features to help support your data resiliency and backup needs.

- If the core loses internet connectivity, client devices can continue to communicate over the local network.
- You can configure the core to store unprocessed messages destined for AWS Cloud targets in a local storage cache instead of in-memory storage. The local storage cache can persist across core restarts (for example, after a group deployment or a device reboot), so AWS IoT Greengrass can continue to process messages destined for AWS IoT Core. For more information, see [the section called “MQTT message queue”](#).
- You can configure the core to establish a persistent session with the AWS IoT Core message broker. This allows the core to receive messages sent while the core is offline. For more information, see [the section called “MQTT persistent sessions with AWS IoT Core”](#).

- You can configure a Greengrass group to write logs to the local file system and to CloudWatch Logs. If the core loses connectivity, local logging can continue, but CloudWatch logs are sent with a limited number of retries. After the retries are exhausted, the event is dropped. You should also be aware of [logging limitations](#).
- You can author Lambda functions that read [stream manager](#) streams and send the data to local storage destinations.

Infrastructure security in AWS IoT Greengrass

As a managed service, AWS IoT Greengrass is protected by AWS global network security. For information about AWS security services and how AWS protects infrastructure, see [AWS Cloud Security](#). To design your AWS environment using the best practices for infrastructure security, see [Infrastructure Protection](#) in *Security Pillar AWS Well-Architected Framework*.

You use AWS published API calls to access AWS IoT Greengrass through the network. Clients must support the following:

- Transport Layer Security (TLS). We require TLS 1.2 and recommend TLS 1.3.
- Cipher suites with perfect forward secrecy (PFS) such as DHE (Ephemeral Diffie-Hellman) or ECDHE (Elliptic Curve Ephemeral Diffie-Hellman). Most modern systems such as Java 7 and later support these modes.

Additionally, requests must be signed by using an access key ID and a secret access key that is associated with an IAM principal. Or you can use the [AWS Security Token Service](#) (AWS STS) to generate temporary security credentials to sign requests.

In an AWS IoT Greengrass environment, devices use X.509 certificates and cryptographic keys to connect and authenticate to the AWS Cloud. For more information, see [the section called “Device authentication and authorization”](#).

Configuration and vulnerability analysis in AWS IoT Greengrass

IoT environments can consist of large numbers of devices that have diverse capabilities, are long-lived, and are geographically distributed. These characteristics make device setup complex and error-prone. And because devices are often constrained in computational power, memory, and storage capabilities, this limits the use of encryption and other forms of security on the devices

themselves. Also, devices often use software with known vulnerabilities. These factors make IoT devices an attractive target for hackers and make it difficult to secure them on an ongoing basis.

AWS IoT Device Defender addresses these challenges by providing tools to identify security issues and deviations from best practices. You can use AWS IoT Device Defender to analyze, audit, and monitor connected devices to detect abnormal behavior, and mitigate security risks. AWS IoT Device Defender can audit devices to ensure they adhere to security best practices and detect abnormal behavior on devices. This makes it possible to enforce consistent security policies across your devices and respond quickly when devices are compromised. In connections with AWS IoT Core, AWS IoT Greengrass generates [predictable client IDs](#) that you can use with AWS IoT Device Defender features. For more information, see [AWS IoT Device Defender](#) in the *AWS IoT Core Developer Guide*.

In AWS IoT Greengrass environments, you should be aware of the following considerations:

- It's your responsibility to secure your physical devices, the file system on your devices, and the local network.
- AWS IoT Greengrass doesn't enforce network isolation for user-defined Lambda functions, whether or not they run in a [Greengrass container](#). Therefore, it's possible for Lambda functions to communicate with any other process running in the system or outside over network.

If you lose control of a Greengrass core device and you want to prevent client devices from transmitting data to the core, do the following:

1. Remove the Greengrass core from the Greengrass group.
2. Rotate the group CA certificate. In the AWS IoT console, you can rotate the CA certificate on the group's **Settings** page. In the AWS IoT Greengrass API, you can use the [CreateGroupCertificateAuthority](#) action.

We also recommend using full disk encryption if the hard drive of your core device is vulnerable to theft.

AWS IoT Greengrass and interface VPC endpoints (AWS PrivateLink)

You can establish a private connection between your VPC and the AWS IoT Greengrass control plane by creating an *interface VPC endpoint*. You can use this endpoint to manage groups,

Lambda functions, deployments, and other resources in the AWS IoT Greengrass service. Interface endpoints are powered by [AWS PrivateLink](#), a technology that enables you to access AWS IoT Greengrass APIs privately without an internet gateway, NAT device, VPN connection, or AWS Direct Connect connection. Instances in your VPC don't need public IP addresses to communicate with AWS IoT Greengrass APIs. Traffic between your VPC and AWS IoT Greengrass does not leave the Amazon network.

 Note

Currently, you can't configure Greengrass core devices to operate completely within your VPC.

Each interface endpoint is represented by one or more [Elastic Network Interfaces](#) in your subnets.

For more information, see [Interface VPC endpoints \(AWS PrivateLink\)](#) in the *Amazon VPC User Guide*.

Topics

- [Considerations for AWS IoT Greengrass VPC endpoints](#)
- [Create an interface VPC endpoint for AWS IoT Greengrass control plane operations](#)
- [Creating a VPC endpoint policy for AWS IoT Greengrass](#)

Considerations for AWS IoT Greengrass VPC endpoints

Before you set up an interface VPC endpoint for AWS IoT Greengrass, review [Interface endpoint properties and limitations](#) in the *Amazon VPC User Guide*. Additionally, be aware of the following considerations:

- AWS IoT Greengrass supports making calls to all of its control plane API actions from your VPC. The control plane includes operations such as [CreateDeployment](#) and [StartBulkDeployment](#). The control plane does *not* include operations such as [GetDeployment](#) and [Discover](#), which are data plane operations.
- VPC endpoints for AWS IoT Greengrass are currently not supported in AWS China Regions.

Create an interface VPC endpoint for AWS IoT Greengrass control plane operations

You can create a VPC endpoint for the AWS IoT Greengrass control plane using either the Amazon VPC console or the AWS Command Line Interface (AWS CLI). For more information, see [Creating an interface endpoint](#) in the *Amazon VPC User Guide*.

Create a VPC endpoint for AWS IoT Greengrass using the following service name:

- `com.amazonaws.region.greengrass`

If you enable private DNS for the endpoint, you can make API requests to AWS IoT Greengrass using its default DNS name for the Region, for example, `greengrass.us-east-1.amazonaws.com`. Private DNS is enabled by default.

For more information, see [Accessing a service through an interface endpoint](#) in the *Amazon VPC User Guide*.

Creating a VPC endpoint policy for AWS IoT Greengrass

You can attach an endpoint policy to your VPC endpoint that controls access to AWS IoT Greengrass control plane operations. The policy specifies the following information:

- The principal that can perform actions.
- The actions that the principal can perform.
- The resources that the principal can perform actions on.

For more information, see [Controlling access to services with VPC endpoints](#) in the *Amazon VPC User Guide*.

Example Example: VPC endpoint policy for AWS IoT Greengrass actions

The following is an example of an endpoint policy for AWS IoT Greengrass. When attached to an endpoint, this policy grants access to the listed AWS IoT Greengrass actions for all principals on all resources.

```
{
```

```
"Statement": [  
  {  
    "Principal": "*",  
    "Effect": "Allow",  
    "Action": [  
      "greengrass:CreateDeployment",  
      "greengrass:StartBulkDeployment"  
    ],  
    "Resource": "*"   
  }  
]
```

Security best practices for AWS IoT Greengrass

This topic contains security best practices for AWS IoT Greengrass.

Grant minimum possible permissions

Follow the principle of least privilege by using the minimum set of permissions in IAM roles. Limit the use of the * wildcard for the Action and Resource properties in your IAM policies. Instead, declare a finite set of actions and resources when possible. For more information about least privilege and other policy best practices, see [the section called “Policy best practices”](#).

The least privilege best practice also applies to AWS IoT policies you attach to your Greengrass core and client devices.

Don't hardcode credentials in Lambda functions

Don't hardcode credentials in your user-defined Lambda functions. To better protect your credentials:

- To interact with AWS services, define permissions for specific actions and resources in the [Greengrass group role](#).
- Use [local secrets](#) to store your credentials. Or, if the function uses the AWS SDK, use credentials from the default credential provider chain.

Don't log sensitive information

You should prevent the logging of credentials and other personally identifiable information (PII). We recommend that you implement the following safeguards even though access to local logs on a core device requires root privileges and access to CloudWatch Logs requires IAM permissions.

- Don't use sensitive information in MQTT topic paths.
- Don't use sensitive information in device (thing) names, types, and attributes in the AWS IoT Core registry.
- Don't log sensitive information in your user-defined Lambda functions.
- Don't use sensitive information in the names and IDs of Greengrass resources:
 - Connectors
 - Cores
 - Devices
 - Functions
 - Groups
 - Loggers
 - Resources (local, machine learning, or secret)
 - Subscriptions

Create targeted subscriptions

Subscriptions control the information flow in a Greengrass group by defining how messages are exchanged between services, devices, and Lambda functions. To ensure that an application can do only what it's intended to do, your subscriptions should allow publishers to send messages to specific topics only, and limit subscribers to receive messages only from topics that are required for their functionality.

Keep your device clock in sync

It's important to have an accurate time on your device. X.509 certificates have an expiry date and time. The clock on your device is used to verify that a server certificate is still valid. Device clocks can drift over time or batteries can get discharged.

For more information, see the [Keep your device's clock in sync](#) best practice in the *AWS IoT Core Developer Guide*.

Manage device authentication with the Greengrass core

Client devices can run [FreeRTOS](#) or use the [AWS IoT Device SDK](#) or [AWS IoT Greengrass Discovery API](#) to get discovery information used to connect and authenticate with the core in the same Greengrass group. Discovery information includes:

- Connectivity information for the Greengrass core that's in the same Greengrass group as the client device. This information includes the host address and port number of each endpoint for the core device.
- The group CA certificate used to sign the local MQTT server certificate. Client devices use the group CA certificate to validate the MQTT server certificate presented by the core.

The following are best practices for client devices to manage mutual authentication with a Greengrass core. These practices can help mitigate your risk if your core device is compromised.

Validate the local MQTT server certificate for each connection.

Client devices should validate the MQTT server certificate presented by the core every time they establish a connection with the core. This validation is the *client device* side of the mutual authentication between a core device and client devices. Client devices must be able to detect a failure and terminate the connection.

Do not hardcode discovery information.

Client devices should rely on discovery operations to get core connectivity information and the group CA certificate, even if the core uses a static IP address. Client devices should not hardcode this discovery information.

Periodically update discovery information.

Client devices should periodically run discovery to update core connectivity information and the group CA certificate. We recommend that client devices update this information before they establish a connection with the core. Because shorter durations between discovery operations can minimize your potential exposure time, we recommend that client devices periodically disconnect and reconnect to trigger the update.

If you lose control of a Greengrass core device and you want to prevent client devices from transmitting data to the core, do the following:

1. Remove the Greengrass core from the Greengrass group.

2. Rotate the group CA certificate. In the AWS IoT console, you can rotate the CA certificate on the group's **Settings** page. In the AWS IoT Greengrass API, you can use the [CreateGroupCertificateAuthority](#) action.

We also recommend using full disk encryption if the hard drive of your core device is vulnerable to theft.

For more information, see [the section called "Device authentication and authorization"](#).

See also

- [Security best practices in AWS IoT Core](#) in the *AWS IoT Developer Guide*
- [Ten security golden rules for Industrial IoT solutions](#) on the *Internet of Things on AWS Official Blog*

Logging and monitoring in AWS IoT Greengrass

Monitoring is an important part of maintaining the reliability, availability, and performance of AWS IoT Greengrass and your AWS solutions. You should collect monitoring data from all parts of your AWS solution so that you can more easily debug a multi-point failure, if one occurs. Before you start monitoring AWS IoT Greengrass, you should create a monitoring plan that includes answers to the following questions:

- What are your monitoring goals?
- Which resources will you monitor?
- How often will you monitor these resources?
- Which monitoring tools will you use?
- Who will perform the monitoring tasks?
- Who should be notified when something goes wrong?

Monitoring tools

AWS provides tools that you can use to monitor AWS IoT Greengrass. You can configure some of these tools to do the monitoring for you. Some of the tools require manual intervention. We recommend that you automate monitoring tasks as much as possible.

You can use the following automated monitoring tools to monitor AWS IoT Greengrass and report issues:

- **Amazon CloudWatch Logs** – Monitor, store, and access your log files from AWS CloudTrail or other sources. For more information, see [Monitoring log files](#) in the *Amazon CloudWatch User Guide*.
- **AWS CloudTrail Log Monitoring** – Share log files between accounts, monitor CloudTrail log files in real time by sending them to CloudWatch Logs, write log processing applications in Java, and validate that your log files have not changed after delivery by CloudTrail. For more information, see [Working with CloudTrail log files](#) in the *AWS CloudTrail User Guide*.
- **Amazon EventBridge** – Use EventBridge event rules to get notifications about state changes for your Greengrass group deployments or API calls logged with CloudTrail. For more information, see [the section called “Get deployment notifications”](#) or [What is Amazon EventBridge?](#) in the *Amazon EventBridge User Guide*.

- **Greengrass system health telemetry** – Subscribe to receive telemetry data sent from the Greengrass core. For more information, see [the section called “Gathering system health telemetry data”](#).
- **Local health check** – Use the health APIs to get a snapshot of the state of local AWS IoT Greengrass processes on the core device. For more information, see [the section called “Calling the local health check API”](#).

See also

- [the section called “Monitoring with AWS IoT Greengrass logs”](#)
- [the section called “Logging AWS IoT Greengrass API calls with AWS CloudTrail”](#)
- [the section called “Get deployment notifications”](#)

Monitoring with AWS IoT Greengrass logs

AWS IoT Greengrass consists of the cloud service and the AWS IoT Greengrass Core software. The AWS IoT Greengrass Core software can write logs to Amazon CloudWatch and to the local file system of your core device. Lambda functions and connectors running on the core can also write logs to CloudWatch Logs and the local file system. You can use logs to monitor events and troubleshoot issues. All AWS IoT Greengrass log entries include a timestamp, log level, and information about the event. Changes to logging settings take effect after you deploy the group.

Logging is configured at the group level. For steps that show how to configure logging for a Greengrass group, see [the section called “Configure logging for AWS IoT Greengrass”](#).

Accessing CloudWatch Logs

If you configure CloudWatch logging, you can view the logs on the **Logs** page of the Amazon CloudWatch console. Log groups for AWS IoT Greengrass logs use the following naming conventions:

```
/aws/greengrass/GreengrassSystem/greengrass-system-component-name  
/aws/greengrass/Lambda/aws-region/account-id/lambda-function-name
```

Each log group contains log streams that use the following naming convention:

date/account-id/greengrass-group-id/name-of-core-that-generated-log

The following considerations apply when you use CloudWatch Logs:

- Logs are sent to CloudWatch Logs with a limited number of retries in case there's no internet connectivity. After the retries are exhausted, the event is dropped.
- Transaction, memory, and other limitations apply. For more information, see [the section called "Logging limitations"](#).
- Your Greengrass group role must allow AWS IoT Greengrass to write to CloudWatch Logs. To grant permissions, [embed the following inline policy](#) in your group role.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "logs:CreateLogGroup",
        "logs:CreateLogStream",
        "logs:PutLogEvents",
        "logs:DescribeLogStreams"
      ],
      "Resource": [
        "arn:aws:logs:*:*:*"
      ]
    }
  ]
}
```

Note

You can grant more granular access to your log resources. For more information, see [Using identity-based policies \(IAM policies\) for CloudWatch Logs](#) in the *Amazon CloudWatch User Guide*.

The group role is an IAM role that you create and attach to your Greengrass group. You can use the console or the AWS IoT Greengrass API to manage the group role.

Using the console

1. In the AWS IoT console navigation pane, under **Manage**, expand **Greengrass devices**, and then choose **Groups (V1)**.
2. Choose the target group.
3. Choose **View settings**. Under **Group role**, you can view, associate, or disassociate the group role.

For steps that show you how to attach the group role, see [group role](#).

Using the CLI

- To find the group role, use the [get-associated-role](#) command.
- To attach the group role, use the [associate-role-to-group](#) command.
- To remove the group role, use the [disassociate-role-from-group](#) command.

To learn how to get the group ID to use with these commands, see [the section called “Getting the group ID”](#).

Accessing file system logs

If you configure file system logging, the log files are stored under *greengrass-root*/ggc/var/log on the core device. The following is the high-level directory structure:

```
greengrass-root/ggc/var/log
- crash.log
- system
  - log files for each Greengrass system component
- user
  - region
    - account-id
      - log files generated by each user-defined Lambda function
  - aws
```

- log files generated by each connector

Note

By default, *greengrass-root* is the /greengrass directory. If a [write directory](#) is configured, then the logs are under that directory.

The following considerations apply when you use file system logs:

- Reading AWS IoT Greengrass logs on the file system requires root permissions.
- AWS IoT Greengrass supports size-based rotation and automatic cleanup when the amount of log data is close to the configured limit.
- The `crash.log` file is available in file system logs only. This log isn't written to CloudWatch Logs.
- Disk usage limitations apply. For more information, see [the section called "Logging limitations"](#).

Note

Logs for AWS IoT Greengrass Core software v1.0 are stored under the *greengrass-root*/var/log directory.

Default logging configuration

If logging settings aren't explicitly configured, AWS IoT Greengrass uses the following default logging configuration after the first group deployment.

AWS IoT Greengrass System Components

- Type - FileSystem
- Component - GreengrassSystem
- Level - INFO
- Space - 128 KB

User-defined Lambda Functions

- Type - FileSystem

- Component - Lambda
- Level - INFO
- Space - 128 KB

Note

Before the first deployment, only system components write logs to the file system because no user-defined Lambda functions are deployed.

Configure logging for AWS IoT Greengrass

You can use the AWS IoT console or the [AWS IoT Greengrass APIs](#) to configure AWS IoT Greengrass logging.

Note

To allow AWS IoT Greengrass to write logs to CloudWatch Logs, your group role must allow the [required CloudWatch Logs actions](#).

Configure logging (console)

You can configure logging on the group's **Settings** page.

1. In the AWS IoT console navigation pane, under **Manage**, expand **Greengrass devices**, and then choose **Groups (V1)**.
2. Choose the group where you want to configure logging.
3. On the group configuration page, choose the **Logs** tab.
4. Choose the logging location, as follows:
 - To configure CloudWatch logging, for **CloudWatch logs configuration**, choose **Edit**.
 - To configure file system logging, for **Local logs configuration**, choose **Edit**.

You can configure logging for one location or both locations.

5. In the edit logs configuration modal, select **Greengrass system log level** or **User Lambda functions log level**. You can choose one component or both components.
6. Choose the lowest level of events that you want to log. Events below this threshold are filtered out and aren't stored.
7. Choose **Save**. Changes take effect after you deploy the group.

Configure logging (API)

You can use AWS IoT Greengrass logger APIs to configure logging programmatically. For example, use the [CreateLoggerDefinition](#) action to create a logger definition based on a [LoggerDefinitionVersion](#) payload, which uses the following syntax:

```
{
  "Loggers": [
    {
      "Id": "string",
      "Type": "FileSystem|AWSCloudWatch",
      "Component": "GreengrassSystem|Lambda",
      "Level": "DEBUG|INFO|WARN|ERROR|FATAL",
      "Space": "integer"
    },
    {
      "Id": "string",
      ...
    }
  ]
}
```

LoggerDefinitionVersion is an array of one or more [Logger](#) objects that have the following properties:

Id

An identifier for the logger.

Type

The storage mechanism for log events. When AWSCloudWatch is used, log events are sent to CloudWatch Logs. When FileSystem is used, log events are stored on the local file system.

Valid values: AWSCloudWatch, FileSystem

Component

The source of the log event. When GreengrassSystem is used, events from Greengrass system components are logged. When Lambda is used, events from user-defined Lambda functions are logged.

Valid values: GreengrassSystem, Lambda

Level

The log-level threshold. Log events below this threshold are filtered out and aren't stored.

Valid values: DEBUG, INFO (recommended), WARN, ERROR, FATAL

Space

The maximum amount of local storage, in KB, to use for storing logs. This field applies only when Type is set to FileSystem.

Configuration example

The following LoggerDefinitionVersion example specifies a logging configuration that:

- Turns on file system ERROR and above logging for AWS IoT Greengrass system components.
- Turns on file system INFO (and above) logging for user-defined Lambda functions.
- Turns on CloudWatch INFO (and above) logging for user-defined Lambda functions.

```
{
  "Name": "LoggingExample",
  "InitialVersion": {
    "Loggers": [
      {
        "Id": "1",
        "Component": "GreengrassSystem",
        "Level": "ERROR",
        "Space": 10240,
        "Type": "FileSystem"
      },
      {
        "Id": "2",
        "Component": "Lambda",
        "Level": "INFO",
```

```
    "Space": 10240,
    "Type": "FileSystem"
  },
  {
    "Id": "3",
    "Component": "Lambda",
    "Level": "INFO",
    "Type": "AWSCloudWatch"
  }
]
```

After you create a logger definition version, you can use its version ARN to create a group version before [deploying the group](#).

Logging limitations

AWS IoT Greengrass has the following logging limitations.

Transactions per second

When logging to CloudWatch is enabled, the logging component batches log events locally before sending them to CloudWatch, so you can log at a rate higher than five requests per second per log stream.

Memory

If AWS IoT Greengrass is configured to send logs to CloudWatch and a Lambda function logs more than 5 MB/second for a prolonged period of time, the internal processing pipeline eventually fills up. The theoretical worst case is 6 MB per Lambda function.

Clock skew

When logging to CloudWatch is enabled, the logging component signs requests to CloudWatch using the normal Signature Version 4 signing process. If the system time on the AWS IoT Greengrass core device is out of sync by more than [15 minutes](#), then the requests are rejected.

Disk usage

Use the following formula to calculate the total maximum amount of disk usage for logging.


```

greengrass-system-component-space * 8    // 7 if automatic IP detection is disabled
+ 128KB                                  // the internal log for the local logging
component
+ lambda-space * lambda-count         // different versions of a Lambda function are
treated as one

```

Where:

`greengrass-system-component-space`

The maximum amount of local storage for the AWS IoT Greengrass system component logs.

`lambda-space`

The maximum amount of local storage for Lambda function logs.

`lambda-count`

The number of deployed Lambda functions.

Log loss

If your AWS IoT Greengrass core device is configured to log only to CloudWatch and there's no internet connectivity, you have no way to retrieve the logs currently in the memory.

When Lambda functions are terminated (for example, during deployment), a few seconds' worth of logs are not written to CloudWatch.

CloudTrail logs

AWS IoT Greengrass runs with AWS CloudTrail, a service that provides a record of actions taken by a user, role, or an AWS service in AWS IoT Greengrass. For more information, see [the section called "Logging AWS IoT Greengrass API calls with AWS CloudTrail"](#).

Logging AWS IoT Greengrass API calls with AWS CloudTrail

AWS IoT Greengrass is integrated with AWS CloudTrail, a service that provides a record of actions taken by a user, role, or an AWS service in AWS IoT Greengrass. CloudTrail captures all API calls for AWS IoT Greengrass as events. The calls captured include calls from the AWS IoT Greengrass console and code calls to the AWS IoT Greengrass API operations. If you create a trail, you can

enable continuous delivery of CloudTrail events to an Amazon S3 bucket, including events for AWS IoT Greengrass. If you don't configure a trail, you can still view the most recent events in the CloudTrail console in **Event history**. Using the information collected by CloudTrail, you can determine the request that was made to AWS IoT Greengrass, the IP address from which the request was made, who made the request, when it was made, and additional details.

To learn more about CloudTrail, see the [AWS CloudTrail User Guide](#).

AWS IoT Greengrass information in CloudTrail

CloudTrail is enabled on your AWS account when you create the account. When activity occurs in AWS IoT Greengrass, that activity is recorded in a CloudTrail event along with other AWS service events in **Event history**. You can view, search, and download recent events in your AWS account. For more information, see [Viewing events with CloudTrail event history](#).

For an ongoing record of events in your AWS account, including events for AWS IoT Greengrass, create a trail. A *trail* enables CloudTrail to deliver log files to an Amazon S3 bucket. By default, when you create a trail in the console, the trail applies to all AWS Regions. The trail logs events from all Regions in the AWS partition and delivers the log files to the Amazon S3 bucket that you specify. Additionally, you can configure other AWS services to further analyze and act upon the event data collected in CloudTrail logs. For more information, see the following:

- [Overview for creating a trail](#)
- [CloudTrail supported services and integrations](#)
- [Configuring Amazon SNS notifications for CloudTrail](#)
- [Receiving CloudTrail log files from multiple regions](#) and [Receiving CloudTrail log files from multiple accounts](#)

All AWS IoT Greengrass actions are logged by CloudTrail and are documented in the [AWS IoT Greengrass API reference](#). For example, calls to the `AssociateServiceRoleToAccount`, `GetGroupVersion`, `GetConnectivityInfo`, and `CreateFunctionDefinition` actions generate entries in the CloudTrail log files.

Every event or log entry contains information about who generated the request. The identity information helps you determine the following:

- Whether the request was made with root or AWS Identity and Access Management (IAM) user credentials.

- Whether the request was made with temporary security credentials for a role or federated user.
- Whether the request was made by another AWS service.

For more information, see the [CloudTrail `userIdentity` element](#).

Understanding AWS IoT Greengrass log file entries

A trail is a configuration that enables delivery of events as log files to an Amazon S3 bucket that you specify. CloudTrail log files contain one or more log entries. An event represents a single request from any source and includes information about the requested action, the date and time of the action, request parameters, and so on. CloudTrail log files aren't an ordered stack trace of the public API calls, so they don't appear in any specific order.

The following example shows a CloudTrail log entry that demonstrates the `AssociateServiceRoleToAccount` action.

```
{
  "eventVersion": "1.05",
  "userIdentity": {
    "type": "IAMUser",
    "principalId": "AIDACKCEVSQ6C2EXAMPLE",
    "arn": "arn:aws:iam::123456789012:user/Mary_Major",
    "accountId": "123456789012",
    "accessKeyId": "AKIAIOSFODNN7EXAMPLE",
    "userName": "Mary_Major"
  },
  "eventTime": "2018-10-17T17:04:02Z",
  "eventSource": "greengrass.amazonaws.com",
  "eventName": "AssociateServiceRoleToAccount",
  "awsRegion": "us-east-1",
  "sourceIPAddress": "203.0.113.12",
  "userAgent": "apimanager.amazonaws.com",
  "errorCode": "BadRequestException",
  "requestParameters": null,
  "responseElements": {
    "Message": "That role ARN is invalid."
  },
  "requestID": "a5990ec6-d22e-11e8-8ae5-c7d2eEXAMPLE",
  "eventID": "b9070ce2-0238-451a-a9db-2dbf1EXAMPLE",
  "readOnly": false,
  "eventType": "AwsApiCall",
```

```
"recipientAccountId": "123456789012"
}
```

The following example shows a CloudTrail log entry that demonstrates the `GetGroupVersion` action.

```
{
  "eventVersion": "1.05",
  "userIdentity": {
    "type": "IAMUser",
    "principalId": "AIDACKCEVSQ6C2EXAMPLE",
    "arn": "arn:aws:iam::123456789012:user/Mary_Major",
    "accountId": "123456789012",
    "accessKeyId": "AKIAIOSFODNN7EXAMPLE",
    "userName": "Mary_Major",
    "sessionContext": {
      "attributes": {
        "mfaAuthenticated": "false",
        "creationDate": "2018-10-17T18:14:57Z"
      }
    },
    "invokedBy": "apimanager.amazonaws.com"
  },
  "eventTime": "2018-10-17T18:15:11Z",
  "eventSource": "greengrass.amazonaws.com",
  "eventName": "GetGroupVersion",
  "awsRegion": "us-east-1",
  "sourceIPAddress": "203.0.113.12",
  "userAgent": "apimanager.amazonaws.com",
  "requestParameters": {
    "GroupVersionId": "6c477753-dbf2-4cb8-acc3-5ba4eEXAMPLE",
    "GroupId": "90fcf6df-413c-4515-93a8-00056EXAMPLE"
  },
  "responseElements": null,
  "requestID": "95dcffce-d238-11e8-9240-a3993EXAMPLE",
  "eventID": "8a608034-82ed-431b-b5e0-87fbdEXAMPLE",
  "readOnly": true,
  "eventType": "AwsApiCall",
  "recipientAccountId": "123456789012"
}
```

The following example shows a CloudTrail log entry that demonstrates the `GetConnectivityInfo` action.

```
{
  "eventVersion": "1.05",
  "userIdentity": {
    "type": "IAMUser",
    "principalId": "AIDACKCEVSQ6C2EXAMPLE",
    "arn": "arn:aws:iam::123456789012:user/Mary_Major",
    "accountId": "123456789012",
    "accessKeyId": "AKIAIOSFODNN7EXAMPLE",
    "userName": "Mary_Major"
  },
  "eventTime": "2018-10-17T17:02:12Z",
  "eventSource": "greengrass.amazonaws.com",
  "eventName": "GetConnectivityInfo",
  "awsRegion": "us-east-1",
  "sourceIPAddress": "203.0.113.12",
  "userAgent": "apimanager.amazonaws.com",
  "requestParameters": {
    "ThingName": "us-east-1_CIS_1539795000000_"
  },
  "responseElements": null,
  "requestID": "63e3ebe3-d22e-11e8-9ddd-5baf3EXAMPLE",
  "eventID": "db2260d1-a8cc-4a65-b92a-13f65EXAMPLE",
  "readOnly": true,
  "eventType": "AwsApiCall",
  "recipientAccountId": "123456789012"
}
```

The following example shows a CloudTrail log entry that demonstrates the `CreateFunctionDefinition` action.

```
{
  "eventVersion": "1.05",
  "userIdentity": {
    "type": "IAMUser",
    "principalId": "AIDACKCEVSQ6C2EXAMPLE",
    "arn": "arn:aws:iam::123456789012:user/Mary_Major",
    "accountId": "123456789012",
    "accessKeyId": "AKIAIOSFODNN7EXAMPLE",
    "userName": "Mary_Major"
  },
  "eventTime": "2018-10-17T18:01:11Z",
  "eventSource": "greengrass.amazonaws.com",
  "eventName": "CreateFunctionDefinition",

```

```

    "awsRegion": "us-east-1",
    "sourceIPAddress": "203.0.113.12",
    "userAgent": "apimanager.amazonaws.com",
    "requestParameters": {
        "InitialVersion": "****"
    },
    "responseElements": {
        "CreationTimestamp": "2018-10-17T18:01:11.449Z",
        "LatestVersion": "dae06a61-c32c-41e9-b983-ee5cfEXAMPLE",
        "LatestVersionArn": "arn:aws:greengrass:us-east-1:123456789012:/greengrass/definition/functions/7a94847d-d4d2-406c-9796-a3529EXAMPLE/versions/dae06a61-c32c-41e9-b983-ee5cfEXAMPLE",
        "LastUpdatedTimestamp": "2018-10-17T18:01:11.449Z",
        "Id": "7a94847d-d4d2-406c-9796-a3529EXAMPLE",
        "Arn": "arn:aws:greengrass:us-east-1:123456789012:/greengrass/definition/functions/7a94847d-d4d2-406c-9796-a3529EXAMPLE"
    },
    "requestID": "a17d4b96-d236-11e8-a74e-3db27EXAMPLE",
    "eventID": "bdbf6677-a47a-4c78-b227-c5f64EXAMPLE",
    "readOnly": false,
    "eventType": "AwsApiCall",
    "recipientAccountId": "123456789012"
}

```

See also

- [What is AWS CloudTrail?](#) in the *AWS CloudTrail User Guide*
- [Creating an EventBridge rule that triggers on an AWS API call using CloudTrail](#) in the *Amazon EventBridge User Guide*
- [AWS IoT Greengrass API reference](#)

Gathering system health telemetry data from AWS IoT Greengrass core devices

System health telemetry data is diagnostic data that can help you monitor the performance of critical operations on your Greengrass core devices. The telemetry agent on the Greengrass core collects local telemetry data and publishes it to Amazon EventBridge without requiring any customer interaction. Core devices publish telemetry data to EventBridge on a best effort basis. For example, core devices might fail to deliver telemetry data while offline.

Note

Amazon EventBridge is an event bus service that you can use to connect your applications with data from a variety of sources, such as Greengrass core devices and [deployment notifications](#). For more information, see [What is Amazon EventBridge?](#) in the *Amazon EventBridge User Guide*.

You can create projects and applications to retrieve, analyze, transform, and report telemetry data from your edge devices. Domain experts, such as process engineers, can use these applications to gain insights into fleet health.

To ensure that the Greengrass edge components function properly, AWS IoT Greengrass uses the data for development and quality improvement purposes. This feature also helps inform new and enhanced edge capabilities. AWS IoT Greengrass only retains telemetry data for up to seven days.

This feature is available in AWS IoT Greengrass Core software v1.11.0 and is enabled by default for all Greengrass cores, including existing cores. You automatically start receiving data as soon as you upgrade to AWS IoT Greengrass Core software v1.11.0 or later.

For information about how to access or manage published telemetry data, see [the section called “Subscribing to receive telemetry data”](#).

The telemetry agent collects and publishes the following system metrics.

Telemetry metrics

Name	Description	Source
SystemMemUsage	The amount of memory currently in use by all applications on the Greengrass core device, including the operating system.	System
CpuUsage	The amount of CPU currently in use by all applications on the Greengrass core device, including the operating system.	System

Name	Description	Source
TotalNumberOfFDs	The number of file descriptors stored by the operating system of the Greengrass core device. One file descriptor uniquely identifies one open file.	System
LambdaOutOfMemory	The number of runs that result in the Lambda function running out of memory.	System
DroppedMessageCount	The number of dropped messages that are destined for AWS IoT Core.	GGCloudSpooler system component
LambdaTimeout	The number of timeouts for running the user-defined Lambda function.	User-defined Lambda function, AWS Cloud, and system
LambdaUngracefully Killed	The number of runs that the user-defined Lambda function fails to complete.	User-defined Lambda function, AWS Cloud, and system
LambdaError	The number of runs that result in the user-defined Lambda function writing error logs.	User-defined Lambda function, AWS Cloud, and system
BytesAppended	The number of bytes of data appended to stream manager.	GGStreamManager system component
BytesUploadedToIoT Analytics	The number of bytes of data that stream manager exports to channels in AWS IoT Analytics.	GGStreamManager system component

Name	Description	Source
BytesUploadedToKinesis	The number of bytes of data that stream manager exports to streams in Amazon Kinesis Data Streams.	GGStreamManager system component
BytesUploadedToIoTSiteWise	The number of bytes of data that stream manager exports to asset properties in AWS IoT SiteWise.	GGStreamManager system component
BytesUploadedToS3ExportTaskExecutor	The number of bytes of data that stream manager exports to objects in Amazon S3.	GGStreamManager system component
BytesUploadedToHTTP	The number of bytes of data that stream manager exports to HTTP.	GGStreamManager system component

Configuring telemetry settings

Greengrass telemetry uses the following settings:

- The telemetry agent aggregates telemetry data every hour.
- The telemetry agent publishes a telemetry message every 24 hours.

Note

The settings are unchangeable.

You can enable or disable the telemetry feature for a Greengrass core device. AWS IoT Greengrass uses [shadows](#) to manage the telemetry configuration. Your changes take effect immediately when the core has a connection to AWS IoT Core.

The telemetry agent publishes data using the MQTT protocol with a quality of service (QoS) level of 0. This means that it doesn't confirm delivery or retry publishing attempts. Telemetry messages share an MQTT connection with other messages for subscriptions destined for AWS IoT Core.

Aside from your data link costs, the data transfer from the core to AWS IoT Core is no charge. This is because the agent publishes to an AWS reserved topic. However, depending on your use case, you might incur costs when you receive or process the data.

Requirements

The following requirements apply, when you configure telemetry settings:

- You must use AWS IoT Greengrass Core software v1.11.0 or later.

Note

If you're running an earlier version and you don't want to use telemetry, you don't have to do anything.

- You must provide IAM permissions to update the core (thing) shadow and to call the configuration APIs before you update telemetry settings.

The following example IAM policy lets you manage the shadow and runtime configuration of a specific core:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "AllowManageShadow",
      "Effect": "Allow",
      "Action": [
        "iot:GetThingShadow",
        "iot:UpdateThingShadow",
        "iot:DeleteThingShadow",
        "iot:DescribeThing"
      ],
      "Resource": [
        "arn:aws:iot:region:account-id:thing/core-name-*"
      ]
    },
    {
```

```
    "Sid": "AllowManageRuntimeConfig",
    "Effect": "Allow",
    "Action": [
        "greengrass:GetCoreRuntimeConfiguration",
        "greengrass:UpdateCoreRuntimeConfiguration"
    ],
    "Resource": [
        "arn:aws:iot:region:account-id:thing/core-name"
    ]
  }
]
```

You can grant granular or conditional access to resources, for example, by using a wildcard * naming scheme. For more information, see [Adding and removing IAM policies](#) in the *IAM User Guide*.

Configure telemetry settings (console)

The following shows how to update the telemetry settings of a Greengrass core in the AWS IoT Greengrass console.

1. In the AWS IoT console navigation pane, under **Manage**, expand **Greengrass devices**, and then choose **Groups (V1)**.
2. Under **Greengrass groups**, choose your target group.
3. On the group configuration page, in the **Overview** section, choose your **Greengrass core**.
4. On the core's configuration page, choose the **Telemetry** tab.
5. In the **System health telemetry** section, choose **Configure**.
6. In **Configure telemetry**, select Telemetry to enable or disable the Telemetry status.

Important

By default, the telemetry feature is enabled for AWS IoT Greengrass Core software v1.11.0 or later.

The changes take effect at runtime. You don't need to deploy the group.

Configure telemetry settings (CLI)

In the AWS IoT Greengrass API, the `TelemetryConfiguration` object represents the telemetry settings of a Greengrass core. This object is part of the `RuntimeConfiguration` object associated with the core. You can use the AWS IoT Greengrass API, AWS CLI, or AWS SDK to manage Greengrass telemetry. The examples in this section use the AWS CLI.

To check telemetry settings

The following command gets the telemetry settings of a Greengrass core.

- Replace *core-thing-name* with the name of the target core.

To get the thing name, you use the [get-core-definition-version](#) command. The command returns the ARN of the thing that contains the thing name.

```
aws greengrass get-thing-runtime-configuration --thing-name core-thing-name
```

The command returns a `GetCoreRuntimeConfigurationResponse` object in the JSON response. For example:

```
{
  "RuntimeConfiguration": {
    "TelemetryConfiguration": {
      "ConfigurationSyncStatus": "OutOfSync",
      "Telemetry": "On"
    }
  }
}
```

To configure telemetry settings

The following command updates the telemetry settings for a Greengrass core.

- Replace *core-thing-name* with the name of the target core.

To get the thing name, you use the [get-core-definition-version](#) command. The command returns the ARN of the thing that contains the thing name.

JSON expanded

```
aws greengrass update-thing-runtime-configuration --thing-name core-thing-name --
telemetry-configuration '{
```

```
"RuntimeConfiguration": {  
  "TelemetryConfiguration": {  
    "ConfigurationSyncStatus": "InSync",  
    "Telemetry": "Off"  
  }  
}
```

JSON single-line

```
aws greengrass update-thing-runtime-configuration --thing-name core-thing-name --  
telemetry-configuration "{\"TelemetryConfiguration\":{\"ConfigurationSyncStatus  
\":\"InSync\",\"Telemetry\":\"Off\"}}"
```

Changes to telemetry settings have been applied if the `ConfigurationSyncStatus` is `InSync`. The changes take effect at runtime. You don't need to deploy the group.

TelemetryConfiguration object

The `TelemetryConfiguration` object has the following properties:

ConfigurationSyncStatus

Checks if telemetry settings are in sync. You might not make changes to this property.

Type: string

Valid values: `InSync` or `OutOfSync`

Telemetry

Turns telemetry on or off. The default is `On`.

Type: string

Valid values: `On` or `Off`

Subscribing to receive telemetry data

You can create rules in Amazon EventBridge that define how to process telemetry data published from the Greengrass core device. When EventBridge receives the data, it invokes the target actions

defined in your rules. For example, you can create event rules that send notifications, store event information, take corrective action, or invoke other events.

Telemetry event

The event for a deployment state change including the telemetry data uses the following format:

```
{
  "version": "0",
  "id": "f70f943b-9ae2-e7a5-fec4-4c22178a3e6a",
  "detail-type": "Greengrass Telemetry Data",
  "source": "aws.greengrass",
  "account": "123456789012",
  "time": "2020-07-28T20:45:53Z",
  "region": "us-west-1",
  "resources": [],
  "detail": {
    "ThingName": "CoolThing",
    "Schema": "2020-06-30",
    "ADP": [
      {
        "TS": 123231546,
        "NS": "StreamManager",
        "M": [
          {
            "N": "BytesAppended|BytesUploadedToKinesis",
            "Sum": 11,
            "U": "Bytes"
          }
        ]
      },
      {
        "TS": 123231546,
        "NS": "StreamManager",
        "M": [
          {
            "N": "BytesAppended|BytesUploadedToS3ExportTaskExecutor",
            "Sum": 11,
            "U": "Bytes"
          }
        ]
      }
    ]
  }
}
```

```

        "TS": 123231546,
        "NS": "StreamManager",
        "M": [
            {
                "N": "BytesAppended|BytesUploadedToHTTP",
                "Sum": 11,
                "U": "Bytes"
            }
        ]
    },
    {
        "TS": 123231546,
        "NS": "StreamManager",
        "M": [
            {
                "N": "BytesAppended|BytesUploadedToIoTAnalytics",
                "Sum": 11,
                "U": "Bytes"
            }
        ]
    },
    {
        "TS": 123231546,
        "NS": "StreamManager",
        "M": [
            {
                "N": "BytesAppended|BytesUploadedToIoTSiteWise",
                "Sum": 11,
                "U": "Bytes"
            }
        ]
    },
    {
        "TS": 123231546,
        "NS": "arn:aws:lambda:us-west-1:123456789012:function:my-function",
        "M": [
            {
                "N": "LambdaTimeout",
                "Sum": 15,
                "U": "Count"
            }
        ]
    },
    {

```

```
"TS": 123231546,
"NS": "CloudSpooler",
"M": [
  {
    "N": "DroppedMessageCount",
    "Sum": 15,
    "U": "Count"
  }
],
{
  "TS": 1593727692,
  "NS": "SystemMetrics",
  "M": [
    {
      "N": "SystemMemUsage",
      "Sum": 11.23,
      "U": "Megabytes"
    },
    {
      "N": "CpuUsage",
      "Sum": 35.63,
      "U": "Percent"
    },
    {
      "N": "TotalNumberOfFDs",
      "Sum": 416,
      "U": "Count"
    }
  ]
},
{
  "TS": 1593727692,
  "NS": "arn:aws:lambda:us-west-1:123456789012:function:my-function",
  "M": [
    {
      "N": "LambdaOutOfMemory",
      "Sum": 12,
      "U": "Count"
    },
    {
      "N": "LambdaUngracefullyKilled",
      "Sum": 100,
      "U": "Count"
    }
  ]
}
```



```
    },  
    {  
      "N": "LambdaError",  
      "Sum": 7,  
      "U": "Count"  
    }  
  ]  
}  
}
```

The ADP array contains a list of aggregated data points that have the following properties:

TS

Required. The timestamp of when the data was aggregated.

NS

Required. The namespace of the system.

M

Required. The list of metrics. A metric contains the following properties:

N

The name of the [metric](#).

Sum

The aggregated metric value. The telemetry agent adds new values to the previous total, so the sum is an ever-increasing value. You can use the timestamp to find the value of a specific aggregation. For example, to find the latest aggregated value, subtract the previous timestamped value from the latest timestamped value.

U

The unit of the metric value.

ThingName

Required. The name of the thing device that you target.

Prerequisites for creating EventBridge rules

Before you create an EventBridge rule for AWS IoT Greengrass, you should do the following:

- Familiarize yourself with events, rules, and targets in EventBridge.
- Create and configure the [targets](#) invoked by your EventBridge rules. Rules can invoke many types of targets, such as Amazon Kinesis streams, AWS Lambda functions, Amazon SNS topics, and Amazon SQS queues.

Your EventBridge rule, and the associated targets must be in the AWS Region where you created your Greengrass resources. For more information, see [Service endpoints and quotas](#) in the *AWS General Reference*.

For more information, see [What is Amazon EventBridge?](#) and [Getting started with Amazon EventBridge](#) in the *Amazon EventBridge User Guide*.

Create an event rule to get telemetry data (console)

Use the following steps to use the AWS Management Console to create an EventBridge rule that receives telemetry data published by the Greengrass core. This allows web servers, email addresses, and other topic subscribers to respond to the event. For more information, see [Creating a EventBridge rule that triggers on an event from an AWS resource](#) in the *Amazon EventBridge User Guide*.

1. Open the [Amazon EventBridge console](#) and choose **Create rule**.
2. Under **Name and description**, enter a name and description for the rule.
3. Choose **Event bus-** and enable the rule on the selected event bus..
4. Select the **Rule type** and choose **Rule with an event pattern**.
5. Choose **Next**.
6. For **Event source**, choose **AWS events or EventBridge partner events**.
7. For **Sample event**, choose **AWS events**, and select **Greengrass Telemetry Data**.
8. In **Event pattern**, make the following selections:
 - a. For **Event source**, choose **AWS services**.
 - b. For **AWS service**, choose **Greengrass**.
 - c. For **Event type**, choose **Greengrass Telemetry Data**.

9. Choose **Next**.
10. For **Target 1**, choose **AWS service**.
11. For **Select a target**, choose **SQS queue**.
12. For **Queue**, choose your function.

Create an event rule to get telemetry data (CLI)

Use the following steps to use the AWS CLI to create an EventBridge rule that receives telemetry data published by the Greengrass core. This allows web servers, email addresses, and other topic subscribers to respond to the event.

1. Create the rule.
 - Replace *thing-name* with the thing name of the core.

To get the thing name, you use the [get-core-definition-version](#) command. The command returns the ARN of the thing that contains the thing name.

```
aws events put-rule \  
  --name TestRule \  
  --event-pattern "{\"source\": [\"aws.greengrass\"], \"detail\": {\"ThingName\":  
    [\"thing-name\"]}}\"
```

Properties that are omitted from the pattern are ignored.

2. Add the topic as a rule target. The following example uses Amazon SQS but you can configure other target types.
 - Replace *queue-arn* with the ARN of your Amazon SQS queue.

```
aws events put-targets \  
  --rule TestRule \  
  --targets "Id"="1", "Arn"="queue-arn"
```

Note

To allow Amazon EventBridge to invoke your target queue, you must add a resource-based policy to your topic. For more information, see [Amazon SQS permissions](#) in the *Amazon EventBridge User Guide*.

For more information, see [Events and event patterns in EventBridge](#) in the *Amazon EventBridge User Guide*.

Troubleshooting AWS IoT Greengrass telemetry

Use the following information to help troubleshoot issues with configuring AWS IoT Greengrass telemetry.

Error: The response contains "ConfigurationStatus": "OutOfSync" after you run the get-thing-runtime-configuration command

Solutions:

- The AWS IoT Device Shadow service takes time to process runtime configuration updates and to deliver the updates to the Greengrass core device. You might wait and check if telemetry settings are in sync later.
- Make sure that your core device is online.
- Enable [Amazon CloudWatch Logs in AWS IoT Core](#) to monitor the shadow.
- Use [AWS IoT metrics](#) to monitor your thing.

Calling the local health check API

AWS IoT Greengrass contains a local HTTP API that provides a snapshot of the current state of local worker processes that were started by AWS IoT Greengrass. This snapshot includes user-defined Lambda functions and system Lambda functions. System Lambda functions are part of the AWS IoT Greengrass Core software. They run as local worker processes on the core device and manage operations such as message routing, local shadow sync, and automatic IP address detection.

The health check API supports the following requests:

- Send a GET request to [get health information for all workers](#).
- Send a POST request to [get health information for specified workers](#).

Requests are sent locally on the device and don't require an internet connection.

Get health information for all workers

Send a GET request to get health information about all running workers.

- Replace *port* with the port number of the IPC.

```
GET http://localhost:port/2016-11-01/health/workers
```

port

The port number of the IPC.

The value can vary between 1024 and 65535. The default value is 8000.

To change this port number, you can update the `ggDaemonPort` property in the `config.json` file. For more information, see [AWS IoT Greengrass core configuration file](#).

Example request

The following example `curl` request gets health information for all workers.

```
curl http://localhost:8000/2016-11-01/health/workers
```

JSON Response

This request returns an array of [worker health information](#) objects.

Example response

The following example response lists health information objects for all worker processes that were started by AWS IoT Greengrass.

```
[  
  {
```

```
[
  {
    "FuncArn": "arn:aws:lambda::function:GGShadowService:1",
    "WorkerId": "65515053-2f70-43dc-7cc0-1712bEXAMPLE",
    "ProcessId": "1234",
    "WorkerState": "Waiting"
  },
  {
    "FuncArn": "arn:aws:lambda::function:GGSecretManager:1",
    "WorkerId": "a9916cc2-1b4d-4f0e-4b12-b1872EXAMPLE",
    "ProcessId": "9798",
    "WorkerState": "Waiting"
  },
  {
    "FuncArn": "arn:aws:lambda:us-west-2:123456789012:function:my-lambda-function:3",
    "WorkerId": "2e6f785e-66a5-42c9-67df-42073EXAMPLE",
    "ProcessId": "11837",
    "WorkerState": "Waiting"
  },
  ...
]
```

Get health information about specified workers

Send a POST request to get health information about specified workers. Replace *port* with the port number of the IPC. The default is 8000.

```
POST http://localhost:port/2016-11-01/health/workers
```

Example request

The following example `curl` request gets health information for specified workers.

```
curl --data "@body.json" http://localhost:8000/2016-11-01/health/workers
```

Here's an example `body.json` request body:

```
{
  "FuncArns": [
    "arn:aws:lambda::function:GGShadowService:1",
    "arn:aws:lambda:us-west-2:123456789012:function:my-lambda-function:3"
  ]
}
```

The request body contains a `FuncArns` array.

FuncArns

A list of Amazon Resource Names (ARNs) for the Lambda functions that represent the target workers.

- For user-defined Lambda functions, specify the ARN of the currently deployed version. If you added Lambda functions to the group using an alias ARN, you can use the GET request to get all workers, and then choose the ARNs you want to query for.
- For system Lambda functions, specify the ARN of the corresponding Lambda function. For more information, see [the section called "System Lambda functions"](#).

Type: array of strings

Minimum length: 1

Maximum length: The total number of workers started by AWS IoT Greengrass on the core device.

JSON Response

This request returns a `Workers` array and an `InvalidArns` array.

Workers

A list of health information objects for the specified workers.

Type: array of [health information objects](#)

InvalidArns

A list of function ARNs that are invalid, including function ARNs that don't have associated workers.

Type: array of strings

Example response

The following example response lists [health information objects](#) for the specified workers.

```
{
```

```

    "Workers": [
      {
        "FuncArn": "arn:aws:lambda::function:GGShadowService:1",
        "WorkerId" : "65515053-2f70-43dc-7cc0-1712bEXAMPLE",
        "ProcessId": "1234",
        "WorkerState": "Waiting"
      },
      {
        "FuncArn": "arn:aws:lambda:us-west-2:123456789012:function:my-lambda-
function:3",
        "WorkerId": "2e6f785e-66a5-42c9-67df-42073ESAMPLE",
        "ProcessId": "11837",
        "WorkerState": "Waiting"
      }
    ],
    "InvalidArns" : [
      "some-malformed-arn",
      "arn:aws:lambda:us-west-2:123456789012:function:some-unknown-function:1"
    ]
  }

```

This request returns the following errors:

400 Invalid request

The request body is malformed. To resolve this issue, use the following format and resend the request:

```

{"FuncArns":["function-1-arn","function-2-arn"]}

```

400 Request exceeds max number of workers

The number of ARNs specified in the FuncArns array exceeds the number of workers.

Worker health information

A health information object contains the following properties:

FuncArn

The ARN of the system Lambda function that represents the worker.

Type: string

WorkerId

The ID of the worker. This property can be useful for debugging. The `runtime.log` file and the Lambda function logs contain the worker ID, so this property can be especially useful to debug an on-demand Lambda function that spins up multiple instances.

Type: string

ProcessId

The process ID (PID) of the worker process.

Type: int

WorkerState

The state of the worker.

Type: string

The following are possible worker states:

Working

Processing a message.

Waiting

Waiting for a message. Applies to long-lived Lambda functions running as a daemon or standalone process.

Starting

Spun up, getting started.

FailedInitialization

Failed to initialize.

Terminated

Stopped by the Greengrass daemon

NotStarted

Failed to start, making another start attempt.

Initialized

Successfully initialized.

System Lambda functions

You can request health information for the following system Lambda functions:

GGCloudSpooler

Manages the queue for MQTT messages that have AWS IoT Core as the source or target.

ARN: `arn:aws:lambda:::function:GGCloudSpooler:1`

GGConnManager

Routes MQTT messages between the Greengrass core and client devices.

ARN: `arn:aws:lambda:::function:GGConnManager`

GGDeviceCertificateManager

Listens to the AWS IoT shadow for changes to the core's IP endpoints and generates the server-side certificate used by GGConnManager for mutual authentication.

ARN: `arn:aws:lambda:::function:GGDeviceCertificateManager`

GGIPDetector

Manages automatic IP address detection that enables devices in the Greengrass group to discover the Greengrass core device. This service isn't applicable when you provide IP addresses manually.

ARN: `arn:aws:lambda:::function:GGIPDetector:1`

GGSecretManager

Manages secure storage of local secrets and access by user-defined Lambda and connectors.

ARN: `arn:aws:lambda:::function:GGSecretManager:1`

GGShadowService

Manages local shadows for client devices.

ARN: `arn:aws:lambda:::function:GGShadowService`

GGShadowSyncManager

Synchronizes local shadows with the AWS Cloud for the core device and client devices, if the device's `syncShadow` property is set to `true`.

ARN: `arn:aws:lambda:::function:GGShadowSyncManager`

GGStreamManager

Processes data streams locally and performs automatic exports to the AWS Cloud.

ARN: `arn:aws:lambda:::function:GGStreamManager:1`

GGTES

The local token exchange service that retrieves IAM credentials defined in the Greengrass group role that local code uses to access AWS services.

ARN: `arn:aws:lambda:::function:GGTES`

Tagging your AWS IoT Greengrass resources

Tags can help you organize and manage your AWS IoT Greengrass groups. You can use tags to assign metadata to groups, bulk deployments, and the cores, devices, and other resources that are added to groups. Tags can also be used in IAM policies to define conditional access to your Greengrass resources.

Note

Currently, Greengrass resource tags are not supported for AWS IoT billing groups or cost allocation reports.

Tag basics

Tags allow you to categorize your AWS IoT Greengrass resources, for example, by purpose, owner, and environment. When you have many resources of the same type, you can quickly identify a resource based on the tags that are attached to it. A tag consists of a key and optional value, both of which you define. We recommend that you design a set of tag keys for each resource type. Using a consistent set of tag keys makes it easier for you to manage your resources. For example, you can define a set of tags for your groups that helps you track the factory location of your core devices. For more information, see [AWS Tagging Strategies](#).

Tagging support in the AWS IoT console

You can create, view, and manage tags for your Greengrass Group resources in the AWS IoT console. Before you create tags, be aware of tagging restrictions. For more information, see [Tag naming and usage conventions](#) in the *Amazon Web Services General Reference*.

To assign tags when you create a group

You can assign tags to a group when you create the group. Choose **Add new tag** under the **Tags** section to show the tagging input fields.

To view and manage tags from the group configuration page

You can view and manage tags from the group configuration page by choosing **View settings**. In the **Tags** section for the group, choose **Manage tags** to add, edit, or remove group tags.

Tagging support in the AWS IoT Greengrass API

You can use the AWS IoT Greengrass API to create, list, and manage tags for AWS IoT Greengrass resources that support tagging. Before you create tags, be aware of tagging restrictions. For more information, see [Tag naming and usage conventions](#) in the *Amazon Web Services General Reference*.

- To add tags during resource creation, define them in the tags property of the resource.
- To add tags after a resource is created, or to update tag values, use the TagResource action.
- To remove tags from a resource, use the UntagResource action.
- To retrieve the tags that are associated with a resource, use the ListTagsForResource action or get the resource and inspect its tags property.

The following table lists resources you can tag in the AWS IoT Greengrass API and their corresponding Create and Get actions.

Resource	Create	Get
Group	CreateGroup	GetGroup
ConnectorDefinition	CreateConnectorDefinition	GetConnectorDefinition
CoreDefinition	CreateCoreDefinition	GetCoreDefinition
DeviceDefinition	CreateDeviceDefinition	GetDeviceDefinition
FunctionDefinition	CreateFunctionDefinition	GetFunctionDefinition
LoggerDefinition	CreateLoggerDefinition	GetLoggerDefinition
ResourceDefinition	CreateResourceDefinition	GetResourceDefinition
SubscriptionDefinition	CreateSubscriptionDefinition	GetSubscriptionDefinition

Resource	Create	Get
BulkDeployment	StartBulkDeployment	GetBulkDeploymentsStatus

Use the following actions to list and manage tags for resources that support tagging:

- [TagResource](#). Adds tags to a resource. Also used to change the value of the tag's key-value pair.
- [ListTagsForResource](#). Lists the tags for a resource.
- [UntagResource](#). Removes tags from a resource.

You can add or remove tags on a resource at any time. To change the value of a tag key, add a tag to the resource that defines the same key and the new value. The new value overwrites the old value. You can set a value to an empty string, but you can't set a value to null.

When you delete a resource, tags that are associated with the resource are also deleted.

Note

Don't confuse resource tags with the attributes that you can assign to AWS IoT things. Although Greengrass cores are AWS IoT things, the resource tags that are described in this topic are attached to a `CoreDefinition`, not the core thing.

Using tags with IAM policies

In your IAM policies, you can use resource tags to control user access and permissions. For example, policies can allow users to create only those resources that have a specific tag. Policies can also restrict users from creating or modifying resources that have certain tags. You can tag resources during creation (called *tag on create*) so you don't have to run custom tagging scripts later. When new environments are launched with tags, the corresponding IAM permissions are applied automatically.

The following condition context keys and values can be used in the `Condition` element (also called the `Condition` block) of the policy.

`greengrass:ResourceTag/tag-key: tag-value`

Allow or deny user actions on resources with specific tags.

`aws:RequestTag/tag-key: tag-value`

Require that a specific tag be used (or not used) when making API requests to create or modify tags on a taggable resource.

`aws:TagKeys: [tag-key, ...]`

Require that a specific set of tag keys be used (or not used) when making an API request to create or modify a taggable resource.

Condition context keys and values can be used only on AWS IoT Greengrass actions that act on a taggable resource. These actions take the resource as a required parameter. For example, you can set conditional access on the `GetGroupVersion`. You can't set conditional access on `AssociateServiceRoleToAccount` because no taggable resource (for example, group, core definition, or device definition) is referenced in the request.

For more information, see [Controlling access using tags](#) and [IAM JSON policy reference](#) in the *IAM User Guide*. The JSON policy reference includes detailed syntax, descriptions and examples of the elements, variables, and evaluation logic of JSON policies in IAM.

Example IAM policies

The following example policy applies tag-based permissions that constrain a beta user to actions on beta resources only.

- The first statement allows an IAM user to act on resources that have the `env=beta` tag only.
- The second statement prevents an IAM user from removing the `env=beta` tag from resources. This protects the user from removing their own access.

Note

If you use tags to control access to resources, you should also manage the permissions that allow users to add tags or remove tags from those same resources. Otherwise, in some cases, it might be possible for users to circumvent your restrictions and gain access to a resource by modifying its tags.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": "greengrass:*",
      "Resource": "*",
      "Condition": {
        "StringEquals": {
          "greengrass:ResourceTag/env": "beta"
        }
      }
    },
    {
      "Effect": "Deny",
      "Action": "greengrass:UntagResource",
      "Resource": "*",
      "Condition": {
        "StringEquals": {
          "aws:ResourceTag/env": "beta"
        }
      }
    }
  ]
}
```

To allow users to tag on create, you must give them appropriate permissions. The following example policy includes the `"aws:RequestTag/env": "beta"` condition on the `greengrass:TagResource` and `greengrass:CreateGroup` actions, which allows users to create a group only if they tag the group with `env=beta`. This effectively forces users to tag new groups.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": "greengrass:TagResource",
      "Resource": "*",
      "Condition": {
        "StringEquals": {
          "aws:RequestTag/env": "beta"
        }
      }
    }
  ]
}
```



```
        }
      },
    {
      "Effect": "Allow",
      "Action": "greengrass:CreateGroup",
      "Resource": "*",
      "Condition": {
        "StringEquals": {
          "aws:RequestTag/env": "beta"
        }
      }
    }
  ]
}
```

The following snippet shows how you can specify multiple tag values for a tag key by enclosing them in a list:

```
"StringEquals" : {
  "greengrass:ResourceTag/env" : ["dev", "test"]
}
```

See also

- [Tagging AWS resources](#) in the *Amazon Web Services General Reference*

AWS CloudFormation support for AWS IoT Greengrass

AWS CloudFormation is a service that can help you create, manage, and replicate your AWS resources. You can use AWS CloudFormation templates to define AWS IoT Greengrass groups and the client devices, subscriptions, and other components that you want to deploy. For an example, see [the section called “Example template”](#).

The resources and infrastructure that you generate from a template is called a *stack*. You can define all of your resources in one template or refer to resources from other stacks. For more information about AWS CloudFormation templates and features, see [What is AWS CloudFormation?](#) in the *AWS CloudFormation User Guide*.

Creating resources

AWS CloudFormation templates are JSON or YAML documents that describe the properties and relationships of AWS resources. The following AWS IoT Greengrass resources are supported:

- Groups
- Cores
- Client devices (devices)
- Lambda functions
- Connectors
- Resources (local, machine learning, and secret)
- Subscriptions
- Loggers (logging configurations)

In AWS CloudFormation templates, the structure and syntax of Greengrass resources are based on the AWS IoT Greengrass API. For example, the [example template](#) associates a top-level DeviceDefinition with a DeviceDefinitionVersion that contains an individual client device. For more information, see [the section called “Overview of the group object model”](#).

The [AWS IoT Greengrass resource types reference](#) in the *AWS CloudFormation User Guide* describes the Greengrass resources that you can manage with AWS CloudFormation. When you use AWS CloudFormation templates to create Greengrass resources, we recommend that you manage them

only from AWS CloudFormation. For example, you should update your template if you want to add, change, or remove a device (instead of using the AWS IoT Greengrass API or AWS IoT console). This allows you to use rollback and other AWS CloudFormation change management features. For more information about using AWS CloudFormation to create and manage your resources and stacks, see [Working with stacks](#) in the *AWS CloudFormation User Guide*.

For a walkthrough that shows how to create and deploy AWS IoT Greengrass resources in an AWS CloudFormation template, see [Automating AWS IoT Greengrass setup with AWS CloudFormation](#) on The Internet of Things on AWS Official Blog.

Deploying resources

After you create an AWS CloudFormation stack that contains your group version, you can use the AWS CLI or AWS IoT console to deploy it.

Note

To deploy a group, you must have a Greengrass service role associated with your AWS account. The service role allows AWS IoT Greengrass to access your resources in AWS Lambda and other AWS services. This role should exist if you already deployed a Greengrass group in the current AWS Region. For more information, see [the section called “Greengrass service role”](#).

To deploy the group (AWS CLI)

- Run the [create-deployment](#) command.

```
aws greengrass create-deployment --group-id GroupId --group-version-id GroupVersionId --deployment-type NewDeployment
```

Note

The `CommandToDeployGroup` statement in the [example template](#) shows how to output the command with your group and group version IDs when you create a stack.

To deploy the group (console)

1. In the AWS IoT console navigation pane, under **Manage**, expand **Greengrass devices**, and then choose **Groups (V1)**.
2. Choose your group.
3. On the group configuration page, choose **Deploy**.

Example template

The following example template creates a Greengrass group that contains a core, client device, function, logger, subscription, and two resources. To do this, the template follows the object model of the AWS IoT Greengrass API. For example, the client devices that you want to add to the group are contained in a `DeviceDefinitionVersion` resource, which is associated with a `DeviceDefinition` resource. To add the devices to the group, the group version references the ARN of the `DeviceDefinitionVersion`.

The template includes parameters that let you specify the certificate ARNs for the core and device and the version ARN of the source Lambda function (which is an AWS Lambda resource). It uses the `Ref` and `GetAtt` intrinsic functions to reference IDs, ARNs, and other attributes that are required to create Greengrass resources.

The template also defines two AWS IoT devices (things), which represent the core and client device that are added to the Greengrass group.

After you create the stack with your Greengrass resources, you can use the AWS CLI or the AWS IoT console to [deploy the group](#).

Note

The `CommandToDeployGroup` statement in the example shows how to output a complete **create-deployment** CLI command that you can use to deploy your group.

JSON

```
{
  "AWSTemplateFormatVersion": "2010-09-09",
  "Description": "AWS IoT Greengrass example template that creates a group version
with a core, device, function, logger, subscription, and resources.",
```

```

"Parameters": {
  "CoreCertificateArn": {
    "Type": "String"
  },
  "DeviceCertificateArn": {
    "Type": "String"
  },
  "LambdaVersionArn": {
    "Type": "String"
  }
},
"Resources": {
  "TestCore1": {
    "Type": "AWS::IoT::Thing",
    "Properties": {
      "ThingName": "TestCore1"
    }
  },
  "TestCoreDefinition": {
    "Type": "AWS::Greengrass::CoreDefinition",
    "Properties": {
      "Name": "DemoTestCoreDefinition"
    }
  },
  "TestCoreDefinitionVersion": {
    "Type": "AWS::Greengrass::CoreDefinitionVersion",
    "Properties": {
      "CoreDefinitionId": {
        "Ref": "TestCoreDefinition"
      },
      "Cores": [
        {
          "Id": "TestCore1",
          "CertificateArn": {
            "Ref": "CoreCertificateArn"
          },
          "SyncShadow": "false",
          "ThingArn": {
            "Fn::Join": [
              ":",
              [
                "arn:aws:iot",
                {
                  "Ref": "AWS::Region"
                }
              ]
            ]
          }
        }
      ]
    }
  }
}

```

```

        },
        {
            "Ref": "AWS::AccountId"
        },
        "thing/TestCore1"
    ]
]
}
}
},
"TestClientDevice1": {
    "Type": "AWS::IoT::Thing",
    "Properties": {
        "ThingName": "TestClientDevice1"
    }
},
"TestDeviceDefinition": {
    "Type": "AWS::Greengrass::DeviceDefinition",
    "Properties": {
        "Name": "DemoTestDeviceDefinition"
    }
},
"TestDeviceDefinitionVersion": {
    "Type": "AWS::Greengrass::DeviceDefinitionVersion",
    "Properties": {
        "DeviceDefinitionId": {
            "Fn::GetAtt": [
                "TestDeviceDefinition",
                "Id"
            ]
        },
        "Devices": [
            {
                "Id": "TestClientDevice1",
                "CertificateArn": {
                    "Ref": "DeviceCertificateArn"
                },
                "SyncShadow": "true",
                "ThingArn": {
                    "Fn::Join": [
                        ":",
                        [

```

```

        "arn:aws:iot",
        {
            "Ref": "AWS::Region"
        },
        {
            "Ref": "AWS::AccountId"
        },
        "thing/TestClientDevice1"
    ]
}
}
}
}
},
"TestFunctionDefinition": {
    "Type": "AWS::Greengrass::FunctionDefinition",
    "Properties": {
        "Name": "DemoTestFunctionDefinition"
    }
},
"TestFunctionDefinitionVersion": {
    "Type": "AWS::Greengrass::FunctionDefinitionVersion",
    "Properties": {
        "FunctionDefinitionId": {
            "Fn::GetAtt": [
                "TestFunctionDefinition",
                "Id"
            ]
        },
        "DefaultConfig": {
            "Execution": {
                "IsolationMode": "GreengrassContainer"
            }
        },
        "Functions": [
            {
                "Id": "TestLambda1",
                "FunctionArn": {
                    "Ref": "LambdaVersionArn"
                },
                "FunctionConfiguration": {
                    "Pinned": "true",
                    "Executable": "run.exe",

```

```

        "ExecArgs": "argument1",
        "MemorySize": "512",
        "Timeout": "2000",
        "EncodingType": "binary",
        "Environment": {
            "Variables": {
                "variable1": "value1"
            },
            "ResourceAccessPolicies": [
                {
                    "ResourceId": "ResourceId1",
                    "Permission": "ro"
                },
                {
                    "ResourceId": "ResourceId2",
                    "Permission": "rw"
                }
            ],
            "AccessSysfs": "false",
            "Execution": {
                "IsolationMode": "GreengrassContainer",
                "RunAs": {
                    "Uid": "1",
                    "Gid": "10"
                }
            }
        }
    },
    "TestLoggerDefinition": {
        "Type": "AWS::Greengrass::LoggerDefinition",
        "Properties": {
            "Name": "DemoTestLoggerDefinition"
        }
    },
    "TestLoggerDefinitionVersion": {
        "Type": "AWS::Greengrass::LoggerDefinitionVersion",
        "Properties": {
            "LoggerDefinitionId": {
                "Ref": "TestLoggerDefinition"
            }
        }
    },

```



```

        "Loggers": [
            {
                "Id": "TestLogger1",
                "Type": "AWS::CloudWatch",
                "Component": "GreengrassSystem",
                "Level": "INFO"
            }
        ]
    },
    "TestResourceDefinition": {
        "Type": "AWS::Greengrass::ResourceDefinition",
        "Properties": {
            "Name": "DemoTestResourceDefinition"
        }
    },
    "TestResourceDefinitionVersion": {
        "Type": "AWS::Greengrass::ResourceDefinitionVersion",
        "Properties": {
            "ResourceDefinitionId": {
                "Ref": "TestResourceDefinition"
            },
            "Resources": [
                {
                    "Id": "ResourceId1",
                    "Name": "LocalDeviceResource",
                    "ResourceDataContainer": {
                        "LocalDeviceResourceData": {
                            "SourcePath": "/dev/TestSourcePath1",
                            "GroupOwnerSetting": {
                                "AutoAddGroupOwner": "false",
                                "GroupOwner": "TestOwner"
                            }
                        }
                    }
                },
                {
                    "Id": "ResourceId2",
                    "Name": "LocalVolumeResourceData",
                    "ResourceDataContainer": {
                        "LocalVolumeResourceData": {
                            "SourcePath": "/dev/TestSourcePath2",
                            "DestinationPath": "/volumes/TestDestinationPath2",
                            "GroupOwnerSetting": {

```

```

        "AutoAddGroupOwner": "false",
        "GroupOwner": "TestOwner"
    }
}
}
}
]
}
},
"TestSubscriptionDefinition": {
    "Type": "AWS::Greengrass::SubscriptionDefinition",
    "Properties": {
        "Name": "DemoTestSubscriptionDefinition"
    }
},
"TestSubscriptionDefinitionVersion": {
    "Type": "AWS::Greengrass::SubscriptionDefinitionVersion",
    "Properties": {
        "SubscriptionDefinitionId": {
            "Ref": "TestSubscriptionDefinition"
        },
        "Subscriptions": [
            {
                "Id": "TestSubscription1",
                "Source": {
                    "Fn::Join": [
                        ":",
                        [
                            "arn:aws:iot",
                            {
                                "Ref": "AWS::Region"
                            },
                            {
                                "Ref": "AWS::AccountId"
                            },
                            "thing/TestClientDevice1"
                        ]
                    ]
                },
                "Subject": "TestSubjectUpdated",
                "Target": {
                    "Ref": "LambdaVersionArn"
                }
            }
        ]
    }
}
}

```

```

    ]
  },
  "TestGroup": {
    "Type": "AWS::Greengrass::Group",
    "Properties": {
      "Name": "DemoTestGroupNewName",
      "RoleArn": {
        "Fn::Join": [
          ":",
          [
            "arn:aws:iam:",
            {
              "Ref": "AWS::AccountId"
            },
            "role/TestUser"
          ]
        ]
      }
    },
    "InitialVersion": {
      "CoreDefinitionVersionArn": {
        "Ref": "TestCoreDefinitionVersion"
      },
      "DeviceDefinitionVersionArn": {
        "Ref": "TestDeviceDefinitionVersion"
      },
      "FunctionDefinitionVersionArn": {
        "Ref": "TestFunctionDefinitionVersion"
      },
      "SubscriptionDefinitionVersionArn": {
        "Ref": "TestSubscriptionDefinitionVersion"
      },
      "LoggerDefinitionVersionArn": {
        "Ref": "TestLoggerDefinitionVersion"
      },
      "ResourceDefinitionVersionArn": {
        "Ref": "TestResourceDefinitionVersion"
      }
    },
    "Tags": {
      "KeyName0": "value",
      "KeyName1": "value",
      "KeyName2": "value"
    }
  }
}

```

```

    }
  },
  "Outputs": {
    "CommandToDeployGroup": {
      "Value": {
        "Fn::Join": [
          " ",
          [
            "groupVersion=$(cut -d'/' -f6 <<<",
            {
              "Fn::GetAtt": [
                "TestGroup",
                "LatestVersionArn"
              ]
            },
            ");",
            "aws --region",
            {
              "Ref": "AWS::Region"
            },
            "greengrass create-deployment --group-id",
            {
              "Ref": "TestGroup"
            },
            "--deployment-type NewDeployment --group-version-id",
            "$groupVersion"
          ]
        ]
      }
    }
  }
}

```

YAML

AWSTemplateFormatVersion: 2010-09-09

Description: >-

AWS IoT Greengrass example template that creates a group version with a core, device, function, logger, subscription, and resources.

Parameters:

CoreCertificateArn:

Type: String

```

DeviceCertificateArn:
  Type: String
LambdaVersionArn:
  Type: String
Resources:
  TestCore1:
    Type: 'AWS::IoT::Thing'
    Properties:
      ThingName: TestCore1
  TestCoreDefinition:
    Type: 'AWS::Greengrass::CoreDefinition'
    Properties:
      Name: DemoTestCoreDefinition
  TestCoreDefinitionVersion:
    Type: 'AWS::Greengrass::CoreDefinitionVersion'
    Properties:
      CoreDefinitionId: !Ref TestCoreDefinition
      Cores:
        - Id: TestCore1
          CertificateArn: !Ref CoreCertificateArn
          SyncShadow: 'false'
          ThingArn: !Join
            - ':'
            - - 'arn:aws:iot'
              - !Ref 'AWS::Region'
              - !Ref 'AWS::AccountId'
              - thing/TestCore1
  TestClientDevice1:
    Type: 'AWS::IoT::Thing'
    Properties:
      ThingName: TestClientDevice1
  TestDeviceDefinition:
    Type: 'AWS::Greengrass::DeviceDefinition'
    Properties:
      Name: DemoTestDeviceDefinition
  TestDeviceDefinitionVersion:
    Type: 'AWS::Greengrass::DeviceDefinitionVersion'
    Properties:
      DeviceDefinitionId: !GetAtt
        - TestDeviceDefinition
        - Id
      Devices:
        - Id: TestClientDevice1
          CertificateArn: !Ref DeviceCertificateArn

```

```

    SyncShadow: 'true'
    ThingArn: !Join
      - ':'
      - - 'arn:aws:iot'
        - !Ref 'AWS::Region'
        - !Ref 'AWS::AccountId'
        - thing/TestClientDevice1
  TestFunctionDefinition:
    Type: 'AWS::Greengrass::FunctionDefinition'
    Properties:
      Name: DemoTestFunctionDefinition
  TestFunctionDefinitionVersion:
    Type: 'AWS::Greengrass::FunctionDefinitionVersion'
    Properties:
      FunctionDefinitionId: !GetAtt
        - TestFunctionDefinition
        - Id
      DefaultConfig:
        Execution:
          IsolationMode: GreengrassContainer
      Functions:
        - Id: TestLambda1
          FunctionArn: !Ref LambdaVersionArn
          FunctionConfiguration:
            Pinned: 'true'
            Executable: run.exe
            ExecArgs: argument1
            MemorySize: '512'
            Timeout: '2000'
            EncodingType: binary
            Environment:
              Variables:
                variable1: value1
            ResourceAccessPolicies:
              - ResourceId: ResourceId1
                Permission: ro
              - ResourceId: ResourceId2
                Permission: rw
            AccessSysfs: 'false'
            Execution:
              IsolationMode: GreengrassContainer
              RunAs:
                Uid: '1'
                Gid: '10'

```

```
TestLoggerDefinition:
  Type: 'AWS::Greengrass::LoggerDefinition'
  Properties:
    Name: DemoTestLoggerDefinition
TestLoggerDefinitionVersion:
  Type: 'AWS::Greengrass::LoggerDefinitionVersion'
  Properties:
    LoggerDefinitionId: !Ref TestLoggerDefinition
    Loggers:
      - Id: TestLogger1
        Type: AWSCloudWatch
        Component: GreengrassSystem
        Level: INFO
TestResourceDefinition:
  Type: 'AWS::Greengrass::ResourceDefinition'
  Properties:
    Name: DemoTestResourceDefinition
TestResourceDefinitionVersion:
  Type: 'AWS::Greengrass::ResourceDefinitionVersion'
  Properties:
    ResourceDefinitionId: !Ref TestResourceDefinition
    Resources:
      - Id: ResourceId1
        Name: LocalDeviceResource
        ResourceDataContainer:
          LocalDeviceResourceData:
            SourcePath: /dev/TestSourcePath1
            GroupOwnerSetting:
              AutoAddGroupOwner: 'false'
              GroupOwner: TestOwner
      - Id: ResourceId2
        Name: LocalVolumeResourceData
        ResourceDataContainer:
          LocalVolumeResourceData:
            SourcePath: /dev/TestSourcePath2
            DestinationPath: /volumes/TestDestinationPath2
            GroupOwnerSetting:
              AutoAddGroupOwner: 'false'
              GroupOwner: TestOwner
TestSubscriptionDefinition:
  Type: 'AWS::Greengrass::SubscriptionDefinition'
  Properties:
    Name: DemoTestSubscriptionDefinition
TestSubscriptionDefinitionVersion:
```

```

Type: 'AWS::Greengrass::SubscriptionDefinitionVersion'
Properties:
  SubscriptionDefinitionId: !Ref TestSubscriptionDefinition
  Subscriptions:
    - Id: TestSubscription1
      Source: !Join
        - ':'
        - - 'arn:aws:iot'
          - !Ref 'AWS::Region'
          - !Ref 'AWS::AccountId'
          - thing/TestClientDevice1
      Subject: TestSubjectUpdated
      Target: !Ref LambdaVersionArn
TestGroup:
  Type: 'AWS::Greengrass::Group'
  Properties:
    Name: DemoTestGroupNewName
    RoleArn: !Join
      - ':'
      - - 'arn:aws:iam:'
        - !Ref 'AWS::AccountId'
        - role/TestUser
    InitialVersion:
      CoreDefinitionVersionArn: !Ref TestCoreDefinitionVersion
      DeviceDefinitionVersionArn: !Ref TestDeviceDefinitionVersion
      FunctionDefinitionVersionArn: !Ref TestFunctionDefinitionVersion
      SubscriptionDefinitionVersionArn: !Ref TestSubscriptionDefinitionVersion
      LoggerDefinitionVersionArn: !Ref TestLoggerDefinitionVersion
      ResourceDefinitionVersionArn: !Ref TestResourceDefinitionVersion
  Tags:
    KeyName0: value
    KeyName1: value
    KeyName2: value
Outputs:
  CommandToDeployGroup:
    Value: !Join
      - ' '
      - - groupVersion=$(cut -d'/' -f6 <<<
        - !GetAtt
          - TestGroup
          - LatestVersionArn
        - );
        - aws --region
        - !Ref 'AWS::Region'

```



```
- greengrass create-deployment --group-id  
- !Ref TestGroup  
- '--deployment-type NewDeployment --group-version-id'  
- $groupVersion
```

Supported AWS Regions

Currently, you can create and manage AWS IoT Greengrass resources only in the following [AWS Regions](#):

- US East (Ohio)
- US East (N. Virginia)
- US West (Oregon)
- Asia Pacific (Mumbai)
- Asia Pacific (Seoul)
- Asia Pacific (Singapore)
- Asia Pacific (Sydney)
- Asia Pacific (Tokyo)
- China (Beijing)
- Europe (Frankfurt)
- Europe (Ireland)
- Europe (London)
- AWS GovCloud (US-West)

Using AWS IoT Device Tester for AWS IoT Greengrass V1

AWS IoT Device Tester (IDT) is a downloadable testing framework that lets you validate IoT devices. Because AWS IoT Greengrass Version 1 has been moved into [maintenance mode](#), IDT for AWS IoT Greengrass V1 no longer generates signed qualification reports. You will no longer be able to qualify new AWS IoT Greengrass V1 devices to list in the [AWS Partner Device Catalog](#) through the [AWS Device Qualification Program](#). However, you can continue to use IDT for AWS IoT Greengrass V1 to test your Greengrass V1 devices. We recommend that you use [IDT for AWS IoT Greengrass V2](#) to qualify and list Greengrass devices in the [AWS Partner Device Catalog](#).

IDT for AWS IoT Greengrass runs on your host computer (Windows, macOS, or Linux) connected to the device to be tested. It runs tests and aggregates results. It also provides a command line interface to manage the testing process.

AWS IoT Greengrass qualification suite

Use IDT for AWS IoT Greengrass to verify that the AWS IoT Greengrass Core software runs on your hardware and can communicate with the AWS Cloud. It also performs end-to-end tests with AWS IoT Core. For example, it verifies that your device can send and receive MQTT messages and process them correctly.



AWS IoT Device Tester for AWS IoT Greengrass organizes tests using the concepts of *test suites* and *test groups*.

- A test suite is the set of test groups used to verify that a device works with particular versions of AWS IoT Greengrass.
- A test group is the set of individual tests related to a particular feature, such as Greengrass group deployments and MQTT messaging.

For more information, see [Use IDT to run the AWS IoT Greengrass qualification suite](#).

Custom test suites

Starting in IDT v4.0.0, IDT for AWS IoT Greengrass combines a standardized configuration setup and result format with a test suite environment that enables you to develop custom test suites for your devices and device software. You can add custom tests for your own internal validation or provide them to your customers for device verification.

How a test writer configures a custom test suite determines the settings configurations that are required to run custom test suites. For more information, see [Use IDT to develop and run your own test suites](#).

Supported versions of AWS IoT Device Tester for AWS IoT Greengrass V1

Because AWS IoT Greengrass Version 1 has been moved into [maintenance mode](#), IDT for AWS IoT Greengrass V1 no longer generates signed qualification reports. We recommend that you use [IDT for AWS IoT Greengrass V2](#).

For information about IDT for AWS IoT Greengrass V2, see [Using AWS IoT Device Tester for AWS IoT Greengrass V2](#) in the *AWS IoT Greengrass V2 Developer Guide*.

Note

You receive a notification when you start a test run if IDT for AWS IoT Greengrass is not compatible with the version of AWS IoT Greengrass you are using.

By downloading the software, you agree to the [AWS IoT Device Tester License Agreement](#).

Unsupported IDT versions for for AWS IoT Greengrass

This topic lists unsupported versions of IDT for AWS IoT Greengrass. Unsupported versions do not receive bug fixes or updates. For more information, see [the section called “Support policy for AWS IoT Device Tester for AWS IoT Greengrass V1”](#).

IDT v4.4.1 for AWS IoT Greengrass versions v1.11.6, v1.10.5

Release notes:

- Enables you to validate and qualify devices running AWS IoT Greengrass core software v1.11.6 and v1.10.5.
- Contains minor bug fixes.

Test suite version:

GGQ_1.3.1

- Released 2021.12.20

IDT v4.1.0 for AWS IoT Greengrass versions v1.11.4, v1.10.4

Release notes:

- Enables you to validate and qualify devices running AWS IoT Greengrass core software v1.11.4 and v1.10.4.
- Fixes an issue that caused the logs that are displayed during a test run to use redundant tags.

Test suite version:

GGQ_1.3.0

- Released 2021.06.23
- Adds retries for API calls to Lambda, IAM, and AWS STS to improve handling for throttling or server issues.
- Adds support for Python 3.8 to the ML and Docker test cases.

IDT v4.0.2 for AWS IoT Greengrass versions v1.11.1, v1.11.0, v1.10.3

Release notes:

- Fixed an issue that caused IDT to mask Hardware Security Integration (HSI) errors.
- Enables you to develop and run your custom test suites using AWS IoT Device Tester for AWS IoT Greengrass. For more information, see [Use IDT to develop and run your own test suites](#).

- Provides code signed IDT applications for macOS and Windows. In macOS, if a security warning message displays, you might need to grant a security exception for IDT. For more information, see [Security exception on macOS](#).

 Note

AWS IoT Greengrass doesn't provide a Dockerfile or a Docker image for version 1.11.1 of the AWS IoT Greengrass core software. To test your device for Docker qualification, use an earlier version of AWS IoT Greengrass core software.

IDT v3.2.0 for AWS IoT Greengrass versions v1.11.0, v1.10.1, v1.10.0

Release notes:

- By default, IDT runs only required tests for qualification. To qualify for additional features, you can modify the [device.json](#) file.
- Added a port number in `device.json` that you can configure for SSH connections.
- Docker supports only [stream manager](#) and machine learning (ML) without containerization. Container, Docker, and Hardware Security Integration (HSI) are not available for Docker devices.
- We merged `device-ml.json` and `device-hsm.json` into `device.json`.

IDT v3.1.3 for AWS IoT Greengrass versions: v1.10.x, v1.9.x, v1.8.x

Release notes:

- Added support for ML feature qualification for AWS IoT Greengrass v1.10.x and v1.9.x. You can now use IDT to validate that your devices can perform ML inference locally with models stored and trained in the cloud.
- Added `--stop-on-first-failure` for the `run-suite` command. You can use this option to configure IDT to stop running on the first failure. We recommend using this option during the debugging stage at the test groups level.
- Added a clock drift check for MQTT tests to ensure that the device under test uses the correct system time. The time used must be within an acceptable time range.

- Added `--update-idt` for the `run-suite` command. You can use this option to set the response for the prompt to update IDT.
- Added `--update-managed-policy` for the `run-suite` command. You can use this option to set the response for the prompt to update the managed policy.
- Added a bug fix for automatic updates of IDT test suite versions. The fix ensures that IDT can run the latest test suites that are available for your AWS IoT Greengrass version.

IDT v3.0.1 for AWS IoT Greengrass

Release notes:

- Added support for AWS IoT Greengrass v1.10.1.
- Automatic updates of IDT test suite versions. IDT can download the latest test suites that are available for your AWS IoT Greengrass version. With this feature:
 - Test suites are versioned using a *major.minor.patch* format. The initial test suite version is `GGQ_1.0.0`.
 - You can download new test suites interactively in the command line interface or set the `upgrade-test-suite` flag when you start IDT.

For more information, see [the section called “Test suite versions”](#).

- Added `list-supported-products`. You can use this command to list the AWS IoT Greengrass and test suite versions that are supported by the installed version of IDT.
- Added `list-test-cases`. You can use this command to list the test cases that are available in a test group.
- Added `test-id` for the `run-suite` command. You can use this option to run individual test cases in a test group.

IDT v2.3.0 for AWS IoT Greengrass v1.10, v1.9.x, and v1.8.x

When testing on a physical device, AWS IoT Greengrass v1.10, v1.9.x, and v1.8.x are supported.

When testing in a Docker container, AWS IoT Greengrass v1.10 and v1.9.x are supported.

Release notes:

- Added support for [the section called “Run AWS IoT Greengrass in a Docker container”](#). You can now use IDT to qualify and validate that your devices can run AWS IoT Greengrass in a Docker container.
- Added an [AWS managed policy](#) (AWSIoTDeviceTesterForGreengrassFullAccess) that defines the permissions required to run AWS IoT Device Tester. If new releases require additional permissions, AWS adds them to this managed policy so you don't have to update your IAM permissions.
- Introduced checks to validate that your environment (for example, device connectivity and internet connectivity) is set up correctly before you run the test cases.
- Improved the Greengrass dependency checker in IDT to make it more flexible while checking for libc on devices.

IDT v2.2.0 for AWS IoT Greengrass v1.10, v1.9.x, and v1.8.x

Release notes:

- Added support for AWS IoT Greengrass v1.10.
- Added support for the [Greengrass Docker application deployment](#) connector.
- Added support for AWS IoT Greengrass [stream manager](#).
- Added support for AWS IoT Greengrass in the China (Beijing) Region.

IDT v2.1.0 for AWS IoT Greengrass v1.9.x, v1.8.x, and v1.7.x

Release notes:

- Added support for AWS IoT Greengrass v1.9.4.
- Added support for Linux-ARMv6l devices.

IDT v2.0.0 for AWS IoT Greengrass v1.9.3, v1.9.2, v1.9.1, v1.9.0, v1.8.4, v1.8.3, and v1.8.2

Release notes:

- Removed dependency on Python for device under test.
- Test suite execution time reduced by more than 50 percent, which makes the qualification process faster.

- Executable size reduced by more than 50 percent, which makes download and installation faster.
- Improved [timeout multiplier support](#) for all test cases.
- Enhanced post-diagnostics messages to troubleshoot errors faster.
- Updated the permissions policy template required to run IDT.
- Added support for AWS IoT Greengrass v1.9.3.

IDT v1.3.3 for AWS IoT Greengrass v1.9.2, v1.9.1, v1.9.0, v1.8.3, and v1.8.2

Release notes:

- Added support for Greengrass v1.9.2 and v1.8.3.
- Added support for Greengrass OpenWrt.
- Added SSH user name and password device sign-in.
- Added native test bug fix for OpenWrt-ARMv7l platform.

IDT v1.2 for AWS IoT Greengrass v1.8.1

Release notes:

- Added a configurable timeout multiplier to address and troubleshoot timeout issues (for example, low bandwidth connections).

IDT v1.1 for AWS IoT Greengrass v1.8.0

Release notes:

- Added support for AWS IoT Greengrass Hardware Security Integration (HSI).
- Added support for AWS IoT Greengrass container and no container.
- Added automated AWS IoT Greengrass service role creation.
- Improved test resource cleanup.
- Added test execution summary report.

IDT v1.1 for AWS IoT Greengrass v1.7.1

Release notes:

- Added support for AWS IoT Greengrass Hardware Security Integration (HSI).
- Added support for AWS IoT Greengrass container and no container.
- Added automated AWS IoT Greengrass service role creation.
- Improved test resource cleanup.
- Added test execution summary report.

IDT v1.0 for AWS IoT Greengrass v1.6.1

Release notes:

- Added OTA test bug fix for future AWS IoT Greengrass version compatibility.

Note

If you're using IDT v1.0 for AWS IoT Greengrass v1.6.1, you must create a [Greengrass service role](#). In later versions, IDT creates the service role for you.

Use IDT to run the AWS IoT Greengrass qualification suite

You can use AWS IoT Device Tester (IDT) for AWS IoT Greengrass to verify that the AWS IoT Greengrass Core software runs on your hardware and can communicate with the AWS Cloud. It also performs end-to-end tests with AWS IoT Core. For example, it verifies that your device can send and receive MQTT messages and process them correctly.

Because AWS IoT Greengrass Version 1 has been moved into [maintenance mode](#), IDT for AWS IoT Greengrass V1 no longer generates signed qualification reports. If you want to add your hardware to the AWS Partner Device Catalog, run the AWS IoT Greengrass V2 qualification suite to generate test reports that you can submit to AWS IoT. For more information, see [AWS Device Qualification Program](#) and [Supported versions of IDT for AWS IoT Greengrass V2](#).

In addition to testing devices, IDT for AWS IoT Greengrass creates resources (for example, AWS IoT things, AWS IoT Greengrass groups, Lambda functions, and so on) in your AWS account to facilitate the qualification process.

To create these resources, IDT for AWS IoT Greengrass uses the AWS credentials configured in the `config.json` file to make API calls on your behalf. These resources are provisioned at various times during a test.

When you use IDT for AWS IoT Greengrass to run the AWS IoT Greengrass qualification suite, IDT performs the following steps:

1. Loads and validates your device and credential configurations.
2. Performs selected tests with the required local and cloud resources.
3. Cleans up local and cloud resources.
4. Generates tests reports that indicate if your device passed the tests required for qualification.

Test suite versions

IDT for AWS IoT Greengrass organizes tests into test suites and test groups.

- A test suite is the set of test groups used to verify that a device works with particular versions of AWS IoT Greengrass.
- A test group is the set of individual tests related to a particular feature, such as Greengrass group deployments and MQTT messaging.

Starting in IDT v3.0.0, test suites are versioned using a *major.minor.patch* format, for example GGQ_1.0.0. When you download IDT, the package includes the latest test suite version.

Important

IDT supports the three latest test suite versions for device qualification. For more information, see [the section called "Support policy for AWS IoT Device Tester for AWS IoT Greengrass V1"](#).

You can run `list-supported-products` to list the versions of AWS IoT Greengrass and test suites that are supported by your current version of IDT. Tests from unsupported test suite versions are not valid for device qualification. IDT doesn't print qualification reports for unsupported versions.

Updates to IDT configuration settings

New tests might introduce new IDT configuration settings.

- If the settings are optional, IDT continues running the tests.

- If the settings are required, IDT notifies you and stops running. After you configure the settings, restart the test run.

Configuration settings are located in the `<device-tester-extract-location>/configs` folder. For more information, see [the section called "Configure IDT settings"](#).

If an updated test suite version adds configuration settings, IDT creates a copy of the original configuration file in `<device-tester-extract-location>/configs`.

Test group descriptions

IDT v2.0.0 and later

Required Test Groups for Core Qualification

These test groups are required to qualify your AWS IoT Greengrass device for the AWS Partner Device Catalog.

AWS IoT Greengrass Core Dependencies

Validates that your device meets all software and hardware requirements for the AWS IoT Greengrass Core software.

The Software Packages Dependencies test case in this test group is not applicable when testing in a [Docker container](#).

Deployment

Validates that Lambda functions can be deployed on your device.

MQTT

Verifies the AWS IoT Greengrass message router functionality by checking local communication between the Greengrass core and client devices, which are local IoT devices.

Over-the-Air (OTA)

Validates that your device can successfully perform an OTA update of the AWS IoT Greengrass Core software.

This test group is not applicable when testing in a [Docker container](#).

Version

Checks that the version of AWS IoT Greengrass provided is compatible with the AWS IoT Device Tester version you are using.

Optional Test Groups

These test groups are optional. If you choose to qualify for optional tests, your device is listed with additional capabilities in the AWS Partner Device Catalog.

Container Dependencies

Validates that the device meets all of the software and hardware requirements to run Lambda functions in container mode on a Greengrass core.

This test group is not applicable when testing in a [Docker container](#).

Deployment Container

Validates that Lambda functions can be deployed on the device and run in container mode on a Greengrass core.

This test group is not applicable when testing in a [Docker container](#).

Docker Dependencies (Supported for IDT v2.2.0 and later)

Validates that the device meets all the required technical dependencies to use the Greengrass Docker application deployment connector to run containers

This test group is not applicable when testing in a [Docker container](#).

Hardware Security Integration (HSI)

Verifies that the provided HSI shared library can interface with the hardware security module (HSM) and implements the required PKCS#11 APIs correctly. The HSM and shared library must be able to sign a CSR, perform TLS operations, and provide the correct key lengths and public key algorithm.

Stream Manager Dependencies (Supported for IDT v2.2.0 and later)

Validates that the device meets all of the required technical dependencies to run AWS IoT Greengrass stream manager.

Machine Learning Dependencies (Supported for IDT v3.1.0 and later)

Validates that the device meets all of the required technical dependencies to perform ML inference locally.

Machine Learning Inference Tests (Supported for IDT v3.1.0 and later)

Validates that ML inference can be performed on the given device under test. For more information, see [the section called “Optional: Configuring your device for ML qualification”](#).

Machine Learning Inference Container Tests (Supported for IDT v3.1.0 and later)

Validates that ML inference can be performed on the given device under test and run in container mode on a Greengrass core. For more information, see [the section called “Optional: Configuring your device for ML qualification”](#).

IDT v1.3.3 and earlier

Required Test Groups for Core Qualification

These tests are required to qualify your AWS IoT Greengrass device for the AWS Partner Device Catalog.

AWS IoT Greengrass Core Dependencies

Validates that your device meets all software and hardware requirements for the AWS IoT Greengrass Core software.

Combination (Device Security Interaction)

Verifies the functionality of the device certificate manager and IP detection on the Greengrass core device by changing connectivity information on the Greengrass group in the cloud. The test group rotates the AWS IoT Greengrass server certificate and verifies that AWS IoT Greengrass allows connections.

Deployment (Required for IDT v1.2 and earlier)

Validates that Lambda functions can be deployed on your device.

Device Certificate Manager (DCM)

Verifies that the AWS IoT Greengrass device certificate manager can generate a server certificate on startup and rotate certificates if they are close to expiration.

IP Detection (IPD)

Verifies that core connectivity information is updated when there are IP address changes in a Greengrass core device. For more information, see [Activate automatic IP detection](#).

Logging

Verifies that the AWS IoT Greengrass logging service can write to a log file using a user Lambda function written in Python.

MQTT

Verifies the AWS IoT Greengrass message router functionality by sending messages on a topic that is routed to two Lambda functions.

Native

Verifies that AWS IoT Greengrass can run native (compiled) Lambda functions.

Over-the-Air (OTA)

Validates that your device can successfully perform a OTA update of the AWS IoT Greengrass Core software.

Penetration

Validates that the AWS IoT Greengrass Core software fails to start if hard link/soft link protection and [seccomp](#) are not enabled. It is also used to verify other security-related features.

Shadow

Verifies local shadow and shadow cloud-syncing functionality.

Spooler

Validates that the MQTT messages are queued with the default spooler configuration.

Token Exchange Service (TES)

Verifies that AWS IoT Greengrass can exchange its core certificate for valid AWS credentials.

Version

Checks that the version of AWS IoT Greengrass provided is compatible with the AWS IoT Device Tester version you are using.

Optional Test Groups

These tests are optional. If you choose to qualify for optional tests, your device is listed with additional capabilities in the AWS Partner Device Catalog.

Container Dependencies

Checks that the device meets all of the required dependencies to run Lambda functions in container mode.

Hardware Security Integration (HSI)

Verifies that the provided HSI shared library can interface with the hardware security module (HSM) and implements the required PKCS#11 APIs correctly. The HSM and shared library must be able to sign a CSR, perform TLS operations, and provide the correct key lengths and public key algorithm.

Local Resource Access

Verifies the local resource access (LRA) feature of AWS IoT Greengrass by providing access to local files and directories owned by various Linux users and groups to containerized Lambda functions through AWS IoT Greengrass LRA APIs. Lambda functions should be allowed or denied access to local resources based on local resource access configuration.

Network

Verifies that socket connections can be established from a Lambda function. These socket connections should be allowed or denied based on the Greengrass core configuration.

Prerequisites for running the AWS IoT Greengrass qualification suite

This section describes the prerequisites for using AWS IoT Device Tester (IDT) for AWS IoT Greengrass to run the AWS IoT Greengrass qualification suite.

Download the latest version of AWS IoT Device Tester for AWS IoT Greengrass

Download the [latest version](#) of IDT and extract the software into a location on your file system where you have read and write permissions.

Note

IDT does not support being run by multiple users from a shared location, such as an NFS directory or a Windows network shared folder. We recommend that you extract the IDT package to a local drive and run the IDT binary on your local workstation.

Windows has a path length limitation of 260 characters. If you are using Windows, extract IDT to a root directory like C:\ or D:\ to keep your paths under the 260 character limit.

Create and configure an AWS account

Before you can use IDT for AWS IoT Greengrass, you must perform the following steps:

1. [Create an AWS account](#). If you already have an AWS account, skip to step 2.
2. [Configure permissions for IDT](#).

These account permissions allow IDT to access AWS services and create AWS resources, such as AWS IoT things, Greengrass groups, and Lambda functions, on your behalf.

To create these resources, IDT for AWS IoT Greengrass uses the AWS credentials configured in the `config.json` file to make API calls on your behalf. These resources are provisioned at various times during a test.

Note

Although most tests qualify for [Amazon Web Services Free Tier](#), you must provide a credit card when you sign up for an AWS account. For more information, see [Why do I need a payment method if my account is covered by the Free Tier?](#).

Step 1: Create an AWS account

In this step, create and configure an AWS account. If you already have an AWS account, skip to [the section called "Step 2: Configure permissions for IDT"](#).

Sign up for an AWS account

If you do not have an AWS account, complete the following steps to create one.

To sign up for an AWS account

1. Open <https://portal.aws.amazon.com/billing/signup>.
2. Follow the online instructions.

Part of the sign-up procedure involves receiving a phone call or text message and entering a verification code on the phone keypad.

When you sign up for an AWS account, an *AWS account root user* is created. The root user has access to all AWS services and resources in the account. As a security best practice, assign administrative access to a user, and use only the root user to perform [tasks that require root user access](#).

AWS sends you a confirmation email after the sign-up process is complete. At any time, you can view your current account activity and manage your account by going to <https://aws.amazon.com/> and choosing **My Account**.

Create a user with administrative access

After you sign up for an AWS account, secure your AWS account root user, enable AWS IAM Identity Center, and create an administrative user so that you don't use the root user for everyday tasks.

Secure your AWS account root user

1. Sign in to the [AWS Management Console](#) as the account owner by choosing **Root user** and entering your AWS account email address. On the next page, enter your password.

For help signing in by using root user, see [Signing in as the root user](#) in the *AWS Sign-In User Guide*.

2. Turn on multi-factor authentication (MFA) for your root user.

For instructions, see [Enable a virtual MFA device for your AWS account root user \(console\)](#) in the *IAM User Guide*.

Create a user with administrative access

1. Enable IAM Identity Center.

For instructions, see [Enabling AWS IAM Identity Center](#) in the *AWS IAM Identity Center User Guide*.

2. In IAM Identity Center, grant administrative access to a user.

For a tutorial about using the IAM Identity Center directory as your identity source, see [Configure user access with the default IAM Identity Center directory](#) in the *AWS IAM Identity Center User Guide*.

Sign in as the user with administrative access

- To sign in with your IAM Identity Center user, use the sign-in URL that was sent to your email address when you created the IAM Identity Center user.

For help signing in using an IAM Identity Center user, see [Signing in to the AWS access portal](#) in the *AWS Sign-In User Guide*.

Assign access to additional users

1. In IAM Identity Center, create a permission set that follows the best practice of applying least-privilege permissions.

For instructions, see [Create a permission set](#) in the *AWS IAM Identity Center User Guide*.

2. Assign users to a group, and then assign single sign-on access to the group.

For instructions, see [Add groups](#) in the *AWS IAM Identity Center User Guide*.

Step 2: Configure permissions for IDT

In this step, configure the permissions that IDT for AWS IoT Greengrass uses to run tests and collect IDT usage data. You can use the AWS Management Console or AWS Command Line Interface (AWS CLI) to create an IAM policy and a test user for IDT, and then attach policies to the user. If you already created a test user for IDT, skip to [the section called "Configure your device to run IDT tests"](#) or [the section called "Optional: Configuring your Docker container"](#).

- [To Configure Permissions for IDT \(Console\)](#)
- [To Configure Permissions for IDT \(AWS CLI\)](#)

To configure permissions for IDT (console)

Follow these steps to use the console to configure permissions for IDT for AWS IoT Greengrass.

1. Sign in to the [IAM console](#).
2. Create a customer managed policy that grants permissions to create roles with specific permissions.
 - a. In the navigation pane, choose **Policies**, and then choose **Create policy**.
 - b. On the **JSON** tab, replace the placeholder content with the following policy.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "ManageRolePoliciesForIDTGreengrass",
      "Effect": "Allow",
      "Action": [
        "iam:DetachRolePolicy",
        "iam:AttachRolePolicy"
      ],
      "Resource": [
        "arn:aws:iam::*:role/idt-*",
        "arn:aws:iam::*:role/GreengrassServiceRole"
      ],
      "Condition": {
        "ArnEquals": {
          "iam:PolicyARN": [
            "arn:aws:iam::aws:policy/service-role/AWSGreengrassResourceAccessRolePolicy",
            "arn:aws:iam::aws:policy/service-role/GreengrassOTAUpdateArtifactAccess",
            "arn:aws:iam::aws:policy/service-role/AWSLambdaBasicExecutionRole"
          ]
        }
      }
    },
    {
      "Sid": "ManageRolesForIDTGreengrass",
      "Effect": "Allow",
      "Action": [
        "iam:CreateRole",
        "iam>DeleteRole",
        "iam:PassRole",
        "iam:GetRole"
      ]
    }
  ]
}
```

```

    ],
    "Resource": [
        "arn:aws:iam::*:role/idt-*",
        "arn:aws:iam::*:role/GreengrassServiceRole"
    ]
  }
]
}

```

Important

The following policy grants permission to create and manage roles required by IDT for AWS IoT Greengrass. This includes permissions to attach the following AWS managed policies:

- [AWSGreengrassResourceAccessRolePolicy](#)
- [GreengrassOTAUpdateArtifactAccess](#)
- [AWSLambdaBasicExecutionRole](#)

- c. Choose **Next: Tags**.
 - d. Choose **Next: Review**.
 - e. For **Name**, enter **IDTGreengrassIAMPermissions**. Under **Summary**, review the permissions granted by your policy.
 - f. Choose **Create policy**.
3. Create an IAM user and attach the permissions required by IDT for AWS IoT Greengrass.
 - a. Create an IAM user. Follow steps 1 through 5 in [Creating IAM users \(console\)](#) in the *IAM User Guide*.
 - b. Attach the permissions to your IAM user:
 - i. On the **Set permissions** page, choose **Attach existing policies directly**.
 - ii. Search for the **IDTGreengrassIAMPermissions** policy that you created in the previous step. Select the check box.
 - iii. Search for the **AWSIoTDeviceTesterForGreengrassFullAccess** policy. Select the check box.

Note

The [AWSIoTDeviceTesterForGreengrassFullAccess](#) is an AWS managed policy that defines the permissions IDT requires to create and access AWS resources used for testing. For more information, see [the section called “AWS managed policy for IDT”](#).

- c. Choose **Next: Tags**.
 - d. Choose **Next: Review** to view a summary of your choices.
 - e. Choose **Create user**.
 - f. To view the user's access keys (access key IDs and secret access keys), choose **Show** next to the password and access key. To save the access keys, choose **Download.csv** and save the file to a secure location. You use this information later to configure your AWS credentials file.
4. Next step: Configure your [physical device](#).

To configure permissions for IDT (AWS CLI)

Follow these steps to use the AWS CLI to configure permissions for IDT for AWS IoT Greengrass. If you already configured permissions in the console, skip to [the section called “Configure your device to run IDT tests”](#) or [the section called “Optional: Configuring your Docker container”](#).

1. On your computer, install and configure the AWS CLI if it's not already installed. Follow the steps in [Installing the AWS CLI](#) in the *AWS Command Line Interface User Guide*.

Note

The AWS CLI is an open source tool that you can use to interact with AWS services from your command-line shell.

2. Create a customer managed policy that grants permissions to manage IDT and AWS IoT Greengrass roles.

Linux, macOS, or Unix

```
aws iam create-policy --policy-name IDTGreengrassIAMPermissions --policy-
document '{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "ManageRolePoliciesForIDTGreengrass",
      "Effect": "Allow",
      "Action": [
        "iam:DetachRolePolicy",
        "iam:AttachRolePolicy"
      ],
      "Resource": [
        "arn:aws:iam::*:role/idt-*",
        "arn:aws:iam::*:role/GreengrassServiceRole"
      ],
      "Condition": {
        "ArnEquals": {
          "iam:PolicyARN": [
            "arn:aws:iam::aws:policy/service-role/
AWSGreengrassResourceAccessRolePolicy",
            "arn:aws:iam::aws:policy/service-role/
GreengrassOTAUpdateArtifactAccess",
            "arn:aws:iam::aws:policy/service-role/
AWSLambdaBasicExecutionRole"
          ]
        }
      }
    },
    {
      "Sid": "ManageRolesForIDTGreengrass",
      "Effect": "Allow",
      "Action": [
        "iam:CreateRole",
        "iam>DeleteRole",
        "iam:PassRole",
        "iam:GetRole"
      ],
      "Resource": [
        "arn:aws:iam::*:role/idt-*",
        "arn:aws:iam::*:role/GreengrassServiceRole"
      ]
    }
  ]
}
```

```
    }
  ]
}'
```

Windows command prompt

```
aws iam create-policy --policy-name IDTGreengrassIAMPermissions --
policy-document '{"Version": "2012-10-17", "Statement": [{"Sid":
"\": "ManageRolePoliciesForIDTGreengrass", "Effect": "Allow",
"Action": ["iam:DetachRolePolicy", "iam:AttachRolePolicy"],
"Resource": ["arn:aws:iam::*:role/idt-*", "arn:aws:iam::*:role/
GreengrassServiceRole"], "Condition": {"ArnEquals": {"iam:PolicyARN":
["arn:aws:iam::aws:policy/service-role/AWSGreengrassResourceAccessRolePolicy
", "arn:aws:iam::aws:policy/service-role/GreengrassOTAUpdateArtifactAccess
", "arn:aws:iam::aws:policy/service-role/AWSLambdaBasicExecutionRole"]}}},
{"Sid": "ManageRolesForIDTGreengrass", "Effect": "Allow", "Action":
["iam:CreateRole", "iam:DeleteRole", "iam:PassRole", "iam:GetRole
"], "Resource": ["arn:aws:iam::*:role/idt-*", "arn:aws:iam::*:role/
GreengrassServiceRole"]}]}'
```

Note

This step includes a Windows command prompt example because it uses a different JSON syntax than Linux, macOS, or Unix terminal commands.

3. Create an IAM user and attach the permissions required by IDT for AWS IoT Greengrass.
 - a. Create an IAM user. In this example setup, the user is named IDTGreengrassUser.

```
aws iam create-user --user-name IDTGreengrassUser
```

- b. Attach the IDTGreengrassIAMPermissions policy you created in step 2 to your IAM user. Replace *<account-id>* in the command with the ID of your AWS account.

```
aws iam attach-user-policy --user-name IDTGreengrassUser --policy-arn
arn:aws:iam::<account-id>:policy/IDTGreengrassIAMPermissions
```

- c. Attach the AWSIoTDeviceTesterForGreengrassFullAccess policy to your IAM user.

```
aws iam attach-user-policy --user-name IDTGreengrassUser --policy-arn
arn:aws:iam::aws:policy/AWSIoTDeviceTesterForGreengrassFullAccess
```

 Note

The [AWSIoTDeviceTesterForGreengrassFullAccess](#) is an AWS managed policy that defines the permissions IDT requires to create and access AWS resources used for testing. For more information, see [the section called “AWS managed policy for IDT”](#).

4. Create a secret access key for the user.

```
aws iam create-access-key --user-name IDTGreengrassUser
```

Store the output in a secure location. You use this information later to configure your AWS credentials file.

5. Next step: Configure your [physical device](#).

AWS managed policy for AWS IoT Device Tester

The [AWSIoTDeviceTesterForGreengrassFullAccess](#) managed policy allows IDT to run operations and collect usage metrics. This policy grants the following IDT permissions:

- `iot-device-tester:CheckVersion`. Check whether a set of AWS IoT Greengrass, test suite, and IDT versions are compatible.
- `iot-device-tester:DownloadTestSuite`. Download test suites.
- `iot-device-tester:LatestIdt`. Get information about the latest IDT version that is available for download.
- `iot-device-tester:SendMetrics`. Publish usage data that IDT collects about your tests.
- `iot-device-tester:SupportedVersion`. Get the list of AWS IoT Greengrass and test suite versions that are supported by IDT. This information is displayed in the command-line window.

Configure your device to run IDT tests

To configure your device you must install AWS IoT Greengrass dependencies, configure the AWS IoT Greengrass Core software, configure your host computer to access your device, and configure user permissions on your device.

Verify AWS IoT Greengrass dependencies on the device under test

Before IDT for AWS IoT Greengrass can test your devices, make sure that you have set up your device as described in [Getting started with AWS IoT Greengrass](#). For information about supported platforms, see [Supported platforms](#).

Configure the AWS IoT Greengrass software

IDT for AWS IoT Greengrass tests your device for compatibility with a specific version of AWS IoT Greengrass. IDT provides two options for testing AWS IoT Greengrass on your devices:

- Download and use a version of the [AWS IoT Greengrass Core software](#). IDT installs the software for you.
- Use a version of the AWS IoT Greengrass Core software already installed on your device.

Note

Each version of AWS IoT Greengrass has a corresponding IDT version. You must download the version of IDT that corresponds to the version of AWS IoT Greengrass you are using.

The following sections describe these options. You only need to do one.

Option 1: Download the AWS IoT Greengrass Core software and configure AWS IoT Device Tester to use it

You can download the AWS IoT Greengrass Core software from the [AWS IoT Greengrass Core Software](#) downloads page.

1. Find the correct architecture and Linux distribution, and then choose **Download**.
2. Copy the tar.gz file to the `<device-tester-extract-location>/products/greengrass/ggc`.

Note

Do not change the name of the AWS IoT Greengrass tar.gz file. Do not place multiple files in this directory for the same operating system and architecture. For example having both `greengrass-linux-armv7l-1.7.1.tar.gz` and `greengrass-linux-armv7l-1.8.1.tar.gz` files in that directory will cause the tests to fail.

Option 2: Use an existing installation of AWS IoT Greengrass with AWS IoT Device Tester

Configure IDT to test the AWS IoT Greengrass Core software installed on your device by adding the `greengrassLocation` attribute to the `device.json` file in the `<device-tester-extract-location>/configs` folder. For example:

```
"greengrassLocation" : "<path-to-greengrass-on-device>"
```

For more information about the `device.json` file, see [Configure device.json](#).

On Linux devices, the default location of the AWS IoT Greengrass Core software is `/greengrass`.

Note

Your device should have an installation of the AWS IoT Greengrass Core software that has not been started.

Make sure you have added the `ggc_user` user and `ggc_group` on your device. For more information, see [Environment setup for AWS IoT Greengrass](#).

Configure your host computer to access your device under test

IDT runs on your host computer and must be able to use SSH to connect to your device. There are two options to allow IDT to gain SSH access to your devices under test:

1. Follow the instructions here to create an SSH key pair and authorize your key to sign in to your device under test without specifying a password.
2. Provide a user name and password for each device in the `device.json` file. For more information, see [Configure device.json](#).

You can use any SSL implementation to create an SSH key. The following instructions show you how to use [SSH-KEYGEN](#) or [PuTTYgen](#) (for Windows). If you are using another SSL implementation, refer to the documentation for that implementation.

IDT uses SSH keys to authenticate with your device under test.

To create an SSH key with SSH-KEYGEN

1. Create an SSH key.

You can use the Open SSH **ssh-keygen** command to create an SSH key pair. If you already have an SSH key pair on your host computer, it is a best practice to create a SSH key pair specifically for IDT. This way, after you have completed testing, your host computer can no longer connect to your device without entering a password. It also allows you to restrict access to the remote device to only those who need it.

Note

Windows does not have an installed SSH client. For information about installing an SSH client on Windows, see [Download SSH Client Software](#).

The **ssh-keygen** command prompts you for a name and path to store the key pair. By default, the key pair files are named `id_rsa` (private key) and `id_rsa.pub` (public key). On macOS and Linux, the default location of these files is `~/.ssh/`. On Windows, the default location is `C:\Users\<user-name>\.ssh`.

When prompted, enter a key phrase to protect your SSH key. For more information, see [Generate a New SSH key](#).

2. Add authorized SSH keys to your device under test.

IDT must use your SSH private key to sign in to your device under test. To authorize your SSH private key to sign in to your device under test, use the **ssh-copy-id** command from your host computer. This command adds your public key into the `~/.ssh/authorized_keys` file on your device under test. For example:

```
$ ssh-copy-id <remote-ssh-user>@<remote-device-ip>
```

Where *remote-ssh-user* is the user name used to sign in to your device under test and *remote-device-ip* is the IP address of the device under test to run tests against. For example:

ssh-copy-id pi@192.168.1.5

When prompted, enter the password for the user name you specified in the **ssh-copy-id** command.

ssh-copy-id assumes the public key is named `id_rsa.pub` and is stored the default location (on macOS and Linux, `~/.ssh/` and on Windows, `C:\Users\<user-name>\.ssh`). If you gave the public key a different name or stored it in a different location, you must specify the fully qualified path to your SSH public key using the **-i** option to **ssh-copy-id** (for example, **ssh-copy-id -i ~/my/path/myKey.pub**). For more information about creating SSH keys and copying public keys, see [SSH-COPY-ID](#).

To create an SSH key using PuTTYgen (Windows only)

1. Make sure you have the OpenSSH server and client installed on your device under test. For more information, see [OpenSSH](#).
2. Install [PuTTYgen](#) on your device under test.
3. Open PuTTYgen.
4. Choose **Generate** and move your mouse cursor inside the box to generate a private key.
5. From the **Conversions** menu, choose **Export OpenSSH key**, and save the private key with a `.pem` file extension.
6. Add the public key to the `/home/<user>/.ssh/authorized_keys` file on device under test.
 - a. Copy the public key text from the PuTTYgen window.
 - b. Use PuTTY to create a session on your device under test.
 - i. From a command prompt or Windows Powershell window, run the following command:
C:/<path-to-putty>/putty.exe -ssh <user>@<dut-ip-address>
 - ii. When prompted, enter your device's password.

- iii. Use vi or another text editor to append the public key to the `/home/<user>/.ssh/authorized_keys` file on your device under test.
7. Update your `device.json` file with your user name, the IP address, and path to the private key file that you just saved on your host computer for each device under test. For more information, see [the section called “Configure device.json”](#). Make sure you provide the full path and file name to the private key and use forward slashes ('/'). For example, for the Windows path `C:\DT\privatekey.pem`, use `C:/DT/privatekey.pem` in the `device.json` file.

Configure user permissions on your device

IDT performs operations on various directories and files in a device under test. Some of these operations require elevated permissions (using **sudo**). To automate these operations, IDT for AWS IoT Greengrass must be able to run commands with sudo without being prompted for a password.

Follow these steps on the device under test to allow sudo access without being prompted for a password.

Note

username refers to the SSH user used by IDT to access the device under test.

To add the user to the sudo group

1. On the device under test, run `sudo usermod -aG sudo <username>`.
2. Sign out and then sign back in for changes to take effect.
3. To verify your user name was added successfully, run **sudo echo test**. If you are not prompted for a password, your user is configured correctly.
4. Open the `/etc/sudoers` file and add the following line to the end of the file:

```
<ssh-username> ALL=(ALL) NOPASSWD: ALL
```

Configure your device to test optional features

The following topics describe how to configure your devices to run IDT tests for optional features. Follow these configuration steps only if you want to test these features. Otherwise, continue to [the section called “Configure IDT settings”](#).

Topics

- [Optional: Configuring your Docker container for IDT for AWS IoT Greengrass](#)
- [Optional: Configuring your device for ML qualification](#)

Optional: Configuring your Docker container for IDT for AWS IoT Greengrass

AWS IoT Greengrass provides a Docker image and Dockerfile that make it easier to run the AWS IoT Greengrass Core software in a Docker container. After you set up the AWS IoT Greengrass container, you can run IDT tests. Currently, only x86_64 Docker architectures are supported to run IDT for AWS IoT Greengrass.

This feature requires IDT v2.3.0 or later.

The process of setting up the Docker container to run IDT tests depends on whether you use the Docker image or Dockerfile provided by AWS IoT Greengrass.

- [Use the Docker image](#). The Docker image has the AWS IoT Greengrass Core software and dependencies installed.
- [Use the Dockerfile](#). The Dockerfile contains source code you can use to build custom AWS IoT Greengrass container images. The image can be modified to run on different platform architectures or to reduce the image size.

Note

AWS IoT Greengrass doesn't provide Dockerfiles or Docker images for AWS IoT Greengrass core software version 1.11.1. To run IDT tests on your own custom container images, your image must include the dependencies defined in the Dockerfile provided by AWS IoT Greengrass.

The following features aren't available when you run AWS IoT Greengrass in a Docker container:

- [Connectors](#) that run in **Greengrass container** mode. To run a connector in a Docker container, the connector must run in **No container** mode. To find connectors that support **No container** mode, see [the section called "AWS-provided Greengrass connectors"](#). Some of these connectors have an isolation mode parameter that you must set to **No container**.
- [Local device and volume resources](#). Your user-defined Lambda functions that run in the Docker container must access devices and volumes on the core directly.

Configure the Docker image provided by AWS IoT Greengrass

Follow these steps to configure the AWS IoT Greengrass Docker image to run IDT tests.

Prerequisites

Before you start this tutorial, you must do the following.

- You must install the following software and versions on your host computer based on the AWS Command Line Interface (AWS CLI) version that you choose.

AWS CLI version 2

- [Docker](#) version 18.09 or later. Earlier versions might also work, but we recommend 18.09 or later.
- AWS CLI version 2.0.0 or later.
 - To install the AWS CLI version 2, see [Installing the AWS CLI version 2](#).
 - To configure the AWS CLI, see [Configuring the AWS CLI](#).

Note

To upgrade to a later AWS CLI version 2 on a Windows computer, you must repeat the [MSI installation](#) process.

AWS CLI version 1

- [Docker](#) version 18.09 or later. Earlier versions might also work, but we recommend 18.09 or later.
- [Python](#) version 3.6 or later.
- [pip](#) version 18.1 or later.
- AWS CLI version 1.17.10 or later
 - To install the AWS CLI version 1, see [Installing the AWS CLI version 1](#).
 - To configure the AWS CLI, see [Configuring the AWS CLI](#).
 - To upgrade to the latest version of the AWS CLI version 1, run the following command.

```
pip install awscli --upgrade --user
```

Note

If you use the [MSI installation](#) of the AWS CLI version 1 on Windows, be aware of the following:

- If the AWS CLI version 1 installation fails to install botocore, try using the [Python and pip installation](#).
- To upgrade to a later AWS CLI version 1, you must repeat the MSI installation process.

- To access Amazon Elastic Container Registry (Amazon ECR) resources, you must grant the following permission.
 - Amazon ECR requires users to grant the `ecr:GetAuthorizationToken` permission through an AWS Identity and Access Management (IAM) policy before they can authenticate to a registry and push or pull images from an Amazon ECR repository. For more information, see [Amazon ECR Repository Policy Examples](#) and [Accessing One Amazon ECR Repository](#) in the *Amazon Elastic Container Registry User Guide*.
1. Download the Docker image and configure the container. You can download the prebuilt image from [Docker Hub](#) or [Amazon Elastic Container Registry](#) (Amazon ECR) and run it on Windows, macOS, and Linux (x86_64) platforms.

To download the Docker image from Amazon ECR, complete all of the steps in [the section called “Get the AWS IoT Greengrass container image from Amazon ECR”](#). Then, return to this topic to continue the configuration.

2. Linux users only: Make sure the user that runs IDT has permission to run Docker commands. For more information, see [Manage Docker as a non-root user](#) in the Docker documentation.
3. To run the AWS IoT Greengrass container, use the command for your operating system:

Linux

```
docker run --rm --init -it -d --name aws-iot-greengrass \
-p 8883:8883 \
-v <host-path-to-kernel-config-file>:<container-path> \
<image-repository>:<tag>
```


- Replace *<host-path-to-kernel-config-file>* with the path to the kernel configuration file on the host and *<container-path>* with the path where the volume is mounted in the container.

The kernel config file on the host is usually located in `/proc/config.gz` or `/boot/config-<kernel-release-date>`. You can run `uname -r` to find the *<kernel-release-date>* value.

Example: To mount the config file from `/boot/config-<kernel-release-date>`

```
-v /boot/config-4.15.0-74-generic:/boot/config-4.15.0-74-generic \
```

Example: To mount the config file from `proc/config.gz`

```
-v /proc/config.gz:/proc/config.gz \
```

- Replace *<image-repository>:<tag>* in the command with the name of the repository and tag of the target image.

Example: To point to the latest version of the AWS IoT Greengrass Core software

```
216483018798.dkr.ecr.us-west-2.amazonaws.com/aws-iot-greengrass:latest
```

To get the list of AWS IoT Greengrass Docker images, run the following command.

```
aws ecr list-images --region us-west-2 --registry-id 216483018798 --  
repository-name aws-iot-greengrass
```

macOS

```
docker run --rm --init -it -d --name aws-iot-greengrass \  
-p 8883:8883 \  
<image-repository>:<tag>
```

- Replace *<image-repository>:<tag>* in the command with the name of the repository and tag of the target image.

Example: To point to the latest version of the AWS IoT Greengrass Core software

```
216483018798.dkr.ecr.us-west-2.amazonaws.com/aws-iot-greengrass:latest
```

To get the list of AWS IoT Greengrass Docker images, run the following command:

```
aws ecr list-images --region us-west-2 --registry-id 216483018798 --  
repository-name aws-iot-greengrass
```

Windows

```
docker run --rm --init -it -d --name aws-iot-greengrass \  
-p 8883:8883 \  
<image-repository>:<tag>
```

- Replace *<image-repository>:<tag>* in the command with the name of the repository and tag of the target image.

Example: To point to the latest version of the AWS IoT Greengrass Core software

```
216483018798.dkr.ecr.us-west-2.amazonaws.com/aws-iot-greengrass:latest
```

To get the list of AWS IoT Greengrass Docker images, run the following command:

```
aws ecr list-images --region us-west-2 --registry-id 216483018798 --  
repository-name aws-iot-greengrass
```

Important

When testing with IDT, do not include the `--entrypoint /greengrass-entrypoint.sh \` argument that's used to run the image for general AWS IoT Greengrass use.

4. Next step: [Configure your AWS credentials and device.json file.](#)

Configure the dockerfile provided by AWS IoT Greengrass

Follow these steps to configure the Docker image built from the AWS IoT Greengrass Dockerfile to run IDT tests.

1. From [the section called “AWS IoT Greengrass Docker software”](#), download the Dockerfile package to your host computer and extract it.
2. Open README.md. The next three steps refer to sections in this file.
3. Make sure that you meet the requirements in the **Prerequisites** section.
4. Linux users only: Complete the **Enable Symlink and Hardlink Protection** and **Enable IPv4 Network Forwarding** steps.
5. To build the Docker image, complete all of the steps in **Step 1. Build the AWS IoT Greengrass Docker Image**. Then, return to this topic to continue the configuration.
6. To run the AWS IoT Greengrass container, use the command for your operating system:

Linux

```
docker run --rm --init -it -d --name aws-iot-greengrass \
-p 8883:8883 \
-v <host-path-to-kernel-config-file>:<container-path> \
<image-repository>:<tag>
```

- Replace *<host-path-to-kernel-config-file>* with the path to the kernel configuration file on the host and *<container-path>* with the path where the volume is mounted in the container.

The kernel config file on the host is usually located in /proc/config.gz or /boot/config-*<kernel-release-date>*. You can run `uname -r` to find the *<kernel-release-date>* value.

Example: To mount the config file from /boot/config-*<kernel-release-date>*

```
-v /boot/config-4.15.0-74-generic:/boot/config-4.15.0-74-generic \
```

Example: To mount the config file from proc/config.gz

```
-v /proc/config.gz:/proc/config.gz \
```

- Replace `<image-repository>:<tag>` in the command with the name of the repository and tag of the target image.

Example: To point to the latest version of the AWS IoT Greengrass Core software

```
216483018798.dkr.ecr.us-west-2.amazonaws.com/aws-iot-greengrass:latest
```

To get the list of AWS IoT Greengrass Docker images, run the following command.

```
aws ecr list-images --region us-west-2 --registry-id 216483018798 --  
repository-name aws-iot-greengrass
```

macOS

```
docker run --rm --init -it -d --name aws-iot-greengrass \  
-p 8883:8883 \  
<image-repository>:<tag>
```

- Replace `<image-repository>:<tag>` in the command with the name of the repository and tag of the target image.

Example: To point to the latest version of the AWS IoT Greengrass Core software

```
216483018798.dkr.ecr.us-west-2.amazonaws.com/aws-iot-greengrass:latest
```

To get the list of AWS IoT Greengrass Docker images, run the following command:

```
aws ecr list-images --region us-west-2 --registry-id 216483018798 --  
repository-name aws-iot-greengrass
```

Windows

```
docker run --rm --init -it -d --name aws-iot-greengrass \  
-p 8883:8883 \  
<image-repository>:<tag>
```

- Replace `<image-repository>:<tag>` in the command with the name of the repository and tag of the target image.

Example: To point to the latest version of the AWS IoT Greengrass Core software

```
216483018798.dkr.ecr.us-west-2.amazonaws.com/aws-iot-greengrass:latest
```

To get the list of AWS IoT Greengrass Docker images, run the following command:

```
aws ecr list-images --region us-west-2 --registry-id 216483018798 --  
repository-name aws-iot-greengrass
```

Important

When testing with IDT, do not include the `--entrypoint /greengrass-entrypoint.sh \` argument that's used to run the image for general AWS IoT Greengrass use.

7. Next step: [Configure your AWS credentials and device.json file.](#)

Troubleshooting your Docker container setup for IDT for AWS IoT Greengrass

Use the following information to help troubleshoot issues with running a Docker container for IDT for AWS IoT Greengrass testing.

WARNING: Error loading config file:/home/user/.docker/config.json - stat /home/<user>/.docker/config.json: permission denied

If you get this error when running `docker` commands on Linux, run the following command. Replace `<user>` in the following command with the user that runs IDT.

```
sudo chown <user>:<user> /home/<user>/.docker -R  
sudo chmod g+rx /home/<user>/.docker -R
```

Optional: Configuring your device for ML qualification

IDT for AWS IoT Greengrass provides machine learning (ML) qualification tests to validate that your devices can perform ML inference locally using cloud-trained models.

To run ML qualification tests, you must first configure your devices as described in [the section called “Configure your device to run IDT tests”](#). Then, follow the steps in this topic to install dependencies for the ML frameworks that you want to run.

IDT v3.1.0 or later is required to run tests for ML qualification.

Installing ML framework dependencies

All ML framework dependencies must be installed under the `/usr/local/lib/python3.x/site-packages` directory. To make sure they are installed under the correct directory, we recommend that you use sudo root permissions when installing the dependencies. Virtual environments are not supported for qualification tests.

Note

If you're testing Lambda functions that run with [containerization](#) (in **Greengrass container mode**), creating symlinks for Python libraries under `/usr/local/lib/python3.x` isn't supported. To avoid errors, you must install the dependencies under the correct directory.

Follow the steps to install the dependencies for your target framework:

- [Install MXNet dependencies](#)
- [the section called “Install TensorFlow dependencies”](#)
- [Install DLR dependencies](#)

Install Apache MXNet dependencies

IDT qualification tests for this framework have the following dependencies:

- Python 3.6 or Python 3.7.

Note

If you're using Python 3.6, you must create a symbolic link from Python 3.7 to Python 3.6 binaries. This configures your device to meet the Python requirement for AWS IoT Greengrass. For example:

```
sudo ln -s path-to-python-3.6/python3.6 path-to-python-3.7/python3.7
```

- Apache MXNet v1.2.1 or later.
- NumPy. The version must be compatible with your MXNet version.

Installing MXNet

Follow the instructions in the MXNet documentation to [install MXNet](#).

Note

If Python 2.x and Python 3.x are both installed on your device, use Python 3.x in the commands that you run to install the dependencies.

Validating the MXNet installation

Choose one of the following options to validate the MXNet installation.

Option 1: SSH into your device and run scripts

1. SSH into your device.
2. Run the following scripts to verify that the dependencies are correctly installed.

```
sudo python3.7 -c "import mxnet; print(mxnet.__version__)"
```

```
sudo python3.7 -c "import numpy; print(numpy.__version__)"
```

The output prints the version number and the script should exit without error.

Option 2: Run the IDT dependency test

1. Make sure that `device.json` is configured for ML qualification. For more information, see [the section called "Configure device.json for ML qualification"](#).
2. Run the dependencies test for the framework.

```
devicetester_[linux | mac | win_x86-64] run-suite --group-id mldependencies --test-id mxnet_dependency_check
```

The test summary displays a PASSED result for mldependencies.

Install TensorFlow dependencies

IDT qualification tests for this framework have the following dependencies:

- Python 3.6 or Python 3.7.

Note

If you're using Python 3.6, you must create a symbolic link from Python 3.7 to Python 3.6 binaries. This configures your device to meet the Python requirement for AWS IoT Greengrass. For example:

```
sudo ln -s path-to-python-3.6/python3.6 path-to-python-3.7/python3.7
```

- TensorFlow 1.x.

Installing TensorFlow

Follow the instructions in the TensorFlow documentation to install TensorFlow 1.x [with pip](#) or [from source](#).

Note

If Python 2.x and Python 3.x are both installed on your device, use Python 3.x in the commands that you run to install the dependencies.

Validating the TensorFlow installation

Choose one of the following options to validate the TensorFlow installation.

Option 1: SSH into your device and run a script

1. SSH into your device.
2. Run the following script to verify that the dependency is correctly installed.

```
sudo python3.7 -c "import tensorflow; print(tensorflow.__version__)"
```

The output prints the version number and the script should exit without error.

Option 2: Run the IDT dependency test

1. Make sure that `device.json` is configured for ML qualification. For more information, see [the section called "Configure device.json for ML qualification"](#).
2. Run the dependencies test for the framework.

```
devicetester_[linux | mac | win_x86-64] run-suite --group-id mldependencies --test-id tensorflow_dependency_check
```

The test summary displays a PASSED result for mldependencies.

Install Amazon SageMaker AI Neo Deep Learning Runtime (DLR) dependencies

IDT qualification tests for this framework have the following dependencies:

- Python 3.6 or Python 3.7.

Note

If you're using Python 3.6, you must create a symbolic link from Python 3.7 to Python 3.6 binaries. This configures your device to meet the Python requirement for AWS IoT Greengrass. For example:

```
sudo ln -s path-to-python-3.6/python3.6 path-to-python-3.7/python3.7
```

- SageMaker AI Neo DLR.

- `numpy`.

After you install the DLR test dependencies, you must [compile the model](#).

Installing DLR

Follow the instructions in the DLR documentation to [install the Neo DLR](#).

Note

If Python 2.x and Python 3.x are both installed on your device, use Python 3.x in the commands that you run to install the dependencies.

Validating the DLR installation

Choose one of the following options to validate the DLR installation.

Option 1: SSH into your device and run scripts

1. SSH into your device.
2. Run the following scripts to verify that the dependencies are correctly installed.

```
sudo python3.7 -c "import dlr; print(dlr.__version__)"
```

```
sudo python3.7 -c "import numpy; print(numpy.__version__)"
```

The output prints the version number and the script should exit without error.

Option 2: Run the IDT dependency test

1. Make sure that `device.json` is configured for ML qualification. For more information, see [the section called “Configure device.json for ML qualification”](#).
2. Run the dependencies test for the framework.

```
devicetester_[linux | mac | win_x86-64] run-suite --group-id mldependencies --test-id dlr_dependency_check
```

The test summary displays a PASSED result for mldependencies.

Compile the DLR model

You must compile the DLR model before you can use it for ML qualification tests. For steps, choose one of the following options.

Option 1: Use Amazon SageMaker AI to compile the model


Follow these steps to use SageMaker AI to compile the ML model provided by IDT. This model is pretrained with Apache MXNet.

1. Verify that your device type is supported by SageMaker AI. For more information, see the [target device options](#) the *Amazon SageMaker AI API Reference*. If your device type is not currently supported by SageMaker AI, follow the steps in [the section called "Option 2: Use TVM to compile the DLR model"](#).

Note

Running the DLR test with a model compiled by SageMaker AI might take 4 or 5 minutes. Don't stop IDT during this time.

2. Download the tarball file that contains the uncompiled, pretrained MXNet model for DLR:
 - [dlr-noncompiled-model-1.0.tar.gz](#)
3. Decompress the tarball. This command generates the following directory structure.



4. Move `synset.txt` out of the `resnet18` directory. Make a note of the new location. You copy this file to compiled model directory later.
5. Compress the contents of the `resnet18` directory.

```
tar cvfz model.tar.gz resnet18v1-symbol.json resnet18v1-0000.params
```

6. Upload the compressed file to an Amazon S3 bucket in your AWS account, and then follow the steps in [Compile a Model \(Console\)](#) to create a compilation job.
 - a. For **Input configuration**, use the following values:
 - For **Data input configuration**, enter `{"data": [1, 3, 224, 224]}`.
 - For **Machine learning framework**, choose MXNet.
 - b. For **Output configuration**, use the following values:
 - For **S3 Output location**, enter the path to the Amazon S3 bucket or folder where you want to store the compiled model.
 - For **Target device**, choose your device type.
7. Download the compiled model from the output location you specified, and then unzip the file.
8. Copy `synset.txt` into the compiled model directory.
9. Change the name of the compiled model directory to `resnet18`.

Your compiled model directory must have the following directory structure.



Option 2: Use TVM to compile the DLR model


Follow these steps to use TVM to compile the ML model provided by IDT. This model is pretrained with Apache MXNet, so you must install MXNet on the computer or device where you compile the model. To install MXNet, follow the instructions in the [MXNet documentation](#).

Note

We recommend that you compile the model on your target device. This practice is optional, but it can help ensure compatibility and mitigate potential issues.

1. Download the tarball file that contains the uncompiled, pretrained MXNet model for DLR:
 - [dlr-noncompiled-model-1.0.tar.gz](#)

2. Decompress the tarball. This command generates the following directory structure.



3. Follow the instructions in the TVM documentation to [build and install TVM from source for your platform](#).
4. After TVM is built, run the TVM compilation for the resnet18 model. The following steps are based on [Quick Start Tutorial for Compiling Deep Learning Models](#) in the TVM documentation.
 - a. Open the `relay_quick_start.py` file from the cloned TVM repository.
 - b. Update the code that [defines a neural network in relay](#). You can use one of following options:
 - Option 1: Use `mxnet.gluon.model_zoo.vision.get_model` to get the relay module and parameters:

```
from mxnet.gluon.model_zoo.vision import get_model
block = get_model('resnet18_v1', pretrained=True)
mod, params = relay.frontend.from_mxnet(block, {"data": data_shape})
```

- Option 2: From the uncompiled model that you downloaded in step 1, copy the following files to the same directory as the `relay_quick_start.py` file. These files contain the relay module and parameters.
 - `resnet18v1-symbol.json`
 - `resnet18v1-0000.params`
- c. Update the code that [saves and loads the compiled module](#) to use the following code.

```
from tvn.contrib import util
path_lib = "deploy_lib.so"
# Export the model library based on your device architecture
lib.export_library("deploy_lib.so", cc="aarch64-linux-gnu-g++")
with open("deploy_graph.json", "w") as fo:
```

```
fo.write(graph)
with open("deploy_param.params", "wb") as fo:
    fo.write(relay.save_param_dict(params))
```

d. Build the model:


```
python3 tutorials/relay_quick_start.py --build-dir ./model
```

This command generates the following files.

- `deploy_graph.json`
- `deploy_lib.so`
- `deploy_param.params`

5. Copy the generated model files into a directory named `resnet18`. This is your compiled model directory.
6. Copy the compiled model directory to your host computer. Then copy `synset.txt` from the uncompiled model that you downloaded in step 1 into the compiled model directory.

Your compiled model directory must have the following directory structure.



Next, [configure your AWS credentials and device.json file](#).

Configure IDT settings to run the AWS IoT Greengrass qualification suite

Before you run tests, you must configure settings for AWS credentials and devices on your host computer.

Configure your AWS credentials

You must configure your IAM user credentials in the `<device-tester-extract-location> / configs/config.json` file. Use the credentials for the IDT for AWS IoT Greengrass user created

in [the section called “Create and configure an AWS account”](#). You can specify your credentials in one of two ways:

- Credentials file
- Environment variables

Configure AWS credentials with a credentials file

IDT uses the same credentials file as the AWS CLI. For more information, see [Configuration and credential files](#).

The location of the credentials file varies, depending on the operating system you are using:

- macOS, Linux: `~/.aws/credentials`
- Windows: `C:\Users\UserName\.aws\credentials`

Add your AWS credentials to the `credentials` file in the following format:

```
[default]
aws_access_key_id = <your_access_key_id>
aws_secret_access_key = <your_secret_access_key>
```

To configure IDT for AWS IoT Greengrass to use AWS credentials from your `credentials` file, edit your `config.json` file as follows:

```
{
  "awsRegion": "us-west-2",
  "auth": {
    "method": "file",
    "credentials": {
      "profile": "default"
    }
  }
}
```

Note

If you do not use the default AWS profile, be sure to change the profile name in your `config.json` file. For more information, see [Named profiles](#).

Configure AWS credentials with environment variables

Environment variables are variables maintained by the operating system and used by system commands. They are not saved if you close the SSH session. IDT for AWS IoT Greengrass can use the `AWS_ACCESS_KEY_ID` and `AWS_SECRET_ACCESS_KEY` environment variables to store your AWS credentials.

To set these variables on Linux, macOS, or Unix, use **export**:

```
export AWS_ACCESS_KEY_ID=<your_access_key_id>
export AWS_SECRET_ACCESS_KEY=<your_secret_access_key>
```

To set these variables on Windows, use **set**:

```
set AWS_ACCESS_KEY_ID=<your_access_key_id>
set AWS_SECRET_ACCESS_KEY=<your_secret_access_key>
```

To configure IDT to use the environment variables, edit the `auth` section in your `config.json` file. Here is an example:

```
{
  "awsRegion": "us-west-2",
  "auth": {
    "method": "environment"
  }
}
```

Configure device.json

In addition to AWS credentials, IDT for AWS IoT Greengrass needs information about the devices that tests are run on (for example, IP address, login information, operating system, and CPU architecture).

You must provide this information using the `device.json` template located in `<device_tester_extract_location>/configs/device.json`:

Physical device

```
[
  {
    "id": "<pool-id>",
    "sku": "<sku>",
    "features": [
      {
        "name": "os",
        "value": "linux | ubuntu | openwrt"
      },
      {
        "name": "arch",
        "value": "x86_64 | armv6l | armv7l | aarch64"
      },
      {
        "name": "container",
        "value": "yes | no"
      },
      {
        "name": "docker",
        "value": "yes | no"
      },
      {
        "name": "streamManagement",
        "value": "yes | no"
      },
      {
        "name": "hsi",
        "value": "yes | no"
      },
      {
        "name": "ml",
        "value": "mxnet | tensorflow | dlr | mxnet,dlr,tensorflow | no"
      },
      ***** Remove the section below if the device is not qualifying for ML
      *****
      {
        "name": "mlLambdaContainerizationMode",
        "value": "container | process | both"
      }
    ]
  }
]
```

```

    },
    {
      "name": "processor",
      "value": "cpu | gpu"
    },
  ],
  ***** Remove the section below if the device is not qualifying for HSI
  *****
  "hsm": {
    "p11Provider": "/path/to/pkcs11ProviderLibrary",
    "slotLabel": "<slot_label>",
    "slotUserPin": "<slot_pin>",
    "privateKeyLabel": "<key_label>",
    "openSSLEngine": "/path/to/openssl/engine"
  },
  ***** Remove the section below if the device is not qualifying for ML
  *****
  "machineLearning": {
    "dlrModelPath": "/path/to/compiled/dlr/model",
    "environmentVariables": [
      {
        "key": "<environment-variable-name>",
        "value": "<Path:$PATH>"
      }
    ],
  },
  "deviceResources": [
    {
      "name": "<resource-name>",
      "path": "<resource-path>",
      "type": "device | volume"
    }
  ],
},

*****
"kernelConfigLocation": "",
"greengrassLocation": "",
"devices": [
  {
    "id": "<device-id>",

```

```

    "connectivity": {
      "protocol": "ssh",
      "ip": "<ip-address>",
      "port": 22,
      "auth": {
        "method": "pki | password",
        "credentials": {
          "user": "<user-name>",
          "privKeyPath": "/path/to/private/key",
          "password": "<password>"
        }
      }
    }
  }
]
}
]

```

Note

Specify `privKeyPath` only if method is set to `pki`.
Specify `password` only if method is set to `password`.

Docker container

```

[
  {
    "id": "<pool-id>",
    "sku": "<sku>",
    "features": [
      {
        "name": "os",
        "value": "linux | ubuntu | openwrt"
      },
      {
        "name": "arch",
        "value": "x86_64"
      },
      {
        "name": "container",
        "value": "no"
      }
    ]
  }
]

```

```

    },
    {
      "name": "docker",
      "value": "no"
    },
    {
      "name": "streamManagement",
      "value": "yes | no"
    },
    {
      "name": "hsi",
      "value": "no"
    },
    {
      "name": "ml",
      "value": "mxnet | tensorflow | dlr | mxnet,dlr,tensorflow | no"
    },
    ***** Remove the section below if the device is not qualifying for ML
    *****
    {
      "name": "mlLambdaContainerizationMode",
      "value": "process"
    },
    {
      "name": "processor",
      "value": "cpu | gpu"
    },
    *****
  ],
  ***** Remove the section below if the device is not qualifying for ML
  *****
  "machineLearning": {
    "dlrModelPath": "/path/to/compiled/dlr/model",
    "environmentVariables": [
      {
        "key": "<environment-variable-name>",
        "value": "<Path:$PATH>"
      }
    ]
  },
  "deviceResources": [
    {
      "name": "<resource-name>",
      "path": "<resource-path>",

```

```

        "type": "device | volume"
      }
    ]
  },
  *****
  "kernelConfigLocation": "",
  "greengrassLocation": "",
  "devices": [
    {
      "id": "<device-id>",
      "connectivity": {
        "protocol": "docker",
        "containerId": "<container-name | container-id>",
        "containerUser": "<user>"
      }
    }
  ]
}
]

```

All fields that contain values are required as described here:

id

A user-defined alphanumeric ID that uniquely identifies a collection of devices called a *device pool*. Devices that belong to a pool must have identical hardware. When you run a suite of tests, devices in the pool are used to parallelize the workload. Multiple devices are used to run different tests.

sku

An alphanumeric value that uniquely identifies the device under test. The SKU is used to track qualified boards.

Note

If you want to list your board in the AWS Partner Device Catalog, the SKU you specify here must match the SKU that you use in the listing process.

features

An array that contains the device's supported features. All features are required.

os and arch

Supported operating system (OS) and architecture combinations:

- linux, x86_64
- linux, armv6l
- linux, armv7l
- linux, aarch64
- ubuntu, x86_64
- openwrt, armv7l
- openwrt, aarch64

Note

If you use IDT to test AWS IoT Greengrass running in a Docker container, only the x86_64 Docker architecture is supported.

container

Validates that the device meets all of the software and hardware requirements to run Lambda functions in container mode on a Greengrass core.

The valid value is yes or no.

docker

Validates that the device meets all the required technical dependencies to use the Greengrass Docker application deployment connector to run containers

The valid value is yes or no.

streamManagement

Validates that the device meets all of the required technical dependencies to run AWS IoT Greengrass stream manager.

The valid value is yes or no.

`hsi`

Verifies that the provided HSI shared library can interface with the hardware security module (HSM) and implements the required PKCS#11 APIs correctly. The HSM and shared library must be able to sign a CSR, perform TLS operations, and provide the correct key lengths and public key algorithm.

The valid value is yes or no.

`ml`

Validates that the device meets all of the required technical dependencies to perform ML inference locally.

The valid value can be any combination of `mxnet`, `tensorflow`, `d1r`, and `no` (for example, `mxnet,mxnet,tensorflow,mxnet,tensorflow,d1r`, or `no`).

`mlLambdaContainerizationMode`

Validates that the device meets all of the required technical dependencies to perform ML inference in container mode on a Greengrass device.

The valid value is `container`, `process`, or `both`.

`processor`

Validates that the device meets all of the hardware requirements for the specified processor type.

The valid value is `cpu` or `gpu`.

 Note

If you don't want to use the `container`, `docker`, `streamManager`, `hsi`, or `ml` feature, you can set the corresponding value to `no`.

Docker only supports feature qualification for `streamManagement` and `ml`.

`machineLearning`

Optional. Configuration information for ML qualification tests. For more information, see [the section called “Configure device.json for ML qualification”](#).

hsm

Optional. Configuration information for testing with an AWS IoT Greengrass Hardware Security Module (HSM). Otherwise, the `hsm` property should be omitted. For more information, see [Hardware security integration](#).

This property applies only if `connectivity.protocol` is set to `ssh`.

`hsm.p11Provider`

The absolute path to the PKCS#11 implementation's libdl-loadable library.

`hsm.slotLabel`

The slot label used to identify the hardware module.

`hsm.slotUserPin`

The user PIN used to authenticate the AWS IoT Greengrass core to the module.

`hsm.privateKeyLabel`

The label used to identify the key in the hardware module.

`hsm.openSSLEngine`

The absolute path to the OpenSSL engine's `.so` file that enables PKCS#11 support on OpenSSL. Used by the AWS IoT Greengrass OTA update agent.

`devices.id`

A user-defined unique identifier for the device being tested.

`connectivity.protocol`

The communication protocol used to communicate with this device. Currently, the only supported values are `ssh` for physical devices and `docker` for Docker containers.

`connectivity.ip`

The IP address of the device being tested.

This property applies only if `connectivity.protocol` is set to `ssh`.

`connectivity.containerId`

The container ID or name of the Docker container being tested.

This property applies only if `connectivity.protocol` is set to `docker`.

`connectivity.auth`

Authentication information for the connection.

This property applies only if `connectivity.protocol` is set to `ssh`.

`connectivity.auth.method`

The authentication method used to access a device over the given connectivity protocol.

Supported values are:

- `pki`
- `password`

`connectivity.auth.credentials`

The credentials used for authentication.

`connectivity.auth.credentials.password`

The password used for signing in to the device being tested.

This value applies only if `connectivity.auth.method` is set to `password`.

`connectivity.auth.credentials.privKeyPath`

The full path to the private key used to sign in to the device under test.

This value applies only if `connectivity.auth.method` is set to `pki`.

`connectivity.auth.credentials.user`

The user name for signing in to the device being tested.

`connectivity.auth.credentials.privKeyPath`

The full path to the private key used to sign in to the device being tested.

`connectivity.port`

Optional. The port number to use for SSH connections.

The default value is 22.

This property only applies if `connectivity.protocol` is set to `ssh`.

`greengrassLocation`

The location of AWS IoT Greengrass Core software on your devices.

For physical devices, this value is only used when you use an existing installation of AWS IoT Greengrass. Use this attribute to tell IDT to use the version of the AWS IoT Greengrass Core software installed on your devices.

When running tests in a Docker container from Docker image or Dockerfile provided by AWS IoT Greengrass, set this value to `/greengrass`.

`kernelConfigLocation`

Optional. The path to the kernel configuration file. AWS IoT Device Tester uses this file to check if the devices have the required kernel features enabled. If not specified, IDT uses the following paths to search for the kernel configuration file: `/proc/config.gz` and `/boot/config-<kernel-version>`. AWS IoT Device Tester uses the first path it finds.

Configure `device.json` for ML qualification

This section describes the optional properties in the device configuration file that apply to ML qualification. If you plan to run tests for ML qualification, you must define the properties that apply to your use case.

You can use the `device-ml.json` template to define the configuration settings for your device. This template contains the optional ML properties. You can also use `device.json` and add the ML qualification properties. These files are located in *<device-tester-extract-location>/configs* and includes ML qualification properties. If you use `device-ml.json`, you must rename the file to `device.json` before you run IDT tests.

For information about device configuration properties that don't apply to ML qualification, see [the section called "Configure device.json"](#).

`ml` in the features array

The ML frameworks that your board supports. This property requires IDT v3.1.0 or later.

- If your board supports only one framework, specify the framework. For example:

```
{
  "name": "ml",
  "value": "mxnet"
}
```

- If your board supports multiple frameworks, specify the frameworks as a comma-separated list. For example:

```
{  
  "name": "ml",  
  "value": "mxnet,tensorflow"  
}
```

`mlLambdaContainerizationMode` in the features array

The [containerization mode](#) that you want to test with. This property requires IDT v3.1.0 or later.

- Choose process to run ML inference code with a non-containerized Lambda function. This option requires AWS IoT Greengrass v1.10.x or later.
- Choose container to run ML inference code with a containerized Lambda function.
- Choose both to run ML inference code with both modes. This option requires AWS IoT Greengrass v1.10.x or later.

`processor` in the features array

Indicates the hardware accelerator that your board supports. This property requires IDT v3.1.0 or later.

- Choose `cpu` if your board uses a CPU as the processor.
- Choose `gpu` if your board uses a GPU as the processor.

`machineLearning`

Optional. Configuration information for ML qualification tests. This property requires IDT v3.1.0 or later.

`d1rModelPath`

Required to use the `d1r` framework. The absolute path to your DLR compiled model directory, which must be named `resnet18`. For more information, see [the section called "Compile the DLR model"](#).

 Note

The following is an example path on macOS: `/Users/<user>/Downloads/resnet18`.

environmentVariables

An array of key-value pairs that can dynamically pass settings to ML inference tests. Optional for CPU devices. You can use this section to add framework-specific environment variables required by your device type. For information about these requirements, see the official website of the framework or the device. For example, to run MXNet inference tests on some devices, the following environment variables might be required.

```
"environmentVariables": [  
  ...  
  {  
    "key": "PYTHONPATH",  
    "value": "$MXNET_HOME/python:$PYTHONPATH"  
  },  
  {  
    "key": "MXNET_HOME",  
    "value": "$HOME/mxnet/"  
  },  
  ...  
]
```

Note

The value field might vary based on your MXNet installation.

If you're testing Lambda functions that run with [containerization](#) on GPU devices, add environment variables for the GPU library. This makes it possible for the GPU to perform computations. To use different GPU libraries, see the official documentation for the library or device.

Note

Configure the following keys if the `m1LambdaContainerizationMode` feature is set to `container` or `both`.

```
"environmentVariables": [  
  {  
    "key": "PATH",
```

```

    "value": "<path/to/software/bin>:$PATH"
  },
  {
    "key": "LD_LIBRARY_PATH",
    "value": "<path/to/ld/lib>"
  },
  ...
]

```

deviceResources

Required by GPU devices. Contains [local resources](#) that can be accessed by Lambda functions. Use this section to add local device and volume resources.

- For device resources, specify "type": "device". For GPU devices, device resources should be GPU-related device files under /dev.

Note

The /dev/shm directory is an exception. It can be configured as a volume resource only.

- For volume resources, specify "type": "volume".

Run the AWS IoT Greengrass qualification suite

After you [set the required configuration](#), you can start the tests. The runtime of the full test suite depends on your hardware. For reference, it takes approximately 30 minutes to complete the full test suite on a Raspberry Pi 3B.

The following example run-suite commands show you how to run the qualification tests for a device pool. A device pool is a set of identical devices.

IDT v3.0.0 and later

Run all test groups in a specified test suite.

```

devicetester_[linux | mac | win_x86-64] run-suite --suite-id GGQ_1.0.0 --pool-id <pool-id>

```

Use the list-suites command to list the test suites that are in the tests folder.

Run a specific test group in a test suite.

```
devicetester_[linux | mac | win_x86-64] run-suite --suite-id GGQ_1.0.0 --group-id <group-id> --pool-id <pool-id>
```

Use the `list-groups` command to list the test groups in a test suite.

Run a specific test case in a test group.

```
devicetester_[linux | mac | win_x86-64] run-suite --group-id <group-id> --test-id <test-id>
```

Run multiple test cases in a test group.

```
devicetester_[linux | mac | win_x86-64] run-suite --group-id <group-id> --test-id <test-id1>,<test-id2>
```

List the test cases in a test group.

```
devicetester_[linux | mac | win_x86-64] list-test-cases --group-id <group-id>
```

The options for the `run-suite` command are optional. For example, you can omit `pool-id` if you have only one device pool defined in your `device.json` file. Or, you can omit `suite-id` if you want to run the latest test suite version in the `tests` folder.

 Note

IDT prompts you if a newer test suite version is available online. For more information, see [the section called “Set the default update behavior”](#).

For more information about `run-suite` and other IDT commands, see [the section called “IDT commands”](#).

IDT v2.3.0 and earlier

Run all test groups in a specified suite.

```
devicetester_[linux | mac | win_x86-64] run-suite --suite-id GGQ_1 --pool-id <pool-id>
```

Run a specific test group.

```
devicetester_[linux | mac | win_x86-64] run-suite --suite-id GGQ_1 --group-id <group-id> --pool-id <pool-id>
```

`suite-id` and `pool-id` are optional if you are running a single test suite on a single device pool. This means that you have only one device pool defined in your `device.json` file.

Check for Greengrass dependencies

We recommend that you run the dependency checker test group to make sure all Greengrass dependencies are installed before you run related test groups. For example:

- Run `ggcdependencies` before running core qualification test groups.
- Run `containerdependencies` before running container-specific test groups.
- Run `dockerdependencies` before running Docker-specific test groups.
- Run `ggcstreammanagementdependencies` before running stream manager-specific test groups.

Set the default update behavior

When you start a test run, IDT checks online for a newer test suite version. If one is available, IDT prompts you to update to the latest available version. You can set the `upgrade-test-suite` (or `u`) flag to control the default update behavior. Valid values are:

- `y`. IDT downloads and uses the latest available version.
- `n` (default). IDT uses the version specified in the `suite-id` option. If `suite-id` is not specified, IDT uses the latest version in the `tests` folder.

If you don't include the `upgrade-test-suite` flag, IDT prompts you when an update is available and waits 30 seconds for your input (`y` or `n`). If no input is entered, it defaults to `n` and continues running the tests.

The following examples show common use cases for this feature:

Automatically use the latest tests available for a test group.

```
devicetester_linux run-suite -u y --group-id mqtt --pool-id DevicePool1
```

Run tests in a specific test suite version.

```
devicetester_linux run-suite -u n --suite-id GGQ_1.0.0 --group-id mqtt --pool-id DevicePool1
```

Prompt for updates at runtime.

```
devicetester_linux run-suite --pool-id DevicePool1
```

IDT for AWS IoT Greengrass commands

The IDT commands are located in the *<device-tester-extract-location>*/bin directory. Use them for the following operations:

IDT v3.0.0 and later

help

Lists information about the specified command.

list-groups

Lists the groups in a given test suite.

list-suites

Lists the available test suites.

list-supported-products

Lists the supported products, in this case AWS IoT Greengrass versions, and test suite versions for the current IDT version.

list-test-cases

Lists the test cases in a given test group. The following option is supported:

- **group-id.** The test group to search for. This option is required and must specify a single group.

run-suite

Runs a suite of tests on a pool of devices. The following are some supported options:

- `suite-id`. The test suite version to run. If not specified, IDT uses the latest version in the `tests` folder.
- `group-id`. The test groups to run, as a comma-separated list. If not specified, IDT runs all test groups in the test suite.
- `test-id`. The test cases to run, as a comma-separated list. When specified, `group-id` must specify a single group.
- `pool-id`. The device pool to test. You must specify a pool if you have multiple device pools defined in your `device.json` file.
- `upgrade-test-suite`. Controls how test suite version updates are handled. Starting in IDT v3.0.0, IDT checks online for updated test suite versions. For more information, see [the section called "Test suite versions"](#).
- `stop-on-first-failure`. Configures IDT to stop execution on the first failure. This option should be used with `group-id` to debug the specified test groups. Do not use this option when running a full test-suite to generate a qualification report.
- `update-idt`. Sets the response for the prompt to update IDT. Y as input stops the test execution if IDT detects there is a newer version. N as input continues the test execution.
- `update-managed-policy`. Y as input stops the test execution if IDT detects that the user's managed policy isn't updated. N as input continues the test execution.

For more information about `run-suite` options, use the `help` option:

```
devicetester_[linux | mac | win_x86-64] run-suite -h
```

IDT v2.3.0 and earlier

help

Lists information about the specified command.

list-groups

Lists the groups in a given test suite.

list-suites

Lists the available test suites.

run-suite

Runs a suite of tests on a pool of devices.

For more information about run-suite options, use the help option:

```
devicetester_[linux | mac | win_x86-64] run-suite -h
```

Understanding results and logs

This section describes how to view and interpret IDT result reports and logs.

Viewing results

While running, IDT writes errors to the console, log files, and test reports. After IDT completes the qualification test suite, it generates two test reports. These reports can be found in *<device-tester-extract-location>/results/<execution-id>/*. Both reports capture the results from the qualification test suite execution.

The `awsiotdevicetester_report.xml` is the qualification test report that you submit to AWS to list your device in the AWS Partner Device Catalog. The report contains the following elements:

- The IDT version.
- The AWS IoT Greengrass version that was tested.
- The SKU and the device pool name specified in the `device.json` file.
- The features of the device pool specified in the `device.json` file.
- The aggregate summary of test results.
- A breakdown of test results by libraries that were tested based on the device features (for example, local resource access, shadow, MQTT, and so on).

The `GGQ_Result.xml` report is in [JUnit XML format](#). You can integrate it into continuous integration and deployment platforms like [Jenkins](#), [Bamboo](#), and so on. The report contains the following elements:

- Aggregate summary of test results.
- Breakdown of test results by the AWS IoT Greengrass functionality that was tested.

Interpreting IDT reports

The report section in `awsiotdevicetester_report.xml` or `awsiotdevicetester_report.xml` lists the tests that were run and the results.

The first XML tag `<testsuites>` contains the summary of the test execution. For example:

```
<testsuites name="GGQ results" time="2299" tests="28" failures="0" errors="0" disabled="0">
```

Attributes used in the `<testsuites>` tag

`name`

The name of the test suite.

`time`

The time, in seconds, it took to run the qualification suite.

`tests`

The number of tests executed.

`failures`

The number of tests that were run, but did not pass.

`errors`

The number of tests that IDT couldn't execute.

`disabled`

This attribute is not used and can be ignored.

The `awsiotdevicetester_report.xml` file contains an `<awsproduct>` tag that contains information about the product being tested and the product features that were validated after running a suite of tests.

Attributes used in the <awsproduct> tag

name

The name of the product being tested.

version

The version of the product being tested.

features

The features validated. Features marked as `required` are required to submit your board for qualification. The following snippet shows how this information appears in the `awsiotdevicetester_report.xml` file.

```
<feature name="aws-iot-greengrass-no-container" value="supported" type="required"></feature>
```

Features marked as `optional` are not required for qualification. The following snippets show optional features.

```
<feature name="aws-iot-greengrass-container" value="supported" type="optional"></feature>
```

```
<feature name="aws-iot-greengrass-hsi" value="not-supported" type="optional"></feature>
```

If there are no test failures or errors for the required features, your device meets the technical requirements to run AWS IoT Greengrass and can interoperate with AWS IoT services. If you want to list your device in the AWS Partner Device Catalog, you can use this report as qualification evidence.

In the event of test failures or errors, you can identify the test that failed by reviewing the `<testsuites>` XML tags. The `<testsuite>` XML tags inside the `<testsuites>` tag show the test result summary for a test group. For example:

```
<testsuite name="combination" package="" tests="1" failures="0" time="161" disabled="0" errors="0" skipped="0">
```

The format is similar to the `<testsuites>` tag, but with a `skipped` attribute that is not used and can be ignored. Inside each `<testsuite>` XML tag, there are `<testcase>` tags for each executed test for a test group. For example:

```
<testcase classname="Security Combination (IPD + DCM) Test Context" name="Security  
Combination IP Change Tests sec4_test_1: Should rotate server cert when IPD disabled  
and following changes are made:Add CIS conn info and Add another CIS conn info"  
attempts="1"></testcase>>
```

Attributes used in the `<testcase>` tag

name

The name of the test.

attempts

The number of times IDT executed the test case.

When a test fails or an error occurs, `<failure>` or `<error>` tags are added to the `<testcase>` tag with information for troubleshooting. For example:

```
<testcase classname="mcu.Full_MQTT" name="AFQP_MQTT_Connect_HappyCase" attempts="1">  
  <failure type="Failure">Reason for the test failure</failure>  
  <error>Reason for the test execution error</error>  
</testcase>
```

Viewing logs

IDT generates logs from test execution in `<devicetester-extract-location>/results/<execution-id>/logs`. Two sets of logs are generated:

`test_manager.log`

Logs generated from the Test Manager component of AWS IoT Device Tester (for example, logs related to configuration, test sequencing, and report generation).

`<test_case_id>.log` (for example, `ota.log`)

Logs of the test group, including logs from the device under test. When a test fails, a `tar.gz` file that contains the logs of the device under test for the test is created (for example, `ota_prod_test_1_ggc_logs.tar.gz`).

For more information, see [IDT for AWS IoT Greengrass troubleshooting](#).

Use IDT to develop and run your own test suites

Starting in IDT v4.0.0, IDT for AWS IoT Greengrass combines a standardized configuration setup and result format with a test suite environment that enables you to develop custom test suites for your devices and device software. You can add custom tests for your own internal validation or provide them to your customers for device verification.

Use IDT to develop and run custom test suites, as follows:

To develop custom test suites

- Create test suites with custom test logic for the Greengrass device that you want to test.
- Provide IDT with your custom test suites to test runners. Include information about specific settings configurations for your test suites.

To run custom test suites

- Set up the device that you want to test.
- Implement the settings configurations as required by the test suites that you want to use.
- Use IDT to run your custom test suites.
- View the test results and execution logs for the tests run by IDT.

Download the latest version of AWS IoT Device Tester for AWS IoT Greengrass

Download the [latest version](#) of IDT and extract the software into a location on your file system where you have read and write permissions.

Note

IDT does not support being run by multiple users from a shared location, such as an NFS directory or a Windows network shared folder. We recommend that you extract the IDT package to a local drive and run the IDT binary on your local workstation. Windows has a path length limitation of 260 characters. If you are using Windows, extract IDT to a root directory like C:\ or D:\ to keep your paths under the 260 character limit.

Test suite creation workflow

Test suites are composed of three types of files:

- JSON configuration files that provide IDT with information on how to execute the test suite.
- Test executable files that IDT uses to run test cases.
- Additional files required to run tests.

Complete the following basic steps to create custom IDT tests:

1. [Create JSON configuration files](#) for your test suite.
2. [Create test case executables](#) that contain the test logic for your test suite.
3. Verify and document the [configuration information required for test runners](#) to run the test suite.
4. Verify that IDT can run your test suite and produce [test results](#) as expected.

To quickly build a sample custom suite and run it, follow the instructions in [Tutorial: Build and run the sample IDT test suite](#).

To get started creating a custom test suite in Python, see [Tutorial: Develop a simple IDT test suite](#).

Tutorial: Build and run the sample IDT test suite

The AWS IoT Device Tester download includes the source code for a sample test suite. You can complete this tutorial to build and run the sample test suite to understand how you can use AWS IoT Device Tester for AWS IoT Greengrass to run custom test suites.

In this tutorial, you will complete the following steps:

1. [Build the sample test suite](#)
2. [Use IDT to run the sample test suite](#)

Prerequisites

To complete this tutorial, you need the following:

- **Host computer requirements**

- Latest version of AWS IoT Device Tester
- [Python](#) 3.7 or later

To check the version of Python installed on your computer, run the following command:

```
python3 --version
```

On Windows, if using this command returns an error, then use `python --version` instead. If the returned version number is 3.7 or greater, then run the following command in a Powershell terminal to set `python3` as an alias for your `python` command.

```
Set-Alias -Name "python3" -Value "python"
```

If no version information is returned or if the version number is less than 3.7, follow the instructions in [Downloading Python](#) to install Python 3.7+. For more information, see the [Python documentation](#).

- [urllib3](#)

To verify that `urllib3` is installed correctly, run the following command:

```
python3 -c 'import urllib3'
```

If `urllib3` is not installed, run the following command to install it:

```
python3 -m pip install urllib3
```

- **Device requirements**

- A device with a Linux operating system and a network connection to the same network as your host computer.

We recommend that you use a [Raspberry Pi](#) with Raspberry Pi OS. Make sure you set up [SSH](#) on your Raspberry Pi to remotely connect to it.

Configure device information for IDT

Configure your device information for IDT to run the test. You must update the `device.json` template located in the `<device-tester-extract-location>/configs` folder with the following information.

```
[
  {
    "id": "pool",
    "sku": "N/A",
    "devices": [
      {
        "id": "<device-id>",
        "connectivity": {
          "protocol": "ssh",
          "ip": "<ip-address>",
          "port": "<port>",
          "auth": {
            "method": "pki | password",
            "credentials": {
              "user": "<user-name>",
              "privKeyPath": "/path/to/private/key",
              "password": "<password>"
            }
          }
        }
      }
    ]
  }
]
```

In the `devices` object, provide the following information:

`id`

A user-defined unique identifier for your device.

`connectivity.ip`

The IP address of your device.

`connectivity.port`

Optional. The port number to use for SSH connections to your device.

`connectivity.auth`

Authentication information for the connection.

This property applies only if `connectivity.protocol` is set to `ssh`.

`connectivity.auth.method`

The authentication method used to access a device over the given connectivity protocol.

Supported values are:

- `pki`
- `password`

`connectivity.auth.credentials`

The credentials used for authentication.

`connectivity.auth.credentials.user`

The user name used to sign in to your device.

`connectivity.auth.credentials.privKeyPath`

The full path to the private key used to sign in to your device.

This value applies only if `connectivity.auth.method` is set to `pki`.

`devices.connectivity.auth.credentials.password`

The password used for signing in to your device.

This value applies only if `connectivity.auth.method` is set to `password`.

Note

Specify `privKeyPath` only if method is set to `pki`.

Specify `password` only if method is set to `password`.

Build the sample test suite

The `<device-tester-extract-location>/samples/python` folder contains sample configuration files, source code, and the IDT Client SDK that you can combine into a test suite using the provided build scripts. The following directory tree shows the location of these sample files:

```
<device-tester-extract-location>
### ...
### tests
### samples
#   ### ...
#   ### python
#       ### configuration
#       ### src
#       ### build-scripts
#           ### build.sh
#           ### build.ps1
### sdks
    ### ...
    ### python
        ### idt_client
```

To build the test suite, run the following commands on your host computer:

Windows

```
cd <device-tester-extract-location>/samples/python/build-scripts
./build.ps1
```

Linux, macOS, or UNIX

```
cd <device-tester-extract-location>/samples/python/build-scripts
./build.sh
```

This creates the sample test suite in the IDTSampleSuitePython_1.0.0 folder within the `<device-tester-extract-location>/tests` folder. Review the files in the IDTSampleSuitePython_1.0.0 folder to understand how the sample test suite is structured and see various examples of test case executables and test configuration JSON files.

Next step: Use IDT to [run the sample test suite](#) that you created.

Use IDT to run the sample test suite

To run the sample test suite, run the following commands on your host computer:

```
cd <device-tester-extract-location>/bin
```

```
./devicetester_[linux | mac | win_x86-64] run-suite --suite-id IDTSampleSuitePython
```

IDT runs the sample test suite and streams the results to the console. When the test has finished running, you see the following information:

```
===== Test Summary =====
Execution Time:           5s
Tests Completed:         4
Tests Passed:            4
Tests Failed:            0
Tests Skipped:           0
-----
Test Groups:
  sample_group:          PASSED
-----
Path to IoT Device Tester Report: /path/to/devicetester/
results/87e673c6-1226-11eb-9269-8c8590419f30/awsiotdevicetester_report.xml
Path to Test Execution Logs: /path/to/devicetester/
results/87e673c6-1226-11eb-9269-8c8590419f30/logs
Path to Aggregated JUnit Report: /path/to/devicetester/
results/87e673c6-1226-11eb-9269-8c8590419f30/IDTSampleSuitePython_Report.xml
```

Troubleshooting

Use the following information to help resolve any issues with completing the tutorial.

Test case does not run successfully

If the test does not run successfully, IDT streams the error logs to the console that can help you troubleshoot the test run. Make sure that you meet all the [prerequisites](#) for this tutorial.

Cannot connect to the device under test

Verify the following:

- Your `device.json` file contains the correct IP address, port, and authentication information.
- You can connect to your device over SSH from your host computer.

Tutorial: Develop a simple IDT test suite

A test suite combines the following:

- Test executables that contain the test logic
- JSON configuration files that describe the test suite

This tutorial shows you how to use IDT for AWS IoT Greengrass to develop a Python test suite that contains a single test case. In this tutorial, you will complete the following steps:

1. [Create a test suite directory](#)
2. [Create JSON configuration files](#)
3. [Create the test case executable](#)
4. [Run the test suite](#)

Prerequisites

To complete this tutorial, you need the following:

- **Host computer requirements**
 - Latest version of AWS IoT Device Tester
 - [Python](#) 3.7 or later

To check the version of Python installed on your computer, run the following command:

```
python3 --version
```

On Windows, if using this command returns an error, then use `python --version` instead. If the returned version number is 3.7 or greater, then run the following command in a Powershell terminal to set `python3` as an alias for your `python` command.

```
Set-Alias -Name "python3" -Value "python"
```

If no version information is returned or if the version number is less than 3.7, follow the instructions in [Downloading Python](#) to install Python 3.7+. For more information, see the [Python documentation](#).

- [urllib3](#)

To verify that `urllib3` is installed correctly, run the following command:

```
python3 -c 'import urllib3'
```

If `urllib3` is not installed, run the following command to install it:

```
python3 -m pip install urllib3
```

• Device requirements

- A device with a Linux operating system and a network connection to the same network as your host computer.

We recommend that you use a [Raspberry Pi](#) with Raspberry Pi OS. Make sure you set up [SSH](#) on your Raspberry Pi to remotely connect to it.

Create a test suite directory

IDT logically separates test cases into test groups within each test suite. Each test case must be inside a test group. For this tutorial, create a folder called `MyTestSuite_1.0.0` and create the following directory tree within this folder:

```
MyTestSuite_1.0.0
### suite
    ### myTestGroup
        ### myTestCase
```

Create JSON configuration files

Your test suite must contain the following required [JSON configuration files](#):

Required JSON files

`suite.json`

Contains information about the test suite. See [Configure suite.json](#).

`group.json`

Contains information about a test group. You must create a `group.json` file for each test group in your test suite. See [Configure group.json](#).

test.json

Contains information about a test case. You must create a `test.json` file for each test case in your test suite. See [Configure test.json](#).

1. In the `MyTestSuite_1.0.0/suite` folder, create a `suite.json` file with the following structure:

```
{
  "id": "MyTestSuite_1.0.0",
  "title": "My Test Suite",
  "details": "This is my test suite.",
  "userDataRequired": false
}
```

2. In the `MyTestSuite_1.0.0/myTestGroup` folder, create a `group.json` file with the following structure:

```
{
  "id": "MyTestGroup",
  "title": "My Test Group",
  "details": "This is my test group.",
  "optional": false
}
```

3. In the `MyTestSuite_1.0.0/myTestGroup/myTestCase` folder, create a `test.json` file with the following structure:

```
{
  "id": "MyTestCase",
  "title": "My Test Case",
  "details": "This is my test case.",
  "execution": {
    "timeout": 300000,
    "linux": {
      "cmd": "python3",
      "args": [
        "myTestCase.py"
      ]
    },
    "mac": {
      "cmd": "python3",

```

```
        "args": [
            "myTestCase.py"
        ],
    },
    "win": {
        "cmd": "python3",
        "args": [
            "myTestCase.py"
        ]
    }
}
```

The directory tree for your `MyTestSuite_1.0.0` folder should now look like the following:

```
MyTestSuite_1.0.0
### suite
### suite.json
### myTestGroup
### group.json
### myTestCase
### test.json
```

Get the IDT client SDK

You use the [IDT client SDK](#) to enable IDT to interact with the device under test and to report test results. For this tutorial, you will use the Python version of the SDK.

From the `<device-tester-extract-location>/sdks/python/` folder, copy the `idt_client` folder to your `MyTestSuite_1.0.0/suite/myTestGroup/myTestCase` folder.

To verify that the SDK was successfully copied, run the following command.

```
cd MyTestSuite_1.0.0/suite/myTestGroup/myTestCase
python3 -c 'import idt_client'
```

Create the test case executable

Test case executables contain the test logic that you want to run. A test suite can contain multiple test case executables. For this tutorial, you will create only one test case executable.

1. Create the test suite file.

In the `MyTestSuite_1.0.0/suite/myTestGroup/myTestCase` folder, create a `myTestCase.py` file with the following content:

```
from idt_client import *

def main():
    # Use the client SDK to communicate with IDT
    client = Client()

if __name__ == "__main__":
    main()
```

2. Use client SDK functions to add the following test logic to your `myTestCase.py` file:

a. Run an SSH command on the device under test.

```
from idt_client import *

def main():
    # Use the client SDK to communicate with IDT
    client = Client()

    # Create an execute on device request
    exec_req = ExecuteOnDeviceRequest(ExecuteOnDeviceCommand("echo 'hello world'"))

    # Run the command
    exec_resp = client.execute_on_device(exec_req)

    # Print the standard output
    print(exec_resp.stdout)

if __name__ == "__main__":
    main()
```

b. Send the test result to IDT.

```
from idt_client import *

def main():
    # Use the client SDK to communicate with IDT
```

```
client = Client()

# Create an execute on device request
exec_req = ExecuteOnDeviceRequest(ExecuteOnDeviceCommand("echo 'hello
world'"))

# Run the command
exec_resp = client.execute_on_device(exec_req)

# Print the standard output
print(exec_resp.stdout)

# Create a send result request
sr_req = SendResultRequest(TestResult(passed=True))

# Send the result
client.send_result(sr_req)

if __name__ == "__main__":
    main()
```

Configure device information for IDT

Configure your device information for IDT to run the test. You must update the `device.json` template located in the *<device-tester-extract-location>/configs* folder with the following information.

```
[
  {
    "id": "pool",
    "sku": "N/A",
    "devices": [
      {
        "id": "<device-id>",
        "connectivity": {
          "protocol": "ssh",
          "ip": "<ip-address>",
          "port": "<port>",
          "auth": {
            "method": "pki | password",
            "credentials": {
              "user": "<user-name>",
```

```
        "privKeyPath": "/path/to/private/key",  
        "password": "<password>"  
      }  
    }  
  }  
]  
}
```

In the `devices` object, provide the following information:

`id`

A user-defined unique identifier for your device.

`connectivity.ip`

The IP address of your device.

`connectivity.port`

Optional. The port number to use for SSH connections to your device.

`connectivity.auth`

Authentication information for the connection.

This property applies only if `connectivity.protocol` is set to `ssh`.

`connectivity.auth.method`

The authentication method used to access a device over the given connectivity protocol.

Supported values are:

- `pki`
- `password`

`connectivity.auth.credentials`

The credentials used for authentication.

`connectivity.auth.credentials.user`

The user name used to sign in to your device.

`connectivity.auth.credentials.privKeyPath`

The full path to the private key used to sign in to your device.

This value applies only if `connectivity.auth.method` is set to `pki`.

`devices.connectivity.auth.credentials.password`

The password used for signing in to your device.

This value applies only if `connectivity.auth.method` is set to `password`.

Note

Specify `privKeyPath` only if method is set to `pki`.
Specify `password` only if method is set to `password`.

Run the test suite

After you create your test suite, you want to make sure that it functions as expected. Complete the following steps to run the test suite with your existing device pool to do so.

1. Copy your `MyTestSuite_1.0.0` folder into *<device-tester-extract-location>/tests*.
2. Run the following commands:

```
cd <device-tester-extract-location>/bin  
./devicetester_[linux | mac | win_x86-64] run-suite --suite-id MyTestSuite
```

IDT runs your test suite and streams the results to the console. When the test has finished running, you see the following information:

```
time="2020-10-19T09:24:47-07:00" level=info msg=Using pool: pool  
time="2020-10-19T09:24:47-07:00" level=info msg=Using test suite "MyTestSuite_1.0.0"  
for execution  
time="2020-10-19T09:24:47-07:00" level=info msg=b'hello world\n'  
suiteId=MyTestSuite groupId=myTestGroup testCaseId=myTestCase deviceId=my-device  
executionId=9a52f362-1227-11eb-86c9-8c8590419f30
```

```
time="2020-10-19T09:24:47-07:00" level=info msg=All tests finished.
  executionId=9a52f362-1227-11eb-86c9-8c8590419f30
time="2020-10-19T09:24:48-07:00" level=info msg=

===== Test Summary =====
Execution Time:          1s
Tests Completed:        1
Tests Passed:           1
Tests Failed:            0
Tests Skipped:           0
-----
Test Groups:
  myTestGroup:          PASSED
-----
Path to IoT Device Tester Report: /path/to/devicetester/
results/9a52f362-1227-11eb-86c9-8c8590419f30/awsiotdevicetester_report.xml
Path to Test Execution Logs: /path/to/devicetester/
results/9a52f362-1227-11eb-86c9-8c8590419f30/logs
Path to Aggregated JUnit Report: /path/to/devicetester/
results/9a52f362-1227-11eb-86c9-8c8590419f30/MyTestSuite_Report.xml
```

Troubleshooting

Use the following information to help resolve any issues with completing the tutorial.

Test case does not run successfully

If the test does not run successfully, IDT streams the error logs to the console that can help you troubleshoot the test run. Before you check the error logs, verify the following:

- The IDT client SDK is in the correct folder as described in [this step](#).
- You meet all the [prerequisites](#) for this tutorial.

Cannot connect to the device under test

Verify the following:

- Your `device.json` file contains the correct IP address, port, and authentication information.
- You can connect to your device over SSH from your host computer.

Create IDT test suite configuration files

This section describes the formats in which you create JSON configuration files that you include when you write a custom test suite.

Required JSON files

`suite.json`

Contains information about the test suite. See [Configure suite.json](#).

`group.json`

Contains information about a test group. You must create a `group.json` file for each test group in your test suite. See [Configure group.json](#).

`test.json`

Contains information about a test case. You must create a `test.json` file for each test case in your test suite. See [Configure test.json](#).

Optional JSON files

`state_machine.json`

Defines how tests are run when IDT runs the test suite. See [Configure state_machine.json](#).

`userdata_schema.json`

Defines the schema for the [userdata.json file](#) that test runners can include in their setting configuration. The `userdata.json` file is used for any additional configuration information that is required to run the test but is not present in the `device.json` file. See [Configure userdata_schema.json](#).

JSON configuration files are placed in your *<custom-test-suite-folder>* as shown here.

```
<custom-test-suite-folder>
### suite
    ### suite.json
    ### state_machine.json
    ### userdata_schema.json
```

```
### <test-group-folder>
### group.json
### <test-case-folder>
### test.json
```

Configure suite.json

The `suite.json` file sets environment variables and determines whether user data is required to run the test suite. Use the following template to configure your `<custom-test-suite-folder>/suite/suite.json` file:

```
{
  "id": "<suite-name>_<suite-version>",
  "title": "<suite-title>",
  "details": "<suite-details>",
  "userDataRequired": true | false,
  "environmentVariables": [
    {
      "key": "<name>",
      "value": "<value>",
    },
    ...
    {
      "key": "<name>",
      "value": "<value>",
    }
  ]
}
```

All fields that contain values are required as described here:

id

A unique user-defined ID for the test suite. The value of `id` must match the name of the test suite folder in which the `suite.json` file is located. The suite name and suite version must also meet the following requirements:

- `<suite-name>` cannot contain underscores.
- `<suite-version>` is denoted as `x.x.x`, where `x` is a number.

The ID is shown in IDT-generated test reports.

title

A user-defined name for the product or feature being tested by this test suite. The name is displayed in the IDT CLI for test runners.

details

A short description of the purpose of the test suite.

userDataRequired

Defines whether test runners need to include custom information in a `userdata.json` file. If you set this value to `true`, you must also include the [userdata_schema.json file](#) in your test suite folder.

environmentVariables

Optional. An array of environment variables to set for this test suite.

environmentVariables.key

The name of the environment variable.

environmentVariables.value

The value of the environment variable.

Configure group.json

The `group.json` file defines whether a test group is required or optional. Use the following template to configure your `<custom-test-suite-folder>/suite/<test-group>/group.json` file:

```
{
  "id": "<group-id>",
  "title": "<group-title>",
  "details": "<group-details>",
  "optional": true | false,
}
```

All fields that contain values are required as described here:

id

A unique user-defined ID for the test group. The value of `id` must match the name of the test group folder in which the `group.json` file is located, and must not contain underscores (`_`). The ID is used in IDT-generated test reports.

title

A descriptive name for the test group. The name is displayed in the IDT CLI for test runners.

details

A short description of the purpose of the test group.

optional

Optional. Set to `true` to display this test group as an optional group after IDT finishes running required tests. Default value is `false`.

Configure test.json

The `test.json` file determines the test case executables and the environment variables that are used by a test case. For more information about creating test case executables, see [Create IDT test case executables](#).

Use the following template to configure your `<custom-test-suite-folder>/suite/<test-group>/<test-case>/test.json` file:

```
{
  "id": "<test-id>",
  "title": "<test-title>",
  "details": "<test-details>",
  "requireDUT": true | false,
  "requiredResources": [
    {
      "name": "<resource-name>",
      "features": [
        {
          "name": "<feature-name>",
          "version": "<feature-version>",
          "jobSlots": <job-slots>
        }
      ]
    }
  ]
}
```

```
],
"execution": {
  "timeout": <timeout>,
  "mac": {
    "cmd": "/path/to/executable",
    "args": [
      "<argument>"
    ],
  },
},
"linux": {
  "cmd": "/path/to/executable",
  "args": [
    "<argument>"
  ],
},
"win": {
  "cmd": "/path/to/executable",
  "args": [
    "<argument>"
  ]
}
},
"environmentVariables": [
  {
    "key": "<name>",
    "value": "<value>",
  }
]
}
```

All fields that contain values are required as described here:

id

A unique user-defined ID for the test case. The value of `id` must match the name of the test case folder in which the `test.json` file is located, and must not contain underscores (`_`). The ID is used in IDT-generated test reports.

title

A descriptive name for the test case. The name is displayed in the IDT CLI for test runners.

details

A short description of the purpose of the test case.

`requireDUT`

Optional. Set to `true` if a device is required to run this test, otherwise set to `false`. Default value is `true`. Test runners will configure the devices they will use to run the test in their `device.json` file.

`requiredResources`

Optional. An array that provides information about resource devices needed to run this test.

`requiredResources.name`

The unique name to give the resource device when this test is running.

`requiredResources.features`

An array of user-defined resource device features.

`requiredResources.features.name`

The name of the feature. The device feature for which you want to use this device. This name is matched against the feature name provided by the test runner in the `resource.json` file.

`requiredResources.features.version`

Optional. The version of the feature. This value is matched against the feature version provided by the test runner in the `resource.json` file. If a version is not provided, then the feature is not checked. If a version number is not required for the feature, leave this field blank.

`requiredResources.features.jobSlots`

Optional. The number of simultaneous tests that this feature can support. The default value is 1. If you want IDT to use distinct devices for individual features, then we recommend that you set this value to 1.

`execution.timeout`

The amount of time (in milliseconds) that IDT waits for the test to finish running. For more information about setting this value, see [Create IDT test case executables](#).

`execution.os`

The test case executables to run based on the operating system of the host computer that runs IDT. Supported values are `linux`, `mac`, and `win`.

`execution.os.cmd`

The path to the test case executable that you want to run for the specified operating system. This location must be in the system path.

`execution.os.args`

Optional. The arguments to provide to run the test case executable.

`environmentVariables`

Optional. An array of environment variables set for this test case.

`environmentVariables.key`

The name of the environment variable.

`environmentVariables.value`

The value of the environment variable.

 Note

If you specify the same environment variable in the `test.json` file and in the `suite.json` file, the value in the `test.json` file takes precedence.

Configure `state_machine.json`

A state machine is a construct that controls the test suite execution flow. It determines the starting state of a test suite, manages state transitions based on user-defined rules, and continues to transition through those states until it reaches the end state.

If your test suite doesn't include a user-defined state machine, IDT will generate a state machine for you. The default state machine performs the following functions:

- Provides test runners with the ability to select and run specific test groups, instead of the entire test suite.
- If specific test groups are not selected, runs every test group in the test suite in a random order.
- Generates reports and prints a console summary that shows the test results for each test group and test case.

For more information about how the IDT state machine functions, see [Configure the IDT state machine](#).

Configure userdata_schema.json

The `userdata_schema.json` file determines the schema in which test runners provide user data. User data is required if your test suite requires information that is not present in the `device.json` file. For example, your tests might need Wi-Fi network credentials, specific open ports, or certificates that a user must provide. This information can be provided to IDT as an input parameter called `userdata`, the value for which is a `userdata.json` file, that users create in their `<device-tester-extract-location>/config` folder. The format of the `userdata.json` file is based on the `userdata_schema.json` file that you include in the test suite.

To indicate that test runners must provide a `userdata.json` file:

1. In the `suite.json` file, set `userDataRequired` to `true`.
2. In your `<custom-test-suite-folder>`, create a `userdata_schema.json` file.
3. Edit the `userdata_schema.json` file to create a valid [IETF Draft v4 JSON Schema](#).

When IDT runs your test suite, it automatically reads the schema and uses it to validate the `userdata.json` file provided by the test runner. If valid, the contents of the `userdata.json` file are available in both the [IDT context](#) and in the [state machine context](#).

Configure the IDT state machine

A state machine is a construct that controls the test suite execution flow. It determines the starting state of a test suite, manages state transitions based on user-defined rules, and continues to transition through those states until it reaches the end state.

If your test suite doesn't include a user-defined state machine, IDT will generate a state machine for you. The default state machine performs the following functions:

- Provides test runners with the ability to select and run specific test groups, instead of the entire test suite.
- If specific test groups are not selected, runs every test group in the test suite in a random order.
- Generates reports and prints a console summary that shows the test results for each test group and test case.

The state machine for an IDT test suite must meet the following criteria:

- Each state corresponds to an action for IDT to take, such as to run a test group or product a report file.
- Transitioning to a state executes the action associated with the state.
- Each state defines the transition rule for the next state.
- The end state must be either Succeed or Fail.

State machine format

You can use the following template to configure your own *<custom-test-suite-folder>/suite/state_machine.json* file:

```
{
  "Comment": "<description>",
  "StartAt": "<state-name>",
  "States": {
    "<state-name>": {
      "Type": "<state-type>",
      // Additional state configuration
    }

    // Required states
    "Succeed": {
      "Type": "Succeed"
    },
    "Fail": {
      "Type": "Fail"
    }
  }
}
```

All fields that contain values are required as described here:

Comment

A description of the state machine.

StartAt

The name of the state at which IDT starts running the test suite. The value of StartAt must be set to one of the states listed in the States object.

States

An object that maps user-defined state names to valid IDT states. Each States.*state-name* object contains the definition of a valid state mapped to the *state-name*.

The States object must include the Succeed and Fail states. For information about valid states, see [Valid states and state definitions](#).

Valid states and state definitions

This section describes the state definitions of all of the valid states that can be used in the IDT state machine. Some of the following states support configurations at the test case level. However, we recommend that you configure state transition rules at the test group level instead of the test case level unless absolutely necessary.

State definitions

- [RunTask](#)
- [Choice](#)
- [Parallel](#)
- [AddProductFeatures](#)
- [Report](#)
- [LogMessage](#)
- [SelectGroup](#)
- [Fail](#)
- [Succeed](#)

RunTask

The RunTask state runs test cases from a test group defined in the test suite.

```
{
```

```
"Type": "RunTask",
"Next": "<state-name>",
"TestGroup": "<group-id>",
"TestCases": [
    "<test-id>"
],
"ResultVar": "<result-name>"
}
```

All fields that contain values are required as described here:

Next

The name of the state to transition to after executing the actions in the current state.

TestGroup

Optional. The ID of the test group to run. If this value is not specified, then IDT runs the test group that the test runner selects.

TestCases

Optional. An array of test case IDs from the group specified in TestGroup. Based on the values of TestGroup and TestCases, IDT determines the test execution behavior as follows:

- When both TestGroup and TestCases are specified, IDT runs the specified test cases from the test group.
- When TestCases are specified but TestGroup is not specified, IDT runs the specified test cases.
- When TestGroup is specified, but TestCases is not specified, IDT runs all of the test cases within the specified test group.
- When neither TestGroup or TestCases is specified, IDT runs all test cases from the test group that the test runner selects from the IDT CLI. To enable group selection for test runners, you must include both RunTask and Choice states in your statemachine.json file. For an example of how this works, see [Example state machine: Run user-selected test groups](#).

For more information about enabling IDT CLI commands for test runners, see [the section called "Enable IDT CLI commands"](#).

ResultVar

The name of the context variable to set with the results of the test run. Do not specify this value if you did not specify a value for TestGroup. IDT sets the value of the variable that you define in ResultVar to true or false based on the following:

- If the variable name is of the form `text_text_passed`, then the value is set to whether all tests in the first test group passed or were skipped.
- In all other cases, the value is set to whether all tests in all test groups passed or were skipped.

Typically, you will use RunTask state to specify a test group ID without specifying individual test case IDs, so that IDT will run all of the test cases in the specified test group. All test cases that are run by this state run in parallel, in a random order. However, if all of the test cases require a device to run, and only a single device is available, then the test cases will run sequentially instead.

Error handling

If any of the specified test groups or test case IDs are not valid, then this state issues the RunTaskError execution error. If the state encounters an execution error, then it also sets the hasExecutionError variable in the state machine context to true.

Choice

The Choice state lets you dynamically set the next state to transition to based on user-defined conditions.

```
{
  "Type": "Choice",
  "Default": "<state-name>",
  "FallthroughOnError": true | false,
  "Choices": [
    {
      "Expression": "<expression>",
      "Next": "<state-name>"
    }
  ]
}
```

All fields that contain values are required as described here:

Default

The default state to transition to if none of the expressions defined in `Choices` can be evaluated to `true`.

FallthroughOnError

Optional. Specifies the behavior when the state encounters an error in evaluating expressions. Set to `true` if you want to skip an expression if the evaluation results in an error. If no expressions match, then the state machine transitions to the `Default` state. If the `FallthroughOnError` value is not specified, it defaults to `false`.

Choices

An array of expressions and states to determine which state to transition to after executing the actions in the current state.

Choices.Expression

An expression string that evaluates to a boolean value. If the expression evaluates to `true`, then the state machine transitions to the state defined in `Choices.Next`. Expression strings retrieve values from the state machine context and then perform operations on them to arrive at a boolean value. For information about accessing the state machine context, see [State machine context](#).

Choices.Next

The name of the state to transition to if the expression defined in `Choices.Expression` evaluates to `true`.

Error handling

The Choice state can require error handling in the following cases:

- Some variables in the choice expressions don't exist in the state machine context.
- The result of an expression is not a boolean value.
- The result of a JSON lookup is not a string, number, or boolean.

You cannot use a `Catch` block to handle errors in this state. If you want to stop executing the state machine when it encounters an error, you must set `FallthroughOnError` to `false`. However, we

recommend that you set `FallthroughOnError` to `true`, and depending on your use case, do one of the following:

- If a variable you are accessing is expected to not exist in some cases, then use the value of `Default` and additional `Choices` blocks to specify the next state.
- If a variable that you are accessing should always exist, then set the `Default` state to `Fail`.

Parallel

The `Parallel` state lets you define and run new state machines in parallel with each other.

```
{
  "Type": "Parallel",
  "Next": "<state-name>",
  "Branches": [
    <state-machine-definition>
  ]
}
```

All fields that contain values are required as described here:

Next

The name of the state to transition to after executing the actions in the current state.

Branches

An array of state machine definitions to run. Each state machine definition must contain its own `StartAt`, `Succeed`, and `Fail` states. The state machine definitions in this array cannot reference states outside of their own definition.

Note

Because each branch state machine shares the same state machine context, setting variables in one branch and then reading those variables from another branch might result in unexpected behavior.

The `Parallel` state moves to the next state only after it runs all of the branch state machines. Each state that requires a device will wait to run until the device is available. If multiple devices

are available, this state runs test cases from multiple groups in parallel. If enough devices are not available, then test cases will run sequentially. Because test cases are run in a random order when they run in parallel, different devices might be used to run tests from the same test group.

Error handling

Make sure that both the branch state machine and the parent state machine transition to the `Fail` state to handle execution errors.

Because branch state machines do not transmit execution errors to the parent state machine, you cannot use a `Catch` block to handle execution errors in branch state machines. Instead, use the `hasExecutionErrors` value in the shared state machine context. For an example of how this works, see [Example state machine: Run two test groups in parallel](#).

AddProductFeatures

The `AddProductFeatures` state lets you add product features to the `awsiotdevicetester_report.xml` file generated by IDT.

A product feature is user-defined information about specific criteria that a device might meet. For example, the MQTT product feature can designate that the device publishes MQTT messages properly. In the report, product features are set as supported, not-supported, or a custom value, based on whether specified tests passed.

Note

The `AddProductFeatures` state does not generate reports by itself. This state must transition to the [Report state](#) to generate reports.

```
{
  "Type": "Parallel",
  "Next": "<state-name>",
  "Features": [
    {
      "Feature": "<feature-name>",
      "Groups": [
        "<group-id>"
      ],
    }
  ],
}
```

```

        "OneOfGroups": [
            "<group-id>"
        ],
        "TestCases": [
            "<test-id>"
        ],
        "IsRequired": true | false,
        "ExecutionMethods": [
            "<execution-method>"
        ]
    }
}

```

All fields that contain values are required as described here:

Next

The name of the state to transition to after executing the actions in the current state.

Features

An array of product features to show in the `awsiotdevicetester_report.xml` file.

Feature

The name of the feature

FeatureValue

Optional. The custom value to use in the report instead of supported. If this value is not specified, then based on test results, the feature value is set to supported or not-supported.

If you use a custom value for FeatureValue, you can test the same feature with different conditions, and IDT concatenates the feature values for the supported conditions. For example, the following excerpt shows the MyFeature feature with two separate feature values:

```

...
{
    "Feature": "MyFeature",
    "FeatureValue": "first-feature-supported",
    "Groups": ["first-feature-group"]
}

```

```
  },  
  {  
    "Feature": "MyFeature",  
    "FeatureValue": "second-feature-supported",  
    "Groups": ["second-feature-group"]  
  },  
  ...  
}
```

If both test groups pass, then the feature value is set to `first-feature-supported`, `second-feature-supported`.

Groups

Optional. An array of test group IDs. All tests within each specified test group must pass for the feature to be supported.

OneOfGroups

Optional. An array of test group IDs. All tests within at least one of the specified test groups must pass for the feature to be supported.

TestCases

Optional. An array of test case IDs. If you specify this value, then the following apply:

- All of the specified test cases must pass for the feature to be supported.
- Groups must contain only one test group ID.
- `OneOfGroups` must not be specified.

IsRequired

Optional. Set to `false` to mark this feature as an optional feature in the report. The default value is `true`.

ExecutionMethods

Optional. An array of execution methods that match the `protocol` value specified in the `device.json` file. If this value is specified, then test runners must specify a `protocol` value that matches one of the values in this array to include the feature in the report. If this value is not specified, then the feature will always be included in the report.

To use the `AddProductFeatures` state, you must set the value of `ResultVar` in the `RunTask` state to one of the following values:

- If you specified individual test case IDs, then set `ResultVar` to `group-id_test-id_passed`.
- If you did not specify individual test case IDs, then set `ResultVar` to `group-id_passed`.

The `AddProductFeatures` state checks for test results in the following manner:

- If you did not specify any test case IDs, then the result for each test group is determined from the value of the `group-id_passed` variable in the state machine context.
- If you did specify test case IDs, then the result for each of the tests is determined from the value of the `group-id_test-id_passed` variable in the state machine context.

Error handling

If a group ID provided in this state is not a valid group ID, then this state results in the `AddProductFeaturesError` execution error. If the state encounters an execution error, then it also sets the `hasExecutionErrors` variable in the state machine context to `true`.

Report

The `Report` state generates the `suite-name_Report.xml` and `awsiotdevicetester_report.xml` files. This state also streams the report to the console.

```
{
  "Type": "Report",
  "Next": "<state-name>"
}
```

All fields that contain values are required as described here:

Next

The name of the state to transition to after executing the actions in the current state.

You should always transition to the `Report` state towards the end of the test execution flow so that test runners can view test results. Typically, the next state after this state is `Succeed`.

Error handling

If this state encounters issues with generating the reports, then it issues the `ReportError` execution error.

LogMessage

The LogMessage state generates the `test_manager.log` file and streams the log message to the console.

```
{
  "Type": "LogMessage",
  "Next": "<state-name>"
  "Level": "info | warn | error"
  "Message": "<message>"
}
```

All fields that contain values are required as described here:

Next

The name of the state to transition to after executing the actions in the current state.

Level

The error level at which to create the log message. If you specify a level that is not valid, this state generates an error message and discards it.

Message

The message to log.

SelectGroup

The SelectGroup state updates the state machine context to indicate which groups are selected. The values set by this state are used by any subsequent Choice states.

```
{
  "Type": "SelectGroup",
  "Next": "<state-name>"
  "TestGroups": [
    <group-id>
  ]
}
```

All fields that contain values are required as described here:

Next

The name of the state to transition to after executing the actions in the current state.

TestGroups

An array of test groups that will be marked as selected. For each test group ID in this array, the *group-id_selected* variable is set to `true` in the context. Make sure that you provide valid test group IDs because IDT does not validate whether the specified groups exist.

Fail

The `Fail` state indicates that the state machine did not execute correctly. This is an end state for the state machine, and each state machine definition must include this state.

```
{
  "Type": "Fail"
}
```

Succeed

The `Succeed` state indicates that the state machine executed correctly. This is an end state for the state machine, and each state machine definition must include this state.

```
{
  "Type": "Succeed"
}
```

State machine context

The state machine context is a read-only JSON document that contains data that is available to the state machine during execution. The state machine context is accessible only from the state machine, and contains information that determines the test flow. For example, you can use information configured by test runners in the `userdata.json` file to determine whether a specific test is required to run.

The state machine context uses the following format:

```
{
  "pool": {
    <device-json-pool-element>
  }
}
```

```
    },
    "userData": {
      <userdata-json-content>
    },
    "config": {
      <config-json-content>
    },
    "suiteFailed": true | false,
    "specificTestGroups": [
      "<group-id>"
    ],
    "specificTestCases": [
      "<test-id>"
    ],
    "hasExecutionErrors": true
  }
}
```

pool

Information about the device pool selected for the test run. For a selected device pool, this information is retrieved from the corresponding top-level device pool array element defined in the `device.json` file.

userData

Information in the `userdata.json` file.

config

Information in the `config.json` file.

suiteFailed

The value is set to `false` when the state machine starts. If a test group fails in a `RunTask` state, then this value is set to `true` for the remaining duration of the state machine execution.

specificTestGroups

If the test runner selects specific test groups to run instead of the entire test suite, this key is created and contains the list of specific test group IDs.

specificTestCases

If the test runner selects specific test cases to run instead of the entire test suite, this key is created and contains the list of specific test case IDs.

hasExecutionErrors

Does not exit when the state machine starts. If any state encounters an execution errors, this variable is created and set to `true` for the remaining duration of the state machine execution.

You can query the context using JSONPath notation. The syntax for JSONPath queries in state definitions is `{{$.query}}`. You can use JSONPath queries as placeholder strings within some states. IDT replaces the placeholder strings with the value of the evaluated JSONPath query from the context. You can use placeholders for the following values:

- The `TestCases` value in `RunTask` states.
- The `Expression` value `Choice` state.

When you access data from the state machine context, make sure the following conditions are met:

- Your JSON paths must begin with `$`.
- Each value must evaluate to a string, a number, or a boolean.

For more information about using JSONPath notation to access data from the context, see [Use the IDT context](#).

Execution errors

Execution errors are errors in the state machine definition that the state machine encounters when executing a state. IDT logs information about each error in the `test_manager.log` file and streams the log message to the console.

You can use the following methods to handle execution errors:

- Add a [Catch block](#) in the state definition.
- Check the value of the [hasExecutionErrors value](#) in the state machine context.

Catch

To use `Catch`, add the following to your state definition:

```
"Catch": [  
  {
```

```
    "ErrorEquals": [
      "<error-type>"
    ]
    "Next": "<state-name>"
  }
]
```

All fields that contain values are required as described here:

`Catch.ErrorEquals`

An array of the error types to catch. If an execution error matches one of the specified values, then the state machine transitions to the state specified in `Catch.Next`. See each state definition for information about the type of error it produces.

`Catch.Next`

The next state to transition to if the current state encounters an execution error that matches one of the values specified in `Catch.ErrorEquals`.

Catch blocks are handled sequentially until one matches. If the no errors match the ones listed in the Catch blocks, then the state machines continues to execute. Because execution errors are a result of incorrect state definitions, we recommend that you transition to the Fail state when a state encounters an execution error.

hasExecutionError

When some states encounter execution errors, in addition to issuing the error, they also set the `hasExecutionError` value to `true` in the state machine context. You can use this value to detect when an error occurs, and then use a Choice state to transition the state machine to the Fail state.

This method has the following characteristics.

- The state machine does not start with any value assigned to `hasExecutionError`, and this value is not available until a particular state sets it. This means that you must explicitly set the `FallthroughOnError` to `false` for the Choice states that access this value to prevent the state machine from stopping if no execution errors occur.
- Once it is set to `true`, `hasExecutionError` is never set to `false` or removed from the context. This means that this value is useful only the first time that it is set to `true`, and for all subsequent states, it does not provide a meaningful value.

- The `hasExecutionError` value is shared with all branch state machines in the `Parallel` state, which can result in unexpected results depending on the order in which it is accessed.

Because of these characteristics, we do not recommend that you use this method if you can use a `Catch` block instead.

Example state machines

This section provides some example state machine configurations.

Examples

- [Example state machine: Run a single test group](#)
- [Example state machine: Run user-selected test groups](#)
- [Example state machine: Run a single test group with product features](#)
- [Example state machine: Run two test groups in parallel](#)

Example state machine: Run a single test group

This state machine:

- Runs the test group with id `GroupA`, which must be present in the suite in a `group.json` file.
- Checks for execution errors and transitions to `Fail` if any are found.
- Generates a report and transitions to `Succeed` if there are no errors, and `Fail` otherwise.

```
{
  "Comment": "Runs a single group and then generates a report.",
  "StartAt": "RunGroupA",
  "States": {
    "RunGroupA": {
      "Type": "RunTask",
      "Next": "Report",
      "TestGroup": "GroupA",
      "Catch": [
        {
          "ErrorEquals": [
            "RunTaskError"
          ],
          "Next": "Fail"
        }
      ]
    }
  }
}
```

```

        }
    ]
},
"Report": {
    "Type": "Report",
    "Next": "Succeed",
    "Catch": [
        {
            "ErrorEquals": [
                "ReportError"
            ],
            "Next": "Fail"
        }
    ]
},
"Succeed": {
    "Type": "Succeed"
},
"Fail": {
    "Type": "Fail"
}
}
}

```

Example state machine: Run user-selected test groups

This state machine:

- Checks if the test runner selected specific test groups. The state machine does not check for specific test cases because test runners cannot select test cases without also selecting a test group.
- If test groups are selected:
 - Runs the test cases within the selected test groups. To do so, the state machine does not explicitly specify any test groups or test cases in the RunTask state.
 - Generates a report after running all tests and exits.
- If test groups are not selected:
 - Runs tests in test group GroupA.
 - Generates reports and exits.

```

{
  "Comment": "Runs specific groups if the test runner chose to do that, otherwise
runs GroupA.",
  "StartAt": "SpecificGroupsCheck",
  "States": {
    "SpecificGroupsCheck": {
      "Type": "Choice",
      "Default": "RunGroupA",
      "FallthroughOnError": true,
      "Choices": [
        {
          "Expression": "{{$.specificTestGroups[0]}} != ''",
          "Next": "RunSpecificGroups"
        }
      ]
    },
    "RunSpecificGroups": {
      "Type": "RunTask",
      "Next": "Report",
      "Catch": [
        {
          "ErrorEquals": [
            "RunTaskError"
          ],
          "Next": "Fail"
        }
      ]
    },
    "RunGroupA": {
      "Type": "RunTask",
      "Next": "Report",
      "TestGroup": "GroupA",
      "Catch": [
        {
          "ErrorEquals": [
            "RunTaskError"
          ],
          "Next": "Fail"
        }
      ]
    },
    "Report": {
      "Type": "Report",

```

```

        "Next": "Succeed",
        "Catch": [
            {
                "ErrorEquals": [
                    "ReportError"
                ],
                "Next": "Fail"
            }
        ],
        "Succeed": {
            "Type": "Succeed"
        },
        "Fail": {
            "Type": "Fail"
        }
    }
}

```

Example state machine: Run a single test group with product features

This state machine:

- Runs the test group GroupA.
- Checks for execution errors and transitions to Fail if any are found.
- Adds the FeatureThatDependsOnGroupA feature to the `awsiotdevicetester_report.xml` file:
 - If GroupA passes, the feature is set to supported.
 - The feature is not marked optional in the report.
- Generates a report and transitions to Succeed if there are no errors, and Fail otherwise

```

{
    "Comment": "Runs GroupA and adds product features based on GroupA",
    "StartAt": "RunGroupA",
    "States": {
        "RunGroupA": {
            "Type": "RunTask",
            "Next": "AddProductFeatures",
            "TestGroup": "GroupA",
            "ResultVar": "GroupA_passed",

```



```

        "Catch": [
            {
                "ErrorEquals": [
                    "RunTaskError"
                ],
                "Next": "Fail"
            }
        ],
    },
    "AddProductFeatures": {
        "Type": "AddProductFeatures",
        "Next": "Report",
        "Features": [
            {
                "Feature": "FeatureThatDependsOnGroupA",
                "Groups": [
                    "GroupA"
                ],
                "IsRequired": true
            }
        ]
    },
    "Report": {
        "Type": "Report",
        "Next": "Succeed",
        "Catch": [
            {
                "ErrorEquals": [
                    "ReportError"
                ],
                "Next": "Fail"
            }
        ]
    },
    "Succeed": {
        "Type": "Succeed"
    },
    "Fail": {
        "Type": "Fail"
    }
}

```

Example state machine: Run two test groups in parallel

This state machine:

- Runs the GroupA and GroupB test groups in parallel. The ResultVar variables stored in the context by the RunTask states in the branch state machines by are available to the AddProductFeatures state.
- Checks for execution errors and transitions to Fail if any are found. This state machine does not use a Catch block because that method does not detect execution errors in branch state machines.
- Adds features to the awsiotdevicetester_report.xml file based on the groups that pass
 - If GroupA passes, the feature is set to supported.
 - The feature is not marked optional in the report.
- Generates a report and transitions to Succeed if there are no errors, and Fail otherwise

If two devices are configured in the device pool, both GroupA and GroupB can run at the same time. However, if either GroupA or GroupB has multiple tests in it, then both devices may be allocated to those tests. If only one device is configured, the test groups will run sequentially.

```
{
  "Comment": "Runs GroupA and GroupB in parallel",
  "StartAt": "RunGroupAAndB",
  "States": {
    "RunGroupAAndB": {
      "Type": "Parallel",
      "Next": "CheckForErrors",
      "Branches": [
        {
          "Comment": "Run GroupA state machine",
          "StartAt": "RunGroupA",
          "States": {
            "RunGroupA": {
              "Type": "RunTask",
              "Next": "Succeed",
              "TestGroup": "GroupA",
              "ResultVar": "GroupA_passed",
              "Catch": [
                {
                  "ErrorEquals": [
```

```

                "RunTaskError"
            ],
            "Next": "Fail"
        }
    ]
},
"Succeed": {
    "Type": "Succeed"
},
"Fail": {
    "Type": "Fail"
}
}
},
{
    "Comment": "Run GroupB state machine",
    "StartAt": "RunGroupB",
    "States": {
        "RunGroupA": {
            "Type": "RunTask",
            "Next": "Succeed",
            "TestGroup": "GroupB",
            "ResultVar": "GroupB_passed",
            "Catch": [
                {
                    "ErrorEquals": [
                        "RunTaskError"
                    ],
                    "Next": "Fail"
                }
            ]
        },
        "Succeed": {
            "Type": "Succeed"
        },
        "Fail": {
            "Type": "Fail"
        }
    }
}
],
},
"CheckForErrors": {
    "Type": "Choice",

```

```

        "Default": "AddProductFeatures",
        "FallthroughOnError": true,
        "Choices": [
            {
                "Expression": "{{$.hasExecutionErrors}} == true",
                "Next": "Fail"
            }
        ]
    },
    "AddProductFeatures": {
        "Type": "AddProductFeatures",
        "Next": "Report",
        "Features": [
            {
                "Feature": "FeatureThatDependsOnGroupA",
                "Groups": [
                    "GroupA"
                ],
                "IsRequired": true
            },
            {
                "Feature": "FeatureThatDependsOnGroupB",
                "Groups": [
                    "GroupB"
                ],
                "IsRequired": true
            }
        ]
    },
    "Report": {
        "Type": "Report",
        "Next": "Succeed",
        "Catch": [
            {
                "ErrorEquals": [
                    "ReportError"
                ],
                "Next": "Fail"
            }
        ]
    },
    "Succeed": {
        "Type": "Succeed"
    },

```

```
    "Fail": {  
      "Type": "Fail"  
    }  
  }  
}
```

Create IDT test case executables

You can create and place test case executables in a test suite folder in the following ways:

- For test suites that use arguments or environment variables from the `test.json` files to determine which tests to run, you can create a single test case executable for the entire test suite, or a test executable for each test group in the test suite.
- For a test suite where you want to run specific tests based on specified commands, you create one test case executable for each test case in the test suite.

As a test writer, you can determine which approach is appropriate for your use case and structure your test case executable accordingly. Make sure that you provide the correct test case executable path in each `test.json` file, and that the specified executable runs correctly.

When all devices are ready for a test case to run, IDT reads the following files:

- The `test.json` for the selected test case determines the processes to start and the environment variables to set.
- The `suite.json` for the test suite determines the environment variables to set.

IDT starts the required test executable process based on the commands and arguments specified in the `test.json` file, and passes the required environment variables to the process.

Use the IDT Client SDK

The IDT Client SDKs let you simplify how you write test logic in your test executable with API commands that you can use to interact with IDT and your devices under test. IDT currently provides the following SDKs:

- IDT Client SDK for Python
- IDT Client SDK for Go

These SDKs are located in the *<device-tester-extract-location>/sdks* folder. When you create a new test case executable, you must copy the SDK that you want to use to the folder that contains your test case executable and reference the SDK in your code. This section provides a brief description of the available API commands that you can use in your test case executables.

In this section

- [Device interaction](#)
- [IDT interaction](#)
- [Host interaction](#)

Device interaction

The following commands enable you to communicate with the device under test without having to implement any additional device interaction and connectivity management functions.

ExecuteOnDevice

Allows test suites to run shell commands on a device that support SSH or Docker shell connections.

CopyToDevice

Allows test suites to copy a local file from the host machine that runs IDT to a specified location on a device that supports SSH or Docker shell connections.

ReadFromDevice

Allows test suites to read from the serial port of devices that support UART connections.

Note

Because IDT does not manage direct connections to devices that are made using device access information from the context, we recommend using these device interaction API commands in your test case executables. However, if these commands do not meet your test case requirements, then you can retrieve device access information from the IDT context and use it to make a direct connection to the device from the test suite.

To make a direct connection, retrieve the information in the `device.connectivity` and the `resource.devices.connectivity` fields for your device under test and for resource

devices, respectively. For more information about using the IDT context, see [Use the IDT context](#).

IDT interaction

The following commands enable your test suites to communicate with IDT.

`PollForNotifications`

Allows test suites to check for notifications from IDT.

`GetContextValue` and `GetContextString`

Allows test suites to retrieve values from the IDT context. For more information, see [Use the IDT context](#).

`SendResult`

Allows test suites to report test case results to IDT. This command must be called at the end of each test case in a test suite.

Host interaction

The following command enable your test suites to communicate with the host machine.

`PollForNotifications`

Allows test suites to check for notifications from IDT.

`GetContextValue` and `GetContextString`

Allows test suites to retrieve values from the IDT context. For more information, see [Use the IDT context](#).

`ExecuteOnHost`

Allows test suites to run commands on the local machine and lets IDT manage the test case executable lifecycle.

Enable IDT CLI commands

The `run-suite` command IDT CLI provides several options that let test runner customize test execution. To allow test runners to use these options to run your custom test suite, you implement

support for the IDT CLI. If you do not implement support, test runners will still be able to run tests, but some CLI options will not function correctly. To provide an ideal customer experience, we recommend that you implement support for the following arguments for the `run-suite` command in the IDT CLI:

`timeout-multiplier`

Specifies a value greater than 1.0 that will be applied to all timeouts while running tests.

Test runners can use this argument to increase the timeout for the test cases that they want to run. When a test runner specifies this argument in their `run-suite` command, IDT uses it to calculate the value of the `IDT_TEST_TIMEOUT` environment variable and sets the `config.timeoutMultiplier` field in the IDT context. To support this argument, you must do the following:

- Instead of directly using the timeout value from the `test.json` file, read the `IDT_TEST_TIMEOUT` environment variable to obtain the correctly calculated timeout value.
- Retrieve the `config.timeoutMultiplier` value from the IDT context and apply it to long running timeouts.

For more information about exiting early because of timeout events, see [Specify exit behavior](#).

`stop-on-first-failure`

Specifies that IDT should stop running all tests if it encounters a failure.

When a test runner specifies this argument in their `run-suite` command, IDT will stop running tests as soon as it encounters a failure. However, if test cases are running in parallel, then this can lead to unexpected results. To implement support, make sure that if IDT encounters this event, your test logic instructs all running test cases to stop, clean up temporary resources, and report a test result to IDT. For more information about exiting early on failures, see [Specify exit behavior](#).

`group-id` and `test-id`

Specifies that IDT should run only the selected test groups or test cases.

Test runners can use these arguments with their `run-suite` command to specify the following test execution behavior:

- Run all tests inside the specified test groups.
- Run a selection of tests from within a specified test group.

To support these arguments, the state machine for your test suite must include a specific set of `RunTask` and `Choice` states in your state machine. If you are not using a custom state machine, then the default IDT state machine includes the required states for you and you do not need to take additional action. However, if you are using a custom state machine, then use [Example state machine: Run user-selected test groups](#) as a sample to add the required states in your state machine.

For more information about IDT CLI commands, see [Debug and run custom test suites](#).

Write event logs

While the test is running, you send data to `stdout` and `stderr` to write event logs and error messages to the console. For information about the format of console messages, see [Console message format](#).

When the IDT finishes running the test suite, this information is also available in the `test_manager.log` file located in the `<devicetester-extract-location>/results/<execution-id>/logs` folder.

You can configure each test case to write the logs from its test run, including logs from the device under test, to the `<group-id>_<test-id>` file located in the `<device-tester-extract-location>/results/<execution-id>/logs` folder. To do this, retrieve the path to the log file from the IDT context with the `testData.logFilePath` query, create a file at that path, and write the content that you want to it. IDT automatically updates the path based on the test case that is running. If you choose not to create the log file for a test case, then no file is generated for that test case.

You can also set up your text executable to create additional log files as needed in the `<device-tester-extract-location>/logs` folder. We recommend that you specify unique prefixes for log file names so your files don't get overwritten.

Report results to IDT

IDT writes test results to the `awsiotdevicetester_report.xml` and the `suite-name_report.xml` files. These report files are located in `<device-tester-extract-location>/results/<execution-id>/`. Both reports capture the results from the test suite execution. For more information about the schemas that IDT uses for these reports, see [Review IDT test results and logs](#)

To populate the contents of the `suite-name_report.xml` file, you must use the `SendResult` command to report test results to IDT before the test execution finishes. If IDT cannot locate the results of a test, it issues an error for the test case. The following Python excerpt shows the commands to send a test result to IDT:

```
request-variable = SendResultRequest(TestResult(result))
client.send_result(request-variable)
```

If you do not report results through the API, IDT looks for test results in the test artifacts folder. The path to this folder is stored in the `testData.testArtifactsPath` field in the IDT context. In this folder, IDT uses the first alphabetically sorted XML file it locates as the test result.

If your test logic produces JUnit XML results, you can write the test results to an XML file in the artifacts folder to directly provide the results to IDT instead of parsing the results and then using the API to submit them to IDT.

If you use this method, make sure that your test logic accurately summarizes the test results and format your result file in the same format as the `suite-name_report.xml` file. IDT does not perform any validation of the data that you provide, with the following exceptions:

- IDT ignores all properties of the `testsuites` tag. Instead, it calculates the tag properties from other reported test group results.
- At least one `testsuite` tag must exist within `testsuites`.

Because IDT uses the same artifacts folder for all test cases and does not delete result files between test runs, this method might also lead to erroneous reporting if IDT reads the incorrect file. We recommend that you use the same name for the generated XML results file across all test cases to overwrite the results for each test case and make sure that the correct results are available for IDT to use. Although you can use a mixed approach to reporting in your test suite, that is, use an XML result file for some test cases and submit results through the API for others, we do not recommend this approach.

Specify exit behavior

Configure your test executable to always exit with an exit code of 0, even if a test case reports a failure or an error result. Use non-zero exit codes only to indicate that a test case did not run or if the test case executable could not communicate any results to IDT. When IDT receives a non-zero exit code, it marks the test case as having encountered an error that prevented it from running.

IDT might request or expect a test case to stop running before it has finished in the following events. Use this information to configure your test case executable to detect each of these events from the test case:

Timeout

Occurs when a test case runs for longer than the timeout value specified in the `test.json` file. If the test runner used the `timeout-multiplier` argument to specify a timeout multiplier, then IDT calculates the timeout value with the multiplier.

To detect this event, use the `IDT_TEST_TIMEOUT` environment variable. When a test runner launches a test, IDT sets the value of the `IDT_TEST_TIMEOUT` environment variable to the calculated timeout value (in seconds) and passes the variable to the test case executable. You can read the variable value to set an appropriate timer.

Interrupt

Occurs when the test runner interrupts IDT. For example, by pressing **Ctrl+C**.

Because terminals propagate signals to all child processes, you can simply configure a signal handler in your test cases to detect interrupt signals.

Alternatively, you can periodically poll the API to check the value of the `CancellationRequested` boolean in the `PollForNotifications` API response. When IDT receives an interrupt signal, it sets the value of the `CancellationRequested` boolean to `true`.

Stop on first failure

Occurs when a test case that is running in parallel with the current test case fails and the test runner used the `stop-on-first-failure` argument to specify that IDT should stop when it encounters any failure.

To detect this event, you can periodically poll the API to check the value of the `CancellationRequested` boolean in the `PollForNotifications` API response. When IDT encounters a failure and is configured to stop on first failure, it sets the value of the `CancellationRequested` boolean to `true`.

When any of these events occur, IDT waits for 5 minutes for any currently running test cases to finish running. If all running test cases do not exit within 5 minutes, IDT forces each of their

processes to stop. If IDT has not received test results before the processes end, it will mark the test cases as having timed out. As a best practice, you should ensure that your test cases perform the following actions when they encounter one of the events:

1. Stop running normal test logic.
2. Clean up any temporary resources, such as test artifacts on the device under test.
3. Report a test result to IDT, such as a test failure or an error.
4. Exit.

Use the IDT context

When IDT runs a test suite, the test suite can access a set of data that can be used to determine how each test runs. This data is called the IDT context. For example, user data configuration provided by test runners in a `userdata.json` file is made available to test suites in the IDT context.

The IDT context can be considered a read-only JSON document. Test suites can retrieve data from and write data to the context using standard JSON data types like objects, arrays, numbers and so on.

Context schema

The IDT context uses the following format:

```
{
  "config": {
    <config-json-content>
    "timeoutMultiplier": timeout-multiplier
  },
  "device": {
    <device-json-device-element>
  },
  "devicePool": {
    <device-json-pool-element>
  },
  "resource": {
    "devices": [
      {
        <resource-json-device-element>
        "name": "<resource-name>"
      }
    ]
  }
}
```

```

    }
  ]
},
"testData": {
  "awsCredentials": {
    "awsAccessKeyId": "<access-key-id>",
    "awsSecretAccessKey": "<secret-access-key>",
    "awsSessionToken": "<session-token>"
  },
  "logFilePath": "/path/to/log/file"
},
"userData": {
  <userdata-json-content>
}
}

```

config

Information from the [config.json file](#). The config field also contains the following additional field:

`config.timeoutMultiplier`

The multiplier for the any timeout value used by the test suite. This value is specified by the test runner from the IDT CLI. The default value is 1.

device

Information about the device selected for the test run. This information is equivalent to the devices array element in the [device.json file](#) for the selected device.

devicePool

Information about the device pool selected for the test run. This information is equivalent to the top-level device pool array element defined in the `device.json` file for the selected device pool.

resource

Information about resource devices from the `resource.json` file.

`resource.devices`

This information is equivalent to the devices array defined in the `resource.json` file. Each devices element includes the following additional field:

`resource.device.name`

The name of the resource device. This value is set to the `requiredResource.name` value in the `test.json` file.

`testData.awsCredentials`

The AWS credentials used by the test to connect to the AWS cloud. This information is obtained from the `config.json` file.

`testData.logFilePath`

The path to the log file to which the test case writes log messages. The test suite creates this file if it doesn't exist.

`userData`

Information provided by the test runner in the [userdata.json file](#).

Access data in the context

You can query the context using JSONPath notation from your JSON files and from your text executable with the `GetContextValue` and `GetContextString` APIs. The syntax for JSONPath strings to access the IDT context varies as follows:

- In `suite.json` and `test.json`, you use `{{query}}`. That is, do not use the root element `$.` to start your expression.
- In `statemachine.json`, you use `{{$.query}}`.
- In API commands, you use `query` or `{{$.query}}`, depending on the command. For more information, see the inline documentation in the SDKs.

The following table describes the operators in a typical JSONPath expression:

Operator	Description
<code>\$</code>	The root element. Because the top-level context value for IDT is an object, you will typically use <code>\$.</code> to start your queries.
<code>.childName</code>	Accesses the child element with name <code>childName</code> from an object. If applied to an

Operator	Description
	array, yields a new array with this operator applied to each element. The element name is case sensitive. For example, the query to access the <code>awsRegion</code> value in the <code>config</code> object is <code>\$.config.awsRegion</code> .
<code>[start:end]</code>	Filters elements from an array, retrieving items beginning from the <code>start</code> index and going up to the end index, both inclusive.
<code>[index1, index2, ... , indexN]</code>	Filters elements from an array, retrieving items from only the specified indices.
<code>[?(expr)]</code>	Filters elements from an array using the <code>expr</code> expression. This expression must evaluate to a boolean value.

To create filter expressions, use the following syntax:

```
<jsonpath> | <value> operator <jsonpath> | <value>
```

In this syntax:

- `jsonpath` is a JSONPath that uses standard JSON syntax.
- `value` is any custom value that uses standard JSON syntax.
- `operator` is one of the following operators:
 - `<` (Less than)
 - `<=` (Less than or equal to)
 - `==` (Equal to)

If the JSONPath or value in your expression is an array, boolean, or object value, then this is the only supported binary operator that you can use.

- `>=` (Greater than or equal to)
- `>` (Greater than)

- `=~` (Regular expression match). To use this operator in a filter expression, the JSONPath or value on the left side of your expression must evaluate to a string and the right side must be a pattern value that follows the [RE2 syntax](#).

You can use JSONPath queries in the form `{{query}}` as placeholder strings within the `args` and `environmentVariables` fields in `test.json` files and within the `environmentVariables` fields in `suite.json` files. IDT performs a context lookup and populates the fields with the evaluated value of the query. For example, in the `suite.json` file, you can use placeholder strings to specify environment variable values that change with each test case and IDT will populate the environment variables with the correct value for each test case. However, when you use placeholder strings in `test.json` and `suite.json` files, the following considerations apply for your queries:

- You must each occurrence of the `devicePool` key in your query in all lower case. That is, use `devicepool` instead.
- For arrays, you can use only arrays of strings. In addition, arrays use a non-standard `item1, item2, ..., itemN` format. If the array contains only one element, then it is serialized as `item`, making it indistinguishable from a string field.
- You cannot use placeholders to retrieve objects from the context.

Because of these considerations, we recommend that whenever possible, you use the API to access the context in your test logic instead of placeholder strings in `test.json` and `suite.json` files. However, in some cases it might be more convenient to use JSONPath placeholders to retrieve single strings to set as environment variables.

Configure settings for test runners

To run custom test suites, test runners must configure their settings based on the test suite that they want to run. Settings are specified based on JSON configuration file templates located in the `<device-tester-extract-location>/configs/` folder. If required, test runners must also set up AWS credentials that IDT will use to connect to the AWS cloud.

As a test writer, you will need to configure these files to [debug your test suite](#). You must provide instructions to test runners so that they can configure the following settings as needed to run your test suites.

Configure device.json

The `device.json` file contains information about the devices that tests are run on (for example, IP address, login information, operating system, and CPU architecture).

Test runners can provide this information using the following template `device.json` file located in the `<device-tester-extract-location>/configs/` folder.

```
[
  {
    "id": "<pool-id>",
    "sku": "<pool-sku>",
    "features": [
      {
        "name": "<feature-name>",
        "value": "<feature-value>",
        "configs": [
          {
            "name": "<config-name>",
            "value": "<config-value>"
          }
        ]
      }
    ],
    "devices": [
      {
        "id": "<device-id>",
        "connectivity": {
          "protocol": "ssh | uart | docker",
          // ssh
          "ip": "<ip-address>",
          "port": "<port-number>",
          "auth": {
            "method": "pki | password",
            "credentials": {
              "user": "<user-name>",
              // pki
              "privKeyPath": "/path/to/private/key",

              // password
              "password": "<password>",
            }
          }
        }
      }
    ]
  }
]
```

```
        // uart
        "serialPort": "<serial-port>",

        // docker
        "containerId": "<container-id>",
        "containerUser": "<container-user-name>",
    }
}
]
}
```

All fields that contain values are required as described here:

id

A user-defined alphanumeric ID that uniquely identifies a collection of devices called a *device pool*. Devices that belong to a pool must have identical hardware. When you run a suite of tests, devices in the pool are used to parallelize the workload. Multiple devices are used to run different tests.

sku

An alphanumeric value that uniquely identifies the device under test. The SKU is used to track qualified devices.

Note

If you want to list your board in the AWS Partner Device Catalog, the SKU you specify here must match the SKU that you use in the listing process.

features

Optional. An array that contains the device's supported features. Device features are user-defined values that you configure in your test suite. You must provide your test runners with information about the feature names and values to include in the `device.json` file. For example, if you want to test a device that functions as an MQTT server for other devices, then you can configure your test logic to validate specific supported levels for a feature named `MQTT_QOS`. Test runners provide this feature name and set the feature value to the QOS

levels supported by their device. You can retrieve the provided information from the [IDT context](#) with the `devicePool.features` query, or from the [state machine context](#) with the `pool.features` query.

`features.name`

The name of the feature.

`features.value`

The supported feature values.

`features.configs`

Configuration settings, if needed, for the feature.

`features.config.name`

The name of the configuration setting.

`features.config.value`

The supported setting values.

`devices`

An array of devices in the pool to be tested. At least one device is required.

`devices.id`

A user-defined unique identifier for the device being tested.

`connectivity.protocol`

The communication protocol used to communicate with this device. Each device in a pool must use the same protocol.

Currently, the only supported values are `ssh` and `uart` for physical devices, and `docker` for Docker containers.

`connectivity.ip`

The IP address of the device being tested.

This property applies only if `connectivity.protocol` is set to `ssh`.

`connectivity.port`

Optional. The port number to use for SSH connections.

The default value is 22.

This property applies only if `connectivity.protocol` is set to `ssh`.

`connectivity.auth`

Authentication information for the connection.

This property applies only if `connectivity.protocol` is set to `ssh`.

`connectivity.auth.method`

The authentication method used to access a device over the given connectivity protocol.

Supported values are:

- `pki`
- `password`

`connectivity.auth.credentials`

The credentials used for authentication.

`connectivity.auth.credentials.password`

The password used for signing in to the device being tested.

This value applies only if `connectivity.auth.method` is set to `password`.

`connectivity.auth.credentials.privKeyPath`

The full path to the private key used to sign in to the device under test.

This value applies only if `connectivity.auth.method` is set to `pki`.

`connectivity.auth.credentials.user`

The user name for signing in to the device being tested.

`connectivity.serialPort`

Optional. The serial port to which the device is connected.

This property applies only if `connectivity.protocol` is set to `uart`.

`connectivity.containerId`

The container ID or name of the Docker container being tested.

This property applies only if `connectivity.protocol` is set to `docker`.

`connectivity.containerUser`

Optional. The name of the user to user inside the container. The default value is the user provided in the Dockerfile.

The default value is 22.

This property applies only if `connectivity.protocol` is set to `docker`.

Note

To check if test runners configure the incorrect device connection for a test, you can retrieve `pool.Devices[0].Connectivity.Protocol` from the state machine context and compare it to the expected value in a Choice state. If an incorrect protocol is used, then print a message using the LogMessage state and transition to the Fail state.

Alternatively, you can use error handling code to report a test failure for incorrect device types.

(Optional) Configure userdata.json

The `userdata.json` file contains any additional information that is required by a test suite but is not specified in the `device.json` file. The format of this file depends on the [userdata_scheme.json file](#) that is defined in the test suite. If you are a test writer, make sure you provide this information to users who will run the test suites that you write.

(Optional) Configure resource.json

The `resource.json` file contains information about any devices that will be used as resource devices. Resource devices are devices that are required to test certain capabilities of a device under test. For example, to test a device's Bluetooth capability, you might use a resource device to test that your device can connect to it successfully. Resource devices are optional, and you can require as many resources devices as you need. As a test writer, you use the [test.json file](#) to define the resource device features that are required for a test. Test runners then use the `resource.json` file to provide a pool of resource devices that have the required features. Make sure you provide this information to users who will run the test suites that you write.

Test runners can provide this information using the following template `resource.json` file located in the `<device-tester-extract-location>/configs/` folder.

```
[
  {
    "id": "<pool-id>",
    "features": [
      {
        "name": "<feature-name>",
        "version": "<feature-value>",
        "jobSlots": <job-slots>
      }
    ],
    "devices": [
      {
        "id": "<device-id>",
        "connectivity": {
          "protocol": "ssh | uart | docker",
          // ssh
          "ip": "<ip-address>",
          "port": <port-number>,
          "auth": {
            "method": "pki | password",
            "credentials": {
              "user": "<user-name>",
              // pki
              "privKeyPath": "/path/to/private/key",

              // password
              "password": "<password>",
            }
          },
        },
        // uart
        "serialPort": "<serial-port>",

        // docker
        "containerId": "<container-id>",
        "containerUser": "<container-user-name>",
      }
    ]
  }
]
```

All fields that contain values are required as described here:

id

A user-defined alphanumeric ID that uniquely identifies a collection of devices called a *device pool*. Devices that belong to a pool must have identical hardware. When you run a suite of tests, devices in the pool are used to parallelize the workload. Multiple devices are used to run different tests.

features

Optional. An array that contains the device's supported features. The information required in this field is defined in the [test.json files](#) in the test suite and determines which tests to run and how to run those tests. If the test suite does not require any features, then this field is not required.

`features.name`

The name of the feature.

`features.version`

The feature version.

`features.jobSlots`

Setting to indicate how many tests can concurrently use the device. The default value is 1.

devices

An array of devices in the pool to be tested. At least one device is required.

`devices.id`

A user-defined unique identifier for the device being tested.

`connectivity.protocol`

The communication protocol used to communicate with this device. Each device in a pool must use the same protocol.

Currently, the only supported values are `ssh` and `uart` for physical devices, and `docker` for Docker containers.

`connectivity.ip`

The IP address of the device being tested.

This property applies only if `connectivity.protocol` is set to `ssh`.

`connectivity.port`

Optional. The port number to use for SSH connections.

The default value is 22.

This property applies only if `connectivity.protocol` is set to `ssh`.

`connectivity.auth`

Authentication information for the connection.

This property applies only if `connectivity.protocol` is set to `ssh`.

`connectivity.auth.method`

The authentication method used to access a device over the given connectivity protocol.

Supported values are:

- `pki`
- `password`

`connectivity.auth.credentials`

The credentials used for authentication.

`connectivity.auth.credentials.password`

The password used for signing in to the device being tested.

This value applies only if `connectivity.auth.method` is set to `password`.

`connectivity.auth.credentials.privKeyPath`

The full path to the private key used to sign in to the device under test.

This value applies only if `connectivity.auth.method` is set to `pki`.

`connectivity.auth.credentials.user`

The user name for signing in to the device being tested.

`connectivity.serialPort`

Optional. The serial port to which the device is connected.

This property applies only if `connectivity.protocol` is set to `uart`.

`connectivity.containerId`

The container ID or name of the Docker container being tested.

This property applies only if `connectivity.protocol` is set to `docker`.

`connectivity.containerUser`

Optional. The name of the user to use inside the container. The default value is the user provided in the Dockerfile.

The default value is 22.

This property applies only if `connectivity.protocol` is set to `docker`.

(Optional) Configure `config.json`

The `config.json` file contains configuration information for IDT. Typically, test runners will not need to modify this file except to provide their AWS user credentials for IDT, and optionally, an AWS region. If AWS credentials with required permissions are provided AWS IoT Device Tester collects and submits usage metrics to AWS. This is an opt-in feature and is used to improve IDT functionality. For more information, see [IDT usage metrics](#).

Test runners can configure their AWS credentials in one of the following ways:

- **Credentials file**

IDT uses the same credentials file as the AWS CLI. For more information, see [Configuration and credential files](#).

The location of the credentials file varies, depending on the operating system you are using:

- macOS, Linux: `~/.aws/credentials`
- Windows: `C:\Users\UserName\.aws\credentials`

- **Environment variables**

Environment variables are variables maintained by the operating system and used by system commands. Variables defined during an SSH session are not available after that session is closed. IDT can use the `AWS_ACCESS_KEY_ID` and `AWS_SECRET_ACCESS_KEY` environment variables to store AWS credentials

To set these variables on Linux, macOS, or Unix, use **export**:

```
export AWS_ACCESS_KEY_ID=<your_access_key_id>
export AWS_SECRET_ACCESS_KEY=<your_secret_access_key>
```

To set these variables on Windows, use **set**:

```
set AWS_ACCESS_KEY_ID=<your_access_key_id>
set AWS_SECRET_ACCESS_KEY=<your_secret_access_key>
```

To configure AWS credentials for IDT, test runners edit the auth section in the `config.json` file located in the `<device-tester-extract-location>/configs/` folder.

```
{
  "log": {
    "location": "logs"
  },
  "configFiles": {
    "root": "configs",
    "device": "configs/device.json"
  },
  "testPath": "tests",
  "reportPath": "results",
  "awsRegion": "<region>",
  "auth": {
    "method": "file | environment",
    "credentials": {
      "profile": "<profile-name>"
    }
  }
}
```

All fields that contain values are required as described here:

Note

All paths in this file are defined relative to the `<device-tester-extract-location>`.

`log.location`

The path to the logs folder in the *<device-tester-extract-location>*.

`configFiles.root`

The path to the folder that contains the configuration files.

`configFiles.device`

The path to the `device.json` file.

`testPath`

The path to the folder that contains test suites.

`reportPath`

The path to the folder that will contain test results after IDT runs a test suite.

`awsRegion`

Optional. The AWS region that test suites will use. If not set, then test suites will use the default region specified in each test suite.

`auth.method`

The method IDT uses to retrieve AWS credentials. Supported values are `file` to retrieve credentials from a credentials file, and `environment` to retrieve credentials using environment variables.

`auth.credentials.profile`

The credentials profile to use from the credentials file. This property applies only if `auth.method` is set to `file`.

Debug and run custom test suites

After the [required configuration](#) is set, IDT can run your test suite. The runtime of the full test suite depends on the hardware and the composition of the test suite. For reference, it takes approximately 30 minutes to complete the full AWS IoT Greengrass qualification test suite on a Raspberry Pi 3B.

As you write your test suite, you can use IDT to run the test suite in debug mode to check your code before you run it or provide it to test runners.

Run IDT in debug mode

Because test suites depend on IDT to interact with devices, provide the context, and receive results, you cannot simply debug your test suites in an IDE without any IDT interaction. To do so, the IDT CLI provides the `debug-test-suite` command that lets you run IDT in debug mode. Run the following command to view the available options for `debug-test-suite`:

```
devicetester_[linux | mac | win_x86-64] debug-test-suite -h
```

When you run IDT in debug mode, IDT does not actually launch the test suite or run the state machine; instead, it interacts with your IDE to respond to requests made from the test suite running in the IDE and prints the logs to the console. IDT does not time out and waits to exit until manually interrupted. In debug mode, IDT also does not run the state machine and will not generate any report files. To debug your test suite, you must use your IDE to provide some information that IDT usually obtains from the configuration JSON files. Make sure you provide the following information:

- Environment variables and arguments for each test. IDT will not read this information from `test.json` or `suite.json`.
- Arguments to select resource devices. IDT will not read this information from `test.json`.

To debug your test suites, complete the following steps:

1. Create the setting configuration files that are required to run the test suite. For example, if your test suite requires the `device.json`, `resource.json`, and `user_data.json`, make sure you configure all of them as needed.
2. Run the following command to place IDT in debug mode and select any devices that are required to run the test.

```
devicetester_[linux | mac | win_x86-64] debug-test-suite [options]
```

After you run this command, IDT waits for requests from the test suite and then responds to them. IDT also generates the environment variables that are required for the case process for the IDT Client SDK.

3. In your IDE, use the `run` or `debug` configuration to do the following:
 - a. Set the values of the IDT-generated environment variables.

- b. Set the value of any environment variables or arguments that you specified in your `test.json` and `suite.json` file.
 - c. Set breakpoints as needed.
4. Run the test suite in your IDE.

You can debug and re-run the test suite as many times as needed. IDT does not time out in debug mode.

5. After you complete debugging, interrupt IDT to exit debug mode.

IDT CLI commands to run tests

The following section describes the IDT CLI commands:

IDT v4.0.0

`help`

Lists information about the specified command.

`list-groups`

Lists the groups in a given test suite.

`list-suites`

Lists the available test suites.

`list-supported-products`

Lists the supported products for your version of IDT, in this case AWS IoT Greengrass versions, and AWS IoT Greengrass qualification test suite versions available for the current IDT version.

`list-test-cases`

Lists the test cases in a given test group. The following option is supported:

- `group-id`. The test group to search for. This option is required and must specify a single group.

`run-suite`

Runs a suite of tests on a pool of devices. The following are some commonly used options:

- `suite-id`. The test suite version to run. If not specified, IDT uses the latest version in the `tests` folder.
- `group-id`. The test groups to run, as a comma-separated list. If not specified, IDT runs all test groups in the test suite.
- `test-id`. The test cases to run, as a comma-separated list. When specified, `group-id` must specify a single group.
- `pool-id`. The device pool to test. Test runners must specify a pool if they have multiple device pools defined in your `device.json` file.
- `timeout-multiplier`. Configures IDT to modify the test execution timeout specified in the `test.json` file for a test with a user-defined multiplier.
- `stop-on-first-failure`. Configures IDT to stop execution on the first failure. This option should be used with `group-id` to debug the specified test groups.
- `userdata`. Sets the file that contains user data information required to run the test suite. This is required only if `userdataRequired` is set to `true` in the `suite.json` file for the test suite.

For more information about `run-suite` options, use the `help` option:

```
devicetester_[linux | mac | win_x86-64] run-suite -h
```

debug-test-suite

Run the test suite in debug mode. For more information, see [Run IDT in debug mode](#).

Review IDT test results and logs

This section describes the format in which IDT generates console logs and test reports.

Console message format

AWS IoT Device Tester uses a standard format for printing messages to the console when it starts a test suite. The following excerpt shows an example of a console message generated by IDT.

```
time="2000-01-02T03:04:05-07:00" level=info msg=Using suite: MyTestSuite_1.0.0  
executionId=9a52f362-1227-11eb-86c9-8c8590419f30
```

Most console messages consist of the following fields:

`time`

A full ISO 8601 timestamp for the logged event.

`level`

The message level for the logged event. Typically, the logged message level is one of `info`, `warn`, or `error`. IDT issues a `fatal` or `panic` message if it encounters an expected event that causes it to exit early.

`msg`

The logged message.

`executionId`

A unique ID string for the current IDT process. This ID is used to differentiate between individual IDT runs.

Console messages generated from a test suite provide additional information about the device under test and the test suite, test group, and test cases that IDT runs. The following excerpt shows an example of a console message generated from a test suite.

```
time="2000-01-02T03:04:05-07:00" level=info msg=Hello world! suiteId=MyTestSuite  
groupId=myTestGroup testCaseId=myTestCase deviceId=my-device  
executionId=9a52f362-1227-11eb-86c9-8c8590419f30
```

The test-suite specific part of the console message contains the following fields:

`suiteId`

The name of the test suite currently running.

`groupId`

The ID of the test group currently running.

`testCaseId`

The ID of the test case current running.

`deviceId`

A ID of the device under test that the current test case is using.

To print a test summary to the console when a IDT finishes running a test, you must include a [Report state](#) in your state machine. The test summary contains information about the test suite, the test results for each group that was run, and the locations of the generated logs and report files. The following example shows a test summary message.

```

===== Test Summary =====
Execution Time:      5m00s
Tests Completed:    4
Tests Passed:       3
Tests Failed:       1
Tests Skipped:      0
-----
Test Groups:
  GroupA:           PASSED
  GroupB:           FAILED
-----
Failed Tests:
  Group Name: GroupB
    Test Name: TestB1
      Reason: Something bad happened
-----
Path to IoT Device Tester Report: /path/to/awsiotdevicetester_report.xml
Path to Test Execution Logs: /path/to/logs
Path to Aggregated JUnit Report: /path/to/MyTestSuite_Report.xml

```

AWS IoT Device Tester report schema

`awsiotdevicetester_report.xml` is a signed report that contains the following information:

- The IDT version.
- The test suite version.
- The report signature and key used to sign the report.
- The device SKU and the device pool name specified in the `device.json` file.
- The product version and the device features that were tested.
- The aggregate summary of test results. This information is the same as that contained in the `suite-name_report.xml` file.

```
<apnreport>
```



```

<awsiotdevicetesterversion>idt-version</awsiotdevicetesterversion>
<testsuiteversion>test-suite-version</testsuiteversion>
<signature>signature</signature>
<keyname>keyname</keyname>
<session>
  <testsession>execution-id</testsession>
  <starttime>start-time</starttime>
  <endtime>end-time</endtime>
</session>
<awsproduct>
  <name>product-name</name>
  <version>product-version</version>
  <features>
    <feature name="<feature-name>" value="supported | not-supported | <feature-value>" type="optional | required"/>
  </features>
</awsproduct>
<device>
  <sku>device-sku</sku>
  <name>device-name</name>
  <features>
    <feature name="<feature-name>" value="<feature-value>"/>
  </features>
  <executionMethod>ssh | uart | docker</executionMethod>
</device>
<devenvironment>
  <os name="<os-name>"/>
</devenvironment>
<report>
  <suite-name-report-contents>
</report>
</apnreport>

```

The `awsiotdevicetester_report.xml` file contains an `<awsproduct>` tag that contains information about the product being tested and the product features that were validated after running a suite of tests.

Attributes used in the `<awsproduct>` tag

name

The name of the product being tested.

version

The version of the product being tested.

features

The features validated. Features marked as `required` are required for the test suite to validate the device. The following snippet shows how this information appears in the `awsiotdevicetester_report.xml` file.

```
<feature name="ssh" value="supported" type="required"></feature>
```

Features marked as `optional` are not required for validation. The following snippets show optional features.

```
<feature name="hsi" value="supported" type="optional"></feature>
```

```
<feature name="mqtt" value="not-supported" type="optional"></feature>
```

Test suite report schema

The `<suite-name>_Result.xml` report is in [JUnit XML format](#). You can integrate it into continuous integration and deployment platforms like [Jenkins](#), [Bamboo](#), and so on. The report contains an aggregate summary of test results.

```
<testsuites name="<suite-name> results" time="<run-duration>" tests="<number-of-test>"
failures="<number-of-tests>" skipped="<number-of-tests>" errors="<number-of-tests>"
disabled="0">
  <testsuite name="<test-group-id>" package="" tests="<number-of-tests>"
failures="<number-of-tests>" skipped="<number-of-tests>" errors="<number-of-tests>"
disabled="0">
    <!--success-->
    <testcase classname="<classname>" name="<name>" time="<run-duration>"/>
    <!--failure-->
    <testcase classname="<classname>" name="<name>" time="<run-duration>">
      <failure type="<failure-type>">
        <reason>
      </failure>
    </testcase>
    <!--skipped-->
    <testcase classname="<classname>" name="<name>" time="<run-duration>">
```

```
        <skipped>
            reason
        </skipped>
    </testcase>
    <!--error-->
    <testcase classname="<classname>" name="<name>" time="<run-duration>">
        <error>
            reason
        </error>
    </testcase>
</testsuite>
</testsuites>
```

The report section in both the `awsiotdevicetester_report.xml` or *suite-name_report.xml* lists the tests that were run and the results.

The first XML tag `<testsuites>` contains the summary of the test execution. For example:

```
<testsuites name="MyTestSuite results" time="2299" tests="28" failures="0" errors="0"
disabled="0">
```

Attributes used in the `<testsuites>` tag

name

The name of the test suite.

time

The time, in seconds, it took to run the test suite.

tests

The number of tests executed.

failures

The number of tests that were run, but did not pass.

errors

The number of tests that IDT couldn't execute.

disabled

This attribute is not used and can be ignored.

In the event of test failures or errors, you can identify the test that failed by reviewing the `<testsuites>` XML tags. The `<testsuite>` XML tags inside the `<testsuites>` tag show the test result summary for a test group. For example:

```
<testsuite name="combination" package="" tests="1" failures="0" time="161" disabled="0"
errors="0" skipped="0">
```

The format is similar to the `<testsuites>` tag, but with a `skipped` attribute that is not used and can be ignored. Inside each `<testsuite>` XML tag, there are `<testcase>` tags for each executed test for a test group. For example:

```
<testcase classname="Security Test" name="IP Change Tests" attempts="1"></testcase>>
```

Attributes used in the `<testcase>` tag

`name`

The name of the test.

`attempts`

The number of times IDT executed the test case.

When a test fails or an error occurs, `<failure>` or `<error>` tags are added to the `<testcase>` tag with information for troubleshooting. For example:

```
<testcase classname="mcu.Full_MQTT" name="MQTT_TestCase" attempts="1">
  <failure type="Failure">Reason for the test failure</failure>
  <error>Reason for the test execution error</error>
</testcase>
```

IDT usage metrics

If you provide AWS credentials with required permissions, AWS IoT Device Tester collects and submits usage metrics to AWS. This is an opt-in feature and is used to improve IDT functionality. IDT collects information such as the following:

- The AWS account ID used to run IDT
- The IDT CLI commands used to run tests

- The test suite that are run
- The test suites in the `<device-tester-extract-location>` folder
- The number of devices configured in the device pool
- Test case names and run times
- Test result information, such as whether tests passed, failed, encountered errors, or were skipped
- Product features tested
- IDT exit behavior, such as unexpected or early exits

All of the information that IDT sends is also logged to a `metrics.log` file in the `<device-tester-extract-location>/results/<execution-id>/` folder. You can view the log file to see the information that was collected during a test run. This file is generated only if you choose to collect usage metrics.

To disable metrics collection, you do not need to take additional action. Simply do not store your AWS credentials, and if you do have stored AWS credentials, do not configure the `config.json` file to access them.

Configure your AWS credentials

If you do not already have an AWS account, you must [create one](#). If you already have an AWS account, you simply need to [configure the required permissions](#) for your account that allow IDT to send usage metrics to AWS on your behalf.

Step 1: Create an AWS account

In this step, create and configure an AWS account. If you already have an AWS account, skip to [the section called “Step 2: Configure permissions for IDT”](#).

Sign up for an AWS account

If you do not have an AWS account, complete the following steps to create one.

To sign up for an AWS account

1. Open <https://portal.aws.amazon.com/billing/signup>.
2. Follow the online instructions.

Part of the sign-up procedure involves receiving a phone call or text message and entering a verification code on the phone keypad.

When you sign up for an AWS account, an *AWS account root user* is created. The root user has access to all AWS services and resources in the account. As a security best practice, assign administrative access to a user, and use only the root user to perform [tasks that require root user access](#).

AWS sends you a confirmation email after the sign-up process is complete. At any time, you can view your current account activity and manage your account by going to <https://aws.amazon.com/> and choosing **My Account**.

Create a user with administrative access

After you sign up for an AWS account, secure your AWS account root user, enable AWS IAM Identity Center, and create an administrative user so that you don't use the root user for everyday tasks.

Secure your AWS account root user

1. Sign in to the [AWS Management Console](#) as the account owner by choosing **Root user** and entering your AWS account email address. On the next page, enter your password.

For help signing in by using root user, see [Signing in as the root user](#) in the *AWS Sign-In User Guide*.

2. Turn on multi-factor authentication (MFA) for your root user.

For instructions, see [Enable a virtual MFA device for your AWS account root user \(console\)](#) in the *IAM User Guide*.

Create a user with administrative access

1. Enable IAM Identity Center.

For instructions, see [Enabling AWS IAM Identity Center](#) in the *AWS IAM Identity Center User Guide*.

2. In IAM Identity Center, grant administrative access to a user.

For a tutorial about using the IAM Identity Center directory as your identity source, see [Configure user access with the default IAM Identity Center directory](#) in the *AWS IAM Identity Center User Guide*.

Sign in as the user with administrative access

- To sign in with your IAM Identity Center user, use the sign-in URL that was sent to your email address when you created the IAM Identity Center user.

For help signing in using an IAM Identity Center user, see [Signing in to the AWS access portal](#) in the *AWS Sign-In User Guide*.

Assign access to additional users

1. In IAM Identity Center, create a permission set that follows the best practice of applying least-privilege permissions.

For instructions, see [Create a permission set](#) in the *AWS IAM Identity Center User Guide*.

2. Assign users to a group, and then assign single sign-on access to the group.

For instructions, see [Add groups](#) in the *AWS IAM Identity Center User Guide*.

Step 2: Configure permissions for IDT

In this step, configure the permissions that IDT uses to run tests and collect IDT usage data. You can use the AWS Management Console or AWS Command Line Interface (AWS CLI) to create an IAM policy and a user for IDT, and then attach policies to the user.

- [To Configure Permissions for IDT \(Console\)](#)
- [To Configure Permissions for IDT \(AWS CLI\)](#)

To configure permissions for IDT (console)

Follow these steps to use the console to configure permissions for IDT for AWS IoT Greengrass.

1. Sign in to the [IAM console](#).
2. Create a customer managed policy that grants permissions to create roles with specific permissions.
 - a. In the navigation pane, choose **Policies**, and then choose **Create policy**.
 - b. On the **JSON** tab, replace the placeholder content with the following policy.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "iot-device-tester:SendMetrics"
      ],
      "Resource": "*"
    }
  ]
}
```

- c. Choose **Next: Tags**.
 - d. Choose **Next: Review**.
 - e. For **Name**, enter **IDTUsageMetricsIAMPermissions**. Under **Summary**, review the permissions granted by your policy.
 - f. Choose **Create policy**.
3. Create an IAM user and attach permissions to the user.
 - a. Create an IAM user. Follow steps 1 through 5 in [Creating IAM users \(console\)](#) in the *IAM User Guide*. If you already created an IAM user, skip to the next step.
 - b. Attach the permissions to your IAM user:
 - i. On the **Set permissions** page, choose **Attach existing policies directly**.
 - ii. Search for the **IDTUsageMetricsIAMPermissions** policy that you created in the previous step. Select the check box.
 - c. Choose **Next: Tags**.
 - d. Choose **Next: Review** to view a summary of your choices.
 - e. Choose **Create user**.
 - f. To view the user's access keys (access key IDs and secret access keys), choose **Show** next to the password and access key. To save the access keys, choose **Download.csv** and save the file to a secure location. You use this information later to configure your AWS credentials file.

To configure permissions for IDT (AWS CLI)

Follow these steps to use the AWS CLI to configure permissions for IDT for AWS IoT Greengrass. If you already configured permissions in the console, skip to [the section called “Configure your device to run IDT tests”](#) or [the section called “Optional: Configuring your Docker container”](#).

1. On your computer, install and configure the AWS CLI if it's not already installed. Follow the steps in [Installing the AWS CLI](#) in the *AWS Command Line Interface User Guide*.

Note

The AWS CLI is an open source tool that you can use to interact with AWS services from your command-line shell.

2. Create the following customer managed policy that grants permissions to manage IDT and AWS IoT Greengrass roles.

Linux, macOS, or Unix

```
aws iam create-policy --policy-name IDTUsageMetricsIAMPermissions --policy-  
document '{  
  "Version": "2012-10-17",  
  "Statement": [  
    {  
      "Effect": "Allow",  
      "Action": [  
        "iot-device-tester:SendMetrics"  
      ],  
      "Resource": "*"   
    }  
  ]  
}'
```

Windows command prompt

```
aws iam create-policy --policy-name IDTUsageMetricsIAMPermissions --policy-  
document '{  
  "Version": "2012-10-17",  
  "Statement": [{  
    "Effect": "Allow",  
    "Action": ["iot-device-  
tester:SendMetrics"],  
    "Resource": "*"   
  }]  
}'
```

Note

This step includes a Windows command prompt example because it uses a different JSON syntax than Linux, macOS, or Unix terminal commands.

3. Create an IAM user and attach the permissions required by IDT for AWS IoT Greengrass.

a. Create an IAM user.

```
aws iam create-user --user-name user-name
```

b. Attach the IDTUsageMetricsIAMPermissions policy you created to your IAM user. Replace *user-name* with your IAM user name and *<account-id>* in the command with the ID of your AWS account.

```
aws iam attach-user-policy --user-name user-name --policy-arn  
arn:aws:iam::<account-id>:policy/IDTGreengrassIAMPermissions
```

4. Create a secret access key for the user.

```
aws iam create-access-key --user-name user-name
```

Store the output in a secure location. You use this information later to configure your AWS credentials file.

Provide AWS credentials to IDT

To allow IDT to access your AWS credentials and submit metrics to AWS, do the following:

1. Store the AWS credentials for your IAM user as environment variables or in a credentials file:

a. To use environment variables, run the following command:

```
AWS_ACCESS_KEY_ID=access-key  
AWS_SECRET_ACCESS_KEY=secret-access-key
```

b. To use the credentials file, add the following information to the `.aws/credentials` file:

```
[profile-name]
aws_access_key_id=access-key
aws_secret_access_key=secret-access-key
```

2. Configure the auth section of the `config.json` file. For more information, see [\(Optional\) Configure config.json](#).

IDT for AWS IoT Greengrass troubleshooting

IDT for AWS IoT Greengrass writes these errors to various locations based on the type of errors. Errors are written to the console, log files, and test reports.

Error codes

The following table lists the error codes generated by IDT for AWS IoT Greengrass.

Error code	Error code name	Possible root cause	Troubleshooting
101	InternalError	An internal error occurred.	Check logs under the <i><device-tester-extract-location></i> /results directory. If you cannot debug the issue, contact AWS Developer Support .
102	TimeoutError	The test cannot be completed in a limited time range. This can happen if: • There is a slow network connection between the test machine and device	• Check the network connection and speed. • Make sure that you did not modify any file under the /test directory.

Error code	Error code name	Possible root cause	Troubleshooting
		(for example, if you are using a VPN network). • A slow network delays the communication between the device and cloud. • The <code>timeout</code> field in test configuration files (<code>test.json</code>) has been mistakenly modified.	• Try running the failed test group manually with <code>--group-id</code> flag. • Try running the test suite by increasing the test timeouts. For more information, see Timeout errors.

Error code	Error code name	Possible root cause	Troubleshooting
103	PlatformNotSupport Error	Incorrect OS/architecture combination specified in <code>device.json</code> .	Change your configuration to one of the supported combinations: • Linux, x86_64 • Linux, ARMv6l • Linux, ARMv7l • Linux, AArch64 • Ubuntu, x86_64 • OpenWRT, ARMv7l • OpenWRT, AArch64 For more information, see Configure device.json.

Error code	Error code name	Possible root cause	Troubleshooting
104	VersionNotSupportError	The AWS IoT Greengrass Core software version is not supported by the version of IDT you are using.	Use the device_tester_bin version command to find the supported version of the AWS IoT Greengrass Core software. For example, if you are using macOS, use ./devicetester_mac_x86_64 version. To find the version of AWS IoT Greengrass Core software that you are using: • If you are running tests with preinstalled AWS IoT Greengrass Core software, use SSH to connect to your AWS IoT Greengrass core device and run <path-to-preinstalled-green-grass-location> /greengrass/ggc/core/greengrassd --version •

Error code	Error code name	Possible root cause	Troubleshooting
			If you are running tests with a different version of the AWS IoT Greengrass Core software, go to the devices directory <code><os>/products/greengrass/gcc</code> directory. The AWS IoT Greengrass Core software version is part of the .zip file name. You can test a different version of the AWS IoT Greengrass Core software. For more information, see Getting started with AWS IoT Greengrass.

Error code	Error code name	Possible root cause	Troubleshooting
105	LanguageNotSupport Error	IDT supports Python for AWS IoT Greengrass libraries and SDKs only.	Make sure: • The SDK package under <code>devicetester_greengrass_<os>/products/greengrass/ggsdk</code> is the Python SDK. • The contents of the <code>bin</code> folder under <code>devicetester_greengrass_<os>/tests/GGQ_1.0.0/suite/resources/run.runtimefarm/bin</code> have not been changed.

Error code	Error code name	Possible root cause	Troubleshooting
106	ValidationError	Some fields in <code>device.json</code> or <code>config.json</code> are invalid.	Check the error message on the right side of the error code in the report. - Invalid auth type for device: Specify the correct method to connect to your device. For more information, see the section called "Configure device.json". - Invalid private key path: Specify the correct path to your private key. For more information, see Configure device.json. - Invalid AWS Region: Specify a valid AWS Region in your <code>config.json</code> file. For more information, see AWS service endpoints. - AWS credentials: Set valid AWS

Error code	Error code name	Possible root cause	Troubleshooting
			credentials on your test machine (by using environment variables or the credentials file). Verify that the auth field is configured correctly. For more information, see the section called “Create and configure an AWS account”. • Invalid HSM input: Check your p11Provider , privateKeyLabel , slotLabel , slotUserPin , and openSSLEngine fields in device.json .

Error code	Error code name	Possible root cause	Troubleshooting
107	SSHConnectionFailed	The test machine cannot connect to the configured device.	Verify that the following fields in your <code>device.json</code> file are correct: • <code>ip</code> • <code>user</code> • <code>privKeyPath</code> • <code>password</code> For more information, see Configure device.json.

Error code	Error code name	Possible root cause	Troubleshooting
108	RunCommandError	A test failed to execute a command on the device under test.	Verify that root access is allowed for the configured user in <code>device.json</code> . A password is required by some devices when executing commands with root access. Make sure root access is allowed without a password. For more information, see the documentation for your device. Try running the failing command manually on your device to see if an error occurs.
109	PermissionDeniedError	No root access.	Set root access for the configured user on your device.
110	CreateFileError	Unable to create a file.	Check your device's disk space and directory permissions.
111	CreateDirError	Unable to create a directory.	Check your device's disk space and directory permissions.

Error code	Error code name	Possible root cause	Troubleshooting
112	InvalidPathError	The path to the AWS IoT Greengrass Core software is incorrect.	Verify that the path in the error message is valid. Do not edit any files under the <code>devicetester_greengrass_<os></code> directory.
113	InvalidFileError	A file is invalid.	Verify that the file in the error message is valid.
114	ReadFileError	The specified file cannot be read.	Verify the following: • File permissions are correct. • <code>limits.config</code> allows enough files to be open. • The file specified in the error message exists and is valid. If you are testing on macOS, increase the open files limit. The default limit is 256, which is enough for testing.

Error code	Error code name	Possible root cause	Troubleshooting
115	FileNotFoundError	A required file was not found.	Verify the following: • A compressed Greengrass file exists under <code>devices/ter_green_grass_<os>/products/greengrass/ggc</code>. You can download the AWS IoT Greengrass Core tar file from the AWS IoT Greengrass Core Software downloads page. • The SDK package exists under <code>devices/ter_green_grass_<os>/products/greengrass/ggsdk</code>. • The files under <code>devices/ter_green_grass_<os>/tests</code> have not been modified.

Error code	Error code name	Possible root cause	Troubleshooting
116	OpenFileFailed	Unable to open the specified file.	Verify the following: • The file specified in the error message exists and is valid. • <code>limits.config</code> allows enough files to be open. If you are testing on macOS, increase the open files limit. The default limit is 256, which is enough for testing.
117	WriteFileFailed	Failed to write file (can be the DUT or test machine).	Verify that the directory specified in the error message exists and that you have write permission.
118	FileCleanUpError	A test failed to remove the specified file or directory or to umount the specified file on the remote device.	If the binary file is still running, the file might be locked. End the process and delete the specified file.
119	InvalidInputError	Invalid configuration.	Verify that your <code>suite.json</code> file is valid.

Error code	Error code name	Possible root cause	Troubleshooting
120	InvalidCredentialError	Invalid AWS credentials.	• Verify your AWS credentials. For more information, see the section called “Configure your AWS credentials”. • Check your network connection and rerun the test group. Network problems can also cause this error.
121	AWSSessionError	Failed to create an AWS session.	This error can occur if AWS credentials are invalid or the internet connection is unstable. Try using the AWS CLI to call an AWS API operation.
122	AWSApiCallError	An AWS API error occurred.	This error might be due to a network issue. Check your network before retrying the test group.

Error code	Error code name	Possible root cause	Troubleshooting
123	IpNotExistError	IP address is not included in connectivity information.	Check your internet connection. You can use the AWS IoT Greengrass console to check the connectivity information for the AWS IoT Greengrass core thing that is being used by the test. If there are 10 endpoints included in the connectivity information, you can remove some or all of them and rerun the test. For more information, see Connectivity information .
124	OTAJobNotCompleteError	An OTA job did not complete.	Check your internet connection and retry the OTA test group.

Error code	Error code name	Possible root cause	Troubleshooting
125	CreateGreengrassServiceRoleError	One of the following occurred: - An error occurred while creating a role. - An error occurred while attaching a policy to the AWS IoT Greengrass service role. - The policy associated with the service role is invalid. - An error occurred when associating a role with an AWS account.	Configure the AWS IoT Greengrass service role. For more information, see the section called "Greengrass service role".
126	DependenciesNotPresentError	One or more dependencies required for the specific test are not present on the device.	Check the test log to see which dependencies are missing on your device: <pre><device-tester-extract-location> /results/ <execution-id>/logs/<test-case-name.log></pre>

Error code	Error code name	Possible root cause	Troubleshooting
127	InvalidHSMConfiguration	The provided HSM/ PKCS configuration is incorrect.	In your <code>device.js</code> on file, provide the configuration required to interact with the HSM using PKCS#11.

Error code	Error code name	Possible root cause	Troubleshooting
128	OTAJobNotSucceededError	The OTA job did not succeed.	• If you ran the ota test group individually, run the ggcdependencies test group to verify that all dependencies (such as wget) are present. Then retry the ota test group. • Review the detailed logs under <i><device-tester-extract-location> / results/ <execution-id>/logs/</i> for troubleshooting and error information. Specifically, check the following logs: • Console log (test_manager.log) • OTA test case log (ota_test.log)

Error code	Error code name	Possible root cause	Troubleshooting
			• GGC daemon log (ota_test_ggc_logs.tar.gz) • OTA agent log (ota_test_ota_logs.tar.gz) • Check your internet connectivity and retry the ota test group. • If the problem persists, contact AWS Developer Support.
129	NoConnectivityError	The host agent is failing to connect to internet.	Check your network connection and firewall settings. Retry the test group after the connectivity issue is resolved.
130	NoPermissionError	The IAM user you are using to run IDT for AWS IoT Greengrass does not have permission to create the AWS resources required to run IDT.	See Permissions policy template for the policy template that grants the permissions required to run IDT for AWS IoT Greengrass.

Error code	Error code name	Possible root cause	Troubleshooting
131	LeftoverAgentExist Error	Your device is running AWS IoT Greengrass processes when you attempt to start IDT for AWS IoT Greengrass.	Make sure there is no existing Greengrass daemon running on your device. - You can use this command to stop daemon: sudo ./<absolute-path-to-greengrass-daemon> /greengrassd stop. - You can also terminate the Greengrass daemon by PID. Note If you are using an existing installation of AWS IoT Greengrass configured to start automatically after reboot, you must stop the daemon after reboot

Error code	Error code name	Possible root cause	Troubleshooting
			and before running the test suite.
132	DeviceTimeOffsetError	The device has the incorrect time.	Set your device to the correct time.
133	InvalidMLConfiguration	The provided ML configuration is incorrect.	In your <code>device.js</code> on file, provide the correct configuration required to run ML inference tests. For more information, see the section called "Optional: Configuring your device for ML qualification" .

Resolving IDT for AWS IoT Greengrass errors

When you use IDT, you must get the correct configuration files in place before you run IDT for AWS IoT Greengrass. If you are getting parsing and configuration errors, your first step is to locate and use a configuration template appropriate for your environment.

If you are still having issues, see the following debugging process.

Topics

- [Where do I look for errors?](#)
- [Parsing errors](#)
- [Required parameter missing error](#)
- [Could not start test error](#)
- [Not authorized to access resource error](#)
- [Permission denied errors](#)

- [SSH connection errors](#)
- [Timeout errors](#)
- [Command not found errors while testing](#)
- [Security exception on macOS](#)

Where do I look for errors?

High-level errors are displayed on the console during execution, and a summary of the failed tests with the error is displayed when all tests are complete. `awsiotdevicetester_report.xml` contains a summary of all the errors that caused a test to fail. The log files for each test run are stored in a directory named with an UUID for the test execution that was displayed on the console during the test run.

The test logs directory is located in `<device-tester-extract-location>/results/<execution-id>/logs/`. This directory contains the following files, which are useful for debugging.

File	Description
<code>test_manager.log</code>	All of the logs that were written to the console during the test execution. A summary of the results is located at the end of this file, which includes a list of which tests failed. The warning and error logs in this file can give you some information about the failures.
<code><test-group-id> __<test-name> .log</code>	Detailed logs for the specific test.
<code><test-name> _ggc_logs.tar.gz</code>	A compressed collection of all the logs the AWS IoT Greengrass core daemon generated during the test. For more information, see Troubleshooting AWS IoT Greengrass .
<code><test-name> _ota_logs.tar.gz</code>	A compressed collection of logs generated by the AWS IoT Greengrass OTA agent during the test. For OTA tests only.

File	Description
<code><test-name> _basic_assertion_publisher_ggad_logs.tar.gz</code>	A compressed collection of logs generated by the AWS IoT publisher device during the test.
<code><test-name> _basic_assertion_subscriber_ggad_logs.tar.gz</code>	A compressed collection of logs generated by the AWS IoT subscriber device during the test.

Parsing errors

Occasionally, a typo in a JSON configuration can lead to parsing errors. Most of the time, the issue is a result of omitting a bracket, comma, or quotation mark from your JSON file. IDT performs JSON validation and prints debugging information. It prints the line where the error occurred, the line number, and the column number of the syntax error. This information should be enough to help you fix the error, but if you still cannot locate the error, you can perform validation manually in your IDE, a text editor such as Atom or Sublime, or through an online tool like JSONLint.

Required parameter missing error

Because new features are being added to IDT, changes to the configuration files might be introduced. Using an old configuration file might break your configuration. If this happens, the `<test_case_id>.log` file under `/results/<execution-id>/logs` explicitly lists all missing parameters. IDT also validates your JSON configuration file schemas to ensure that the latest supported version has been used.

Could not start test error

You might encounter errors that point to failures during test start. There are several possible causes, so do the following:

- Make sure that the pool name you included in your execution command actually exists. The pool name is referenced directly from your `device.json` file.
- Make sure that the devices in your pool have correct configuration parameters.

Not authorized to access resource error

You might see the `<user or role> is not authorized to access this resource` error message in the terminal output or in the `test_manager.log`

file under `/results/<execution-id>/logs`. To resolve this issue, attach the `AWSIoTDeviceTesterForGreengrassFullAccess` managed policy to your test user. For more information, see [the section called "Create and configure an AWS account"](#).

Permission denied errors

IDT performs operations on various directories and files in a device under test. Some of these operations require root access. To automate these operations, IDT must be able to run commands with `sudo` without typing a password.

Follow these steps to allow `sudo` access without typing a password.

Note

`user` and `username` refer to the SSH user used by IDT to access the device under test.

1. Use `sudo usermod -aG sudo <ssh-username>` to add your SSH user to the `sudo` group.
2. Sign out and then sign in for changes to take effect.
3. Open `/etc/sudoers` file and add the following line to the end of the file: `<ssh-username> ALL=(ALL) NOPASSWD: ALL`

Note

As a best practice, we recommend that you use `sudo visudo` when you edit `/etc/sudoers`.

SSH connection errors

When IDT cannot connect to a device under test, connection failures are logged in `/results/<execution-id>/logs/<test-case-id>.log`. SSH failure messages appear at the top of this log file because connecting to a device under test is one of the first operations that IDT performs.

Most Windows setups use the PuTTY terminal application to connect to Linux hosts. This application requires that standard PEM private key files are converted into a proprietary Windows format called PPK. When IDT is configured in your `device.json` file, use PEM files only. If you use

a PPK file, IDT cannot create an SSH connection with the AWS IoT Greengrass device and cannot run tests.

Timeout errors

You can increase the timeout for each test by specifying a timeout multiplier, which is applied to the default value of each test's timeout. Any value configured for this flag must be greater than or equal to 1.0.

To use the timeout multiplier, use the flag `--timeout-multiplier` when running the tests. For example:

```
./devicetester_linux run-suite --suite-id GGQ_1.0.0 --pool-id DevicePool1 --timeout-multiplier 2.5
```

For more information, run `run-suite --help`.

Command not found errors while testing

You need an older version of the OpenSSL library (libssl1.0.0) to run tests on AWS IoT Greengrass devices. Most current Linux distributions use libssl version 1.0.2 or later (v1.1.0).

For example, on a Raspberry Pi, run the following commands to install the required version of libssl:

1.

```
wget http://ftp.us.debian.org/debian/pool/main/o/openssl/libssl1.0.0_1.0.2l-1~bpo8+1_armhf.deb
```
2.

```
sudo dpkg -i libssl1.0.0_1.0.2l-1~bpo8+1_armhf.deb
```

Security exception on macOS

When you run IDT on host machine that uses macOS 10.15, the notarization ticket for IDT is not correctly detected and IDT is blocked from being run. To run IDT, you will need to grant a security exception to the `devicetester_mac_x86-64` executable.

To grant a security exception to the IDT executable

1. Launch **System Preferences** from the Apple menu.

2. Choose **Security & Privacy**, then on the **General** tab, click the lock icon to make changes to security settings.
3. Look for the message "devicetester_mac_x86-64" was blocked from use because it is not from an identified developer. and choose **Allow Anyway**.
4. Accept the security warning.

If you have questions about the IDT support policy, contact [AWS Customer Support](#).

Support policy for AWS IoT Device Tester for AWS IoT Greengrass V1

AWS IoT Device Tester (IDT) for AWS IoT Greengrass is a downloadable testing framework that lets you validate and [qualify](#) your AWS IoT Greengrass devices for inclusion in the [AWS Partner Device Catalog](#). We recommend that you use the most recent version of AWS IoT Greengrass and IDT to test or qualify your devices. For more information, see [Supported versions of IDT for AWS IoT Greengrass V2](#) in the *AWS IoT Greengrass Version 2 Developer Guide*.

You can also use any of the supported versions of AWS IoT Greengrass and IDT to test or qualify your devices. Although you can continue to use [unsupported versions of IDT](#), those versions do not receive bug fixes or updates.

Important

As of April 4, 2022, AWS IoT Device Tester (IDT) for AWS IoT Greengrass V1 no longer generates signed qualification reports. You can no longer qualify new AWS IoT Greengrass V1 devices to list in the [AWS Partner Device Catalog](#) through the [AWS Device Qualification Program](#). While you can't qualify Greengrass V1 devices, you can continue to use IDT for AWS IoT Greengrass V1 to test your Greengrass V1 devices. We recommend that you use [IDT for AWS IoT Greengrass V2](#) to qualify and list Greengrass devices in the [AWS Partner Device Catalog](#).

If you have questions about the support policy, contact [AWS Customer Support](#).

Troubleshooting AWS IoT Greengrass

This section provides troubleshooting information and possible solutions to help resolve issues with AWS IoT Greengrass.

For information about AWS IoT Greengrass quotas (limits), see [Service Quotas](#) in the *Amazon Web Services General Reference*.

AWS IoT Greengrass Core issues

If the AWS IoT Greengrass Core software does not start, try the following general troubleshooting steps:

- Make sure that you install the binaries that are appropriate for your architecture. For more information, see [AWS IoT Greengrass Core Software](#).
- Make sure that your core device has local storage available. For more information, see [the section called “Troubleshooting storage issues”](#).
- Check `runtime.log` and `crash.log` for error messages. For more information, see [the section called “Troubleshooting with logs”](#).

Search the following symptoms and errors to find information to help troubleshoot issues with an AWS IoT Greengrass core.

Issues

- [Error: The configuration file is missing the CaPath, CertPath or KeyPath. The Greengrass daemon process with \[pid = <pid>\] died.](#)
- [Error: Failed to parse /<greengrass-root>/config/config.json.](#)
- [Error: Error occurred while generating TLS config: ErrUnknownURIScheme](#)
- [Error: Runtime failed to start: unable to start workers: container test timed out.](#)
- [Error: Failed to invoke PutLogEvents on local Cloudwatch, logGroup: /GreengrassSystem/connection_manager, error: RequestError: send request failed caused by: Post http://<path>/cloudwatch/logs/: dial tcp <address>: getsockopt: connection refused, response: { }.](#)
- [Error: Unable to create server due to: failed to load group: chmod /<greengrass-root>/ggc/deployment/lambda/arn:aws:lambda:<region>:<account-id>:function:<function-name>:<version>/<file-name>: no such file or directory.](#)

- [The AWS IoT Greengrass Core software doesn't start after you changed from running with no containerization to running in a Greengrass container.](#)
- [Error: Spool size should be at least 262144 bytes.](#)
- [Error: \[ERROR\]-Cloud messaging error: Error occurred while trying to publish a message. {"errorString": "operation timed out"}](#)
- [Error: container_linux.go:344: starting container process caused "process_linux.go:424: container init caused \"rootfs_linux.go:64: mounting \\\"/greengrass/ggc/socket/greengrass_ipc.sock\\\" to rootfs \\\"/greengrass/ggc/packages/<version>/rootfs/merged\\\" at \\\"/greengrass_ipc.sock\\\" caused \\\"stat /greengrass/ggc/socket/greengrass_ipc.sock: permission denied\\\"\".](#)
- [Error: Greengrass daemon running with PID: <process-id>. Some system components failed to start. Check 'runtime.log' for errors.](#)
- [Device shadow does not sync with the cloud.](#)
- [ERROR: unable to accept TCP connection. accept tcp \[::\]:8000: accept4: too many open files.](#)
- [Error: Runtime execution error: unable to start lambda container. container_linux.go:259: starting container process caused "process_linux.go:345: container init caused \"rootfs_linux.go:50: preparing rootfs caused \\\"permission denied\\\"\".](#)
- [Warning: \[WARN\]-\[5\]GK Remote: Error retrieving public key data: ErrPrincipalNotConfigured: private key for MqttCertificate is not set.](#)
- [Error: Permission denied when attempting to use role arn:aws:iam::<account-id>:role/<role-name> to access s3 url https://<region>-greengrass-updates.s3.<region>.amazonaws.com/core/<architecture>/greengrass-core-<distribution-version>.tar.gz.](#)
- [The AWS IoT Greengrass core is configured to use a network proxy and your Lambda function can't make outgoing connections.](#)
- [The core is in an infinite connect-disconnect loop. The runtime.log file contains a continuous series of connect and disconnect entries.](#)
- [Error: unable to start lambda container. container_linux.go:259: starting container process caused "process_linux.go:345: container init caused \"rootfs_linux.go:62: mounting \\\"proc\\\" to rootfs \\\"\"](#)
- [\[ERROR\]-runtime execution error: unable to start lambda container. {"errorString": "failed to initialize container mounts: failed to mask greengrass root in overlay upper dir: failed to create mask device at directory <ggc-path>: file exists"}](#)
- [\[ERROR\]-Deployment failed. {"deploymentId": "<deployment-id>", "errorString": "container test process with pid <pid> failed: container process state: exit status 1"}](#)

- [Error: \[ERROR\]-runtime execution error: unable to start lambda container. {"errorString": "failed to initialize container mounts: failed to create overlay fs for container: mounting overlay at /greengrass/ggc/packages/<ggc-version>/rootfs/merged failed: failed to mount with args source=\"/no_source\" dest=\"/greengrass/ggc/packages/<ggc-version>/rootfs/merged\" fstype=\"overlay\" flags=\"0\" data=\"lowerdir=/greengrass/ggc/packages/<ggc-version>/dns:/,upperdir=/greengrass/ggc/packages/<ggc-version>/rootfs/upper,workdir=/greengrass/ggc/packages/<ggc-version>/rootfs/work\": too many levels of symbolic links\"}](#)
- [Error: \[DEBUG\]-Failed to get routes. Discarding message.](#)
- [Error: \[Errno 24\] Too many open <lambda-function>,\[Errno 24\] Too many open files](#)
- [Error: ds server failed to start listening to socket: listen unix <ggc-path>/ggc/socket/greengrass_ipc.sock: bind: invalid argument](#)
- [\[INFO\] \(Copier\) aws.greengrass.StreamManager: stdout. Caused by: com.fasterxml.jackson.databind.JsonMappingException: Instant exceeds minimum or maximum instant](#)
- [GPG error: https://dnw9lb6lzp2d8.cloudfront.net stable InRelease: The following signatures were invalid: EXPKEYSIG 68D644ABD2327D47 AWS Greengrass Master Key](#)

Error: The configuration file is missing the CaPath, CertPath or KeyPath. The Greengrass daemon process with [pid = <pid>] died.

Solution: You might see this error in `crash.log` when the AWS IoT Greengrass Core software does not start. This can occur if you're running v1.6 or earlier. Do one of the following:

- Upgrade to v1.7 or later. We recommend that you always run the latest version of the AWS IoT Greengrass Core software. For download information, see [AWS IoT Greengrass Core Software](#).
- Use the correct `config.json` format for your AWS IoT Greengrass Core software version. For more information, see [the section called "AWS IoT Greengrass core configuration file"](#).

Note

To find which version of the AWS IoT Greengrass Core software is installed on the core device, run the following commands in your device terminal.

```
cd /greengrass-root/ggc/core/
```

```
sudo ./greengrassd --version
```

Error: Failed to parse /<greengrass-root>/config/config.json.

Solution: You might see this error when the AWS IoT Greengrass Core software does not start. Make sure the [Greengrass configuration file](#) is using valid JSON format.

Open `config.json` (located in `/greengrass-root/config`) and validate the JSON format. For example, make sure that commas are used correctly.

Error: Error occurred while generating TLS config: ErrUnknownURIScheme

Solution: You might see this error when the AWS IoT Greengrass Core software does not start. Make sure the properties in the [crypto](#) section of the Greengrass configuration file are valid. The error message should provide more information.

Open `config.json` (located in `/greengrass-root/config`) and check the `crypto` section. For example, certificate and key paths must use the correct URI format and point to the correct location.

Error: Runtime failed to start: unable to start workers: container test timed out.

Solution: You might see this error when the AWS IoT Greengrass Core software does not start. Set the `postStartHealthCheckTimeout` property in the [Greengrass configuration file](#). This optional property configures the amount of time (in milliseconds) that the Greengrass daemon waits for the post-start health check to finish. The default value is 30 seconds (30000 ms).

Open `config.json` (located in `/greengrass-root/config`). In the `runtime` object, add the `postStartHealthCheckTimeout` property and set the value to a number greater than 30000. Add a comma where needed to create a valid JSON document. For example:


```
...
"runtime" : {
  "cgroup" : {
    "useSystemd" : "yes"
  },
  "postStartHealthCheckTimeout" : 40000
},
...
```

Error: Failed to invoke PutLogEvents on local Cloudwatch, logGroup: /GreengrassSystem/connection_manager, error: RequestError: send request failed caused by: Post http://<path>/cloudwatch/logs/: dial tcp <address>: getsockopt: connection refused, response: { }.

Solution: You might see this error when the AWS IoT Greengrass Core software does not start. This can occur if you're running AWS IoT Greengrass on a Raspberry Pi and the required memory setup has not been completed. For more information, see [this step](#).

Error: Unable to create server due to: failed to load group: chmod /<greengrass-root>/ggc/deployment/lambda/arn:aws:lambda:<region>:<account-id>:function:<function-name>:<version>/<file-name>: no such file or directory.

Solution: You might see this error when the AWS IoT Greengrass Core software does not start. If you deployed a [Lambda executable](#) to the core, check the function's Handler property in the group.json file (located in */greengrass-root/ggc/deployment/group*). If the handler is not the exact name of your compiled executable, replace the contents of the group.json file with an empty JSON object (`{}`), and run the following commands to start AWS IoT Greengrass:

```
cd /greengrass/ggc/core/
sudo ./greengrassd start
```

Then, use the [AWS Lambda API](#) to update the function configuration's handler parameter, publish a new function version, and update the alias. For more information, see [AWS Lambda function versioning and aliases](#).

Assuming that you added the function to your Greengrass group by alias (recommended), you can now redeploy your group. (If not, you must point to the new function version or alias in your group definition and subscriptions before you deploy the group.)

The AWS IoT Greengrass Core software doesn't start after you changed from running with no containerization to running in a Greengrass container.

Solution: Check that you are not missing any container dependencies.

Error: Spool size should be at least 262144 bytes.

Solution: You might see this error when the AWS IoT Greengrass Core software does not start. Open the `group.json` file (located in `/greengrass-root/ggc/deployment/group`), replace the contents of the file with an empty JSON object (`{}`), and run the following commands to start AWS IoT Greengrass:

```
cd /greengrass/ggc/core/  
sudo ./greengrassd start
```

Then follow the steps in the [the section called "To cache messages in local storage"](#) procedure. For the `GGCloudSpooler` function, make sure to specify a `GG_CONFIG_MAX_SIZE_BYTES` value that's greater than or equal to 262144.

Error: [ERROR]-Cloud messaging error: Error occurred while trying to publish a message. {"errorString": "operation timed out"}

Solution: You might see this error in `GGCloudSpooler.log` when the Greengrass core is unable to send MQTT messages to AWS IoT Core. This can occur if the core environment has limited bandwidth and high latency. If you're running AWS IoT Greengrass v1.10.2 or later, try increasing

the `mqttOperationTimeout` value in the [config.json](#) file. If the property is not present, add it to the `coreThing` object. For example:

```
{
  "coreThing": {
    "mqttOperationTimeout": 10,
    "caPath": "root-ca.pem",
    "certPath": "hash.cert.pem",
    "keyPath": "hash.private.key",
    ...
  },
  ...
}
```

The default value is 5 and the minimum value is 5.

Error: container_linux.go:344: starting container process caused "process_linux.go:424: container init caused \"rootfs_linux.go:64: mounting \"\\\"/greengrass/ggc/socket/greengrass_ipc.sock\\\" to rootfs \"\\\"/greengrass/ggc/packages/<version>/rootfs/merged\\\" at \"\\\"/greengrass_ipc.sock\\\" caused \"\\\"stat /greengrass/ggc/socket/greengrass_ipc.sock: permission denied\\\"\"\".

Solution: You might see this error in `runtime.log` when the AWS IoT Greengrass Core software does not start. This occurs if your `umask` is higher than `0022`. To resolve this issue, you must set the `umask` to `0022` or lower. A value of `0022` grants everyone read permission to new files by default.

Error: Greengrass daemon running with PID: <process-id>. Some system components failed to start. Check 'runtime.log' for errors.

Solution: You might see this error when the AWS IoT Greengrass Core software does not start. Check `runtime.log` and `crash.log` for specific error information. For more information, see [the section called "Troubleshooting with logs"](#).

Device shadow does not sync with the cloud.

Solution: Make sure that AWS IoT Greengrass has permissions for `iot:UpdateThingShadow` and `iot:GetThingShadow` actions in the [Greengrass service role](#). If the service role uses the `AWSGreengrassResourceAccessRolePolicy` managed policy, these permissions are included by default.

See [Troubleshooting shadow synchronization timeout issues](#).

ERROR: unable to accept TCP connection. accept tcp [::]:8000: accept4: too many open files.

Solution: You might see this error in the `greengrassd` script output. This can occur if the file descriptor limit for the AWS IoT Greengrass Core software has reached the threshold and must be increased.

Use the following command and then restart the AWS IoT Greengrass Core software.

```
ulimit -n 2048
```

Note

In this example, the limit is increased to 2048. Choose a value appropriate for your use case.

Error: Runtime execution error: unable to start lambda container. container_linux.go:259: starting container process caused "process_linux.go:345: container init caused \"rootfs_linux.go:50: preparing rootfs caused \\\"permission denied\\\"\"\".

Solution: Either install AWS IoT Greengrass directly under the root directory, or make sure that the directory where the AWS IoT Greengrass Core software is installed and its parent directories have execute permissions for everyone.

Warning: [WARN]-[5]GK Remote: Error retrieving public key data: ErrPrincipalNotConfigured: private key for MqttCertificate is not set.

Solution: AWS IoT Greengrass uses a common handler to validate the properties of all security principals. This warning in `runtime.log` is expected unless you specified a custom private key for the local MQTT server. For more information, see [the section called "Security principals"](#).

Error: Permission denied when attempting to use role `arn:aws:iam::<account-id>:role/<role-name>` to access s3 url `https://<region>-greengrass-updates.s3.<region>.amazonaws.com/core/<architecture>/greengrass-core-<distribution-version>.tar.gz`.

Solution: You might see this error when an over-the-air (OTA) update fails. In the signer role policy, add the target AWS Region as a Resource. This signer role is used to presign the S3 URL for the AWS IoT Greengrass software update. For more information, see [S3 URL signer role](#).

The AWS IoT Greengrass core is configured to use a [network proxy](#) and your Lambda function can't make outgoing connections.

Solution: Depending on your runtime and the executables used by the Lambda function to create connections, you might also receive connection timeout errors. Make sure your Lambda functions use the appropriate proxy configuration to connect through the network proxy. AWS IoT Greengrass passes the proxy configuration to user-defined Lambda functions through the `http_proxy`, `https_proxy`, and `no_proxy` environment variables. They can be accessed as shown in the following Python snippet.

```
import os
print(os.environ['http_proxy'])
```

Use the same case as the variable defined in your environment, for example, all lower case `http_proxy` or all upper case `HTTP_PROXY`. For these variables, AWS IoT Greengrass supports both.

Note

Most common libraries used to make connections (such as boto3 or cURL and python requests packages) use these environment variables by default.

The core is in an infinite connect-disconnect loop. The runtime.log file contains a continuous series of connect and disconnect entries.

Solution: This can happen when another device is hard-coded to use the core thing name as the client ID for MQTT connections to AWS IoT. Simultaneous connections in the same AWS Region and AWS account must use unique client IDs. By default, the core uses the core thing name as the client ID for these connections.

To resolve this issue, you can change the client ID used by the other device for the connection (recommended) or override the default value for the core.

To override the default client ID for the core device

1. Run the following command to stop the Greengrass daemon:

```
cd /greengrass-root/ggc/core/  
sudo ./greengrassd stop
```

2. Open `greengrass-root/config/config.json` for editing as the su user.
3. In the `coreThing` object, add the `coreClientId` property, and set the value to your custom client ID. The value must be between 1 and 128 characters. It must be unique in the current AWS Region for the AWS account.

```
"coreClientId": "MyCustomClientId"
```

4. Start the daemon.

```
cd /greengrass-root/ggc/core/  
sudo ./greengrassd start
```

Error: unable to start lambda container. container_linux.go:259: starting container process caused "process_linux.go:345: container init caused \"rootfs_linux.go:62: mounting \"\\\"proc\\\"\" to rootfs \"\\\""

Solution: On some platforms, you might see this error in `runtime.log` when AWS IoT Greengrass tries to mount the `/proc` file system to create a Lambda container. Or, you might see similar errors, such as `operation not permitted` or `EPERM`. These errors can occur even if tests run on the platform by the dependency checker script pass.

Try one of the following possible solutions:

- Enable the `CONFIG_DEVPTS_MULTIPLE_INSTANCES` option in the Linux kernel.
- Set the `/proc` mount options on the host to `rw,relatime` only.
- Upgrade the Linux kernel to 4.9 or later.

Note

This issue is not related to mounting `/proc` for local resource access.

[ERROR]-runtime execution error: unable to start lambda container. {"errorString": "failed to initialize container mounts: failed to mask greengrass root in overlay upper dir: failed to create mask device at directory <ggc-path>: file exists"}

Solution: You might see this error in `runtime.log` when the deployment fails. This error occurs if a Lambda function in the AWS IoT Greengrass group cannot access the `/usr` directory in the core's file system.

To resolve this issue, add a local volume resource to the group and then deploy the group. This resource must:

- Specify `/usr` as the **Source path** and **Destination path**.
- Automatically add OS group permissions of the Linux group that owns the resource.
- Be affiliated with the Lambda function and allow read-only access.

[ERROR]-Deployment failed. {"deploymentId": "<deployment-id>", "errorString": "container test process with pid <pid> failed: container process state: exit status 1"}

Solution: You might see this error in `runtime.log` when the deployment fails. This error occurs if a Lambda function in the AWS IoT Greengrass group cannot access the `/usr` directory in the core's file system.

You can confirm that this is the case by checking `GGCanary.log` for additional errors. If the Lambda function cannot access the `/usr` directory, `GGCanary.log` will contain the following error:

```
[ERROR]-standard_init_linux.go:207: exec user process caused "no such file or directory"
```

To resolve this issue, add a local volume resource to the group and then deploy the group. This resource must:

- Specify `/usr` as the **Source path** and **Destination path**.
- Automatically add OS group permissions of the Linux group that owns the resource.
- Be affiliated with the Lambda function and allow read-only access.

Error: [ERROR]-runtime execution error: unable to start lambda container. {"errorString": "failed to initialize container mounts: failed to create overlay fs for container: mounting overlay at /greengrass/ggc/packages/<ggc-version>/rootfs/merged failed: failed to mount with args source=\"no_source\" dest=\"/greengrass/ggc/packages/<ggc-version>/rootfs/merged\" fstype=\"overlay\" flags=\"0\" data=\"lowerdir=/greengrass/ggc/packages/<ggc-version>/dns:/,upperdir=/greengrass/ggc/packages/<ggc-version>/rootfs/upper,workdir=/greengrass/ggc/packages/<ggc-version>/rootfs/work\": too many levels of symbolic links"}

Solution: You might see this error in the `runtime.log` file when the AWS IoT Greengrass Core software doesn't start. This issue might be more common on Debian operating systems.

To resolve this issue, do the following:

1. Upgrade the AWS IoT Greengrass Core software to v1.9.3 or later. This should automatically resolve this issue.
2. If you still get this error after you upgrade the AWS IoT Greengrass Core software, set the `system.useOverlayWithTmpfs` property to `true` in the [config.json](#) file.

Example Example

```
{
  "system": {
    "useOverlayWithTmpfs": true
  },
  "coreThing": {
    "caPath": "root-ca.pem",
    "certPath": "cloud.pem.crt",
    "keyPath": "cloud.pem.key",
    ...
  },
  ...
}
```

Note

Your AWS IoT Greengrass Core software version is shown in the error message. To find your Linux kernel version, run `uname -r`.

Error: [DEBUG]-Failed to get routes. Discarding message.

Solution: Check the subscriptions in your group and make sure that the subscription listed in the [DEBUG] message exists.

Error: [Errno 24] Too many open <lambda-function>,[Errno 24] Too many open files

Solution: You might see this error in your Lambda function log file if the function instantiates `StreamManagerClient` in the function handler. We recommend that you create the client outside the handler. For more information, see [the section called “Use StreamManagerClient to work with streams”](#).

Error: ds server failed to start listening to socket: listen unix <ggc-path>/ggc/socket/greengrass_ipc.sock: bind: invalid argument

Solution: You might see this error when the AWS IoT Greengrass Core software doesn't start. This error occurs when the AWS IoT Greengrass Core software is installed to a folder with a long file path. Reinstall the AWS IoT Greengrass Core software to a folder with a file path that has fewer than 79 bytes, if you don't use a [write directory](#), or 83 bytes, if you do use a write directory.

[INFO] (Copier) aws.greengrass.StreamManager: stdout. Caused by: com.fasterxml.jackson.databind.JsonMappingException: Instant exceeds minimum or maximum instant

When you upgrade AWS IoT Greengrass core software to v1.11.3, you might see the following error in the stream manager logs if stream manager fails to start.

```
2021-07-16T00:54:58.568Z [INFO] (Copier) aws.greengrass.StreamManager:
stdout. Caused by: com.fasterxml.jackson.databind.JsonMappingException:
Instant exceeds minimum or maximum instant (through reference chain:
com.amazonaws.iot.greengrass.streammanager.export.PersistedSuccessExportStatesV1["lastExportTime"])
{scriptName=services.aws.greengrass.StreamManager.lifecycle.startup.script,
serviceName=aws.greengrass.StreamManager, currentState=STARTING}
2021-07-16T00:54:58.579Z [INFO] (Copier) aws.greengrass.StreamManager: stdout.
Caused by: java.time.DateTimeException: Instant exceeds minimum or maximum instant.
{scriptName=services.aws.greengrass.StreamManager.lifecycle.startup.script,
serviceName=aws.greengrass.StreamManager, currentState=STARTING}
```

If you're using a version of AWS IoT Greengrass core software older than v1.11.3, and you want to upgrade to a later version, use an OTA update to upgrade to v1.11.4.

GPG error: <https://dnw9lb6lzp2d8.cloudfront.net> stable InRelease: The following signatures were invalid: EXPKEYSIG 68D644ABD2327D47 AWS Greengrass Master Key

When you run `apt update` on a device where you [installed the AWS IoT Greengrass core software from an APT repository](#), you might see the following error.

```
Err:4 https://dnw9lb6lzp2d8.cloudfront.net stable InRelease
  The following signatures were invalid: EXPKEYSIG 68D644ABD2327D47 AWS Greengrass
  Master Key
Reading package lists... Done
W: GPG error: https://dnw9lb6lzp2d8.cloudfront.net stable InRelease: The following
  signatures were invalid: EXPKEYSIG 68D644ABD2327D47 AWS Greengrass Master Key
```

This error occurs because AWS IoT Greengrass no longer offers the option to install or update the AWS IoT Greengrass core software from the APT repository. To successfully run `apt update`, remove the AWS IoT Greengrass repository from the device's sources list.

```
sudo rm /etc/apt/sources.list.d/greengrass.list
sudo apt update
```

Deployment issues

Use the following information to help troubleshoot deployment issues.

Issues

- [Your current deployment does not work and you want to revert to a previous working deployment.](#)
- [You see a 403 Forbidden error on deployment in the logs.](#)
- [A ConcurrentDeployment error occurs when you run the create-deployment command for the first time.](#)
- [Error: Greengrass is not authorized to assume the Service Role associated with this account, or the error: Failed: TES service role is not associated with this account.](#)
- [Error: unable to execute download step in deployment. error while downloading: error while downloading the Group definition file: ... x509: certificate has expired or is not yet valid](#)
- [The deployment doesn't finish.](#)
- [Error: Unable to find java or java8 executables, or the error: Deployment <deployment-id> of type NewDeployment for group <group-id> failed error: worker with <worker-id> failed to initialize with reason Installed Java version must be greater than or equal to 8](#)
- [The deployment doesn't finish, and runtime.log contains multiple "wait 1s for container to stop" entries.](#)
- [The deployment doesn't finish, and runtime.log contains "\[ERROR\]-Greengrass deployment error: failed to report deployment status back to cloud {"deploymentId": "<deployment-id>", "errorString": "Failed to initiate PUT, endpoint: https://<deployment-status>, error: Put https://<deployment-status>: proxyconnect tcp: x509: certificate signed by unknown authority"}"](#)
- [Error: Deployment <deployment-id> of type NewDeployment for group <group-id> failed error: Error while processing. group config is invalid: 112 or \[119 0\] don't have rw permission on the file: <path>.](#)
- [Error: <list-of-function-arns> are configured to run as root but Greengrass is not configured to run Lambda functions with root permissions.](#)
- [Error: Deployment <deployment-id> of type NewDeployment for group <group-id> failed error: Greengrass deployment error: unable to execute download step in deployment. error while processing: unable to load the group file downloaded: could not find UID based on user name, userName: ggc_user: user: unknown user ggc_user.](#)
- [Error: \[ERROR\]-runtime execution error: unable to start lambda container. {"errorString": "failed to initialize container mounts: failed to mask greengrass root in overlay upper dir: failed to create mask device at directory <ggc-path>: file exists"}](#)

- [Error: Deployment <deployment-id> of type NewDeployment for group <group-id> failed error: process start failed: container_linux.go:259: starting container process caused "process_linux.go:250: running exec setns process for init caused \"wait: no child processes\"".](#)
- [Error: \[WARN\]-MQTT\[client\] dial tcp: lookup <host-prefix>-ats.iot.<region>.amazonaws.com: no such host ... \[ERROR\]-Greengrass deployment error: failed to report deployment status back to cloud ... net/http: request canceled while waiting for connection \(Client.Timeout exceeded while awaiting headers\)](#)

Your current deployment does not work and you want to revert to a previous working deployment.

Solution: Use the AWS IoT console or AWS IoT Greengrass API to redeploy a previous working deployment. This deploys the corresponding group version to your core device.

To redeploy a deployment (console)

1. On the group configuration page, choose the **Deployments** tab. This page displays the deployment history for the group, including the date and time, group version, and status of each deployment attempt.
2. Find the row that contains the deployment you want to redeploy. Select the deployment you want to redeploy and choose **Redeploy**.

Deployments	Group history overview			By deployment ▾
Subscriptions	Deployed	Version	Status	
Cores	Jul 1, 2019 1:56:49 PM -0700	8dd1d899-4ac9-4f5d-afe4-22de086efc62	● Successfully complet...	...
Devices	Jul 1, 2019 1:41:47 PM -0700	4ad66e5d-3808-446b-940a-b1a788898382	● Successfully complet...	...
Lambdas	Jun 18, 2019 8:16:02 AM -0700	1f3870b6-850e-4c97-8018-c872e17b235b	● Failed	...
Resources				
Connectors				

To redeploy a deployment (CLI)

1. Use [ListDeployments](#) to find the ID of the deployment you want to redeploy. For example:

```
aws greengrass list-deployments --group-id 74d0b623-c2f2-4cad-9acc-ef92f61fcaf7
```

The command returns the list of deployments for the group.

```
{
  "Deployments": [
    {
      "DeploymentId": "8d179428-f617-4a77-8a0c-3d61fb8446a6",
      "DeploymentType": "NewDeployment",
      "GroupArn": "arn:aws:greengrass:us-west-2:123456789012:/greengrass/
groups/74d0b623-c2f2-4cad-9acc-ef92f61fcaf7/versions/8dd1d899-4ac9-4f5d-
afe4-22de086efc62",
      "CreatedAt": "2019-07-01T20:56:49.641Z"
    },
    {
      "DeploymentId": "f8e4c455-8ac4-453a-8252-512dc3e9c596",
      "DeploymentType": "NewDeployment",
      "GroupArn": "arn:aws:greengrass:us-west-2:123456789012:/greengrass/
groups/74d0b623-c2f2-4cad-9acc-ef92f61fcaf7/versions/4ad66e5d-3808-446b-940a-
b1a788898382",
      "CreatedAt": "2019-07-01T20:41:47.048Z"
    },
    {
      "DeploymentId": "e4aca044-bbd8-41b4-b697-930ca7c40f3e",
      "DeploymentType": "NewDeployment",
      "GroupArn": "arn:aws:greengrass:us-west-2:123456789012:/greengrass/
groups/74d0b623-c2f2-4cad-9acc-ef92f61fcaf7/versions/1f3870b6-850e-4c97-8018-
c872e17b235b",
      "CreatedAt": "2019-06-18T15:16:02.965Z"
    }
  ]
}
```

Note

These AWS CLI commands use example values for the group and deployment ID. When you run the commands, make sure to replace the example values.

2. Use [CreateDeployment](#) to redeploy the target deployment. Set the deployment type to Redeployment. For example:

```
aws greengrass create-deployment --deployment-type Redeployment \  
  --group-id 74d0b623-c2f2-4cad-9acc-ef92f61fcdf7 \  
  --deployment-id f8e4c455-8ac4-453a-8252-512dc3e9c596
```

The command returns the ARN and ID of the new deployment.

```
{  
  "DeploymentId": "f9ed02b7-c28e-4df6-83b1-e9553ddd0fc2",  
  "DeploymentArn": "arn:aws:greengrass:us-west-2::123456789012:/greengrass/  
groups/74d0b623-c2f2-4cad-9acc-ef92f61fcdf7/deployments/f9ed02b7-c28e-4df6-83b1-  
e9553ddd0fc2"  
}
```

3. Use [GetDeploymentStatus](#) to get the status of the deployment.

You see a 403 Forbidden error on deployment in the logs.

Solution: Make sure the policy of the AWS IoT Greengrass core in the cloud includes "greengrass:*" as an allowed action.

A ConcurrentDeployment error occurs when you run the create-deployment command for the first time.

Solution: A deployment might be in progress. You can run [get-deployment-status](#) to see if a deployment was created. If not, try creating the deployment again.

Error: Greengrass is not authorized to assume the Service Role associated with this account, or the error: Failed: TES service role is not associated with this account.

Solution: You might see this error when the deployment fails. Check that a Greengrass service role is associated with your AWS account in the current AWS Region. For more information, see

[the section called “Manage the service role \(CLI\)”](#) or [the section called “Manage the service role \(console\)”](#).

Error: unable to execute download step in deployment. error while downloading: error while downloading the Group definition file: ... x509: certificate has expired or is not yet valid

Solution: You might see this error in `runtime.log` when the deployment fails. If you receive a `Deployment failed` error that contains the message `x509: certificate has expired or is not yet valid`, check the device clock. TLS and X.509 certificates provide a secure foundation for building IoT systems, but they require accurate times on servers and clients. IoT devices should have the correct time (within 15 minutes) before they attempt to connect to AWS IoT Greengrass or other TLS services that use server certificates. For more information, see [Using Device Time to Validate AWS IoT Server Certificates](#) on *The Internet of Things on AWS Official Blog*.

The deployment doesn't finish.

Solution: Do the following:

- Make sure that the AWS IoT Greengrass daemon is running on your core device. In your core device terminal, run the following commands to check whether the daemon is running and start it, if needed.

1. To check whether the daemon is running:

```
ps aux | grep -E 'greengrass.*daemon'
```

If the output contains a root entry for `/greengrass/ggc/packages/1.11.6/bin/daemon`, then the daemon is running.

The version in the path depends on the AWS IoT Greengrass Core software version that's installed on your core device.

2. To start the daemon:

```
cd /greengrass/ggc/core/
```



```
sudo ./greengrassd start
```

- Make sure that your core device is connected and the core connection endpoints are configured properly.

Error: Unable to find java or java8 executables, or the error: Deployment <deployment-id> of type NewDeployment for group <group-id> failed error: worker with <worker-id> failed to initialize with reason Installed Java version must be greater than or equal to 8

Solution: If stream manager is enabled for the AWS IoT Greengrass core, you must install the Java 8 runtime on the core device before you deploy the group. For more information, see the [requirements](#) for stream manager. Stream manager is enabled by default when you use the **Default Group creation** workflow in the AWS IoT console to create a group.

Or, disable stream manager and then deploy the group. For more information, see [the section called "Configure settings \(console\)"](#).

The deployment doesn't finish, and runtime.log contains multiple "wait 1s for container to stop" entries.

Solution: Run the following commands in your core device terminal to restart the AWS IoT Greengrass daemon.

```
cd /greengrass/ggc/core/  
sudo ./greengrassd stop  
sudo ./greengrassd start
```

The deployment doesn't finish, and `runtime.log` contains "[ERROR]-Greengrass deployment error: failed to report deployment status back to cloud {"deploymentId": "<deployment-id>", "errorString": "Failed to initiate PUT, endpoint: https://<deployment-status>, error: Put https://<deployment-status>: proxyconnect tcp: x509: certificate signed by unknown authority"}"

Solution: You might see this error in `runtime.log` when the Greengrass core is configured to use an HTTPS proxy connection and the proxy server certificate chain isn't trusted on the system. To try to resolve this issue, add the certificate chain to the root CA certificate. The Greengrass core adds the certificates from this file to the certificate pool used for TLS authentication in HTTPS and MQTT connections with AWS IoT Greengrass.

The following example shows a proxy server CA certificate added to the root CA certificate file:

```
# My proxy CA
-----BEGIN CERTIFICATE-----
MIIEFTCCA v2gAwIQWgIVAMHSAzWG/5YVRYtRQ0xXUTEpHuEmApzGCSqGSIb3DQEK
\nCwUAhUL9MQswCQwJVUzEPMAVUzEYMBYGA1UECgwP1hem9uLmNvbSBjbMUMRww
... content of proxy CA certificate ...
+vHIRlt0e5JAm5\noTIZGoFbK82A0/n07f/t5PSIDAim9V3Gc3pSXxCCAQoFYnui
GaPUlGk1gCE84a0X\n7Rp/1ND/PuMZ/s8Yj1kY2NmYmNjMCAXDTE5MTEyN2cM216
gJMIADggEPADf2/m45hzEXAMPLE=
-----END CERTIFICATE-----

# Amazon Root CA 1
-----BEGIN CERTIFICATE-----
MIIDQTCCAimgF6AwIBAgITBmyfz/5mjAo54vB4ikPm1jZKyjANJmApzyMZFo6qBg
ADA5MQswCQYDVQQGEwJVUzEPMA0tMVT8QtPHRh8jrdkGA1UEChMGRGV33QDExBBKw
... content of root CA certificate ...
o/ufQJQWUCyziar1hem9uMRkwFwYVPSHCb2XV4cdFyQzR1KldZwgJcIQ6XUDgHaa
5MsI+yMRQ+hDaXJioblXgUka642M4UwtBV8oK2xJNDd2ZhwLnoQdeXeGADKkpy
rqXRfKoQnoZsG4q5WTP46EXAMPLE
-----END CERTIFICATE-----
```

By default, the root CA certificate file is located in `/greengrass-root/certs/root.ca.pem`. To find the location on your core device, check the `crypto.caPath` property in [config.json](#).

Note

greengrass-root represents the path where the AWS IoT Greengrass Core software is installed on your device. Typically, this is the /greengrass directory.

Error: Deployment <deployment-id> of type NewDeployment for group <group-id> failed error: Error while processing. group config is invalid: 112 or [119 0] don't have rw permission on the file: <path>.

Solution: Make sure that the owner group of the *<path>* directory has read and write permissions to the directory.

Error: <list-of-function-arns> are configured to run as root but Greengrass is not configured to run Lambda functions with root permissions.

Solution: You might see this error in `runtime.log` when the deployment fails. Make sure that you have configured AWS IoT Greengrass to allow Lambda functions to run with root permissions. Either change the value of `allowFunctionsToRunAsRoot` in `greengrass_root/config/config.json` to `yes` or change the Lambda function to run as another user/group. For more information, see [the section called "Running a Lambda function as root"](#).

Error: Deployment <deployment-id> of type NewDeployment for group <group-id> failed error: Greengrass deployment error: unable to execute download step in deployment. error while processing: unable to load the group file downloaded: could not find UID based on user name, userName: ggc_user: user: unknown user ggc_user.

Solution: If the [default access identity](#) of the AWS IoT Greengrass group uses the standard system accounts, the `ggc_user` user and `ggc_group` group must be present on the device. For

instructions that show how to add the user and group, see this [step](#). Make sure to enter the names exactly as shown.

Error: [ERROR]-runtime execution error: unable to start lambda container. {"errorString": "failed to initialize container mounts: failed to mask greengrass root in overlay upper dir: failed to create mask device at directory <ggc-path>: file exists"}

Solution: You might see this error in `runtime.log` when the deployment fails. This error occurs if a Lambda function in the Greengrass group can't access the `/usr` directory in the core's file system. To resolve this issue, add a [local volume resource](#) to the group and then deploy the group. The resource must:

- Specify `/usr` as the **Source path** and **Destination path**.
- Automatically add OS group permissions of the Linux group that owns the resource.
- Be affiliated with the Lambda function and allow read-only access.

Error: Deployment <deployment-id> of type NewDeployment for group <group-id> failed error: process start failed: container_linux.go:259: starting container process caused "process_linux.go:250: running exec setns process for init caused \"wait: no child processes\"".

Solution: You might see this error when the deployment fails. Retry the deployment.

Error: [WARN]-MQTT[client] dial tcp: lookup <host-prefix>-ats.iot.<region>.amazonaws.com: no such host ... [ERROR]-Greengrass deployment error: failed to report deployment status back to cloud ... net/http: request canceled while waiting for connection (Client.Timeout exceeded while awaiting headers)

Solution: You might see this error if you're using `systemd-resolved`, which enables the DNSSEC setting by default. As a result, many public domains are not recognized. Attempts to reach the AWS IoT Greengrass endpoint fail to find the host, so your deployments remain in the `In Progress` state.

You can use the following commands and output to test for this issue. Replace the *region* placeholder in the endpoints with your AWS Region.

```
$ ping greengrass-ats.iot.region.amazonaws.com
ping: greengrass-ats.iot.region.amazonaws.com: Name or service not known
```

```
$ systemd-resolve greengrass-ats.iot.region.amazonaws.com
greengrass-ats.iot.region.amazonaws.com: resolve call failed: DNSSEC validation failed:
failed-auxiliary
```

One possible solution is to disable DNSSEC. When DNSSEC is false, DNS lookups are not DNSSEC validated. For more information, see this [known issue](#) for `systemd`.

1. Add `DNSSEC=false` to `/etc/systemd/resolved.conf`.
2. Restart `systemd-resolved`.

For information about `resolved.conf` and DNSSEC, run `man resolved.conf` in your terminal.

Create group and create function issues

Use the following information to help troubleshoot issues with creating an AWS IoT Greengrass group or Greengrass Lambda function.

Issues

- [Error: Your 'IsolationMode' configuration for the group is invalid.](#)
- [Error: Your 'IsolationMode' configuration for function with arn <function-arn> is invalid.](#)
- [Error: MemorySize configuration for function with arn <function-arn> is not allowed in IsolationMode=NoContainer.](#)
- [Error: Access Sysfs configuration for function with arn <function-arn> is not allowed in IsolationMode=NoContainer.](#)
- [Error: MemorySize configuration for function with arn <function-arn> is required in IsolationMode=GreengrassContainer.](#)
- [Error: Function <function-arn> refers to resource of type <resource-type> that is not allowed in IsolationMode=NoContainer.](#)
- [Error: Execution configuration for function with arn <function-arn> is not allowed.](#)

Error: Your 'IsolationMode' configuration for the group is invalid.

Solution: This error occurs when the IsolationMode value in the DefaultConfig of function-definition-version is not supported. Supported values are GreengrassContainer and NoContainer.

Error: Your 'IsolationMode' configuration for function with arn <function-arn> is invalid.

Solution: This error occurs when the IsolationMode value in the <function-arn> of the function-definition-version is not supported. Supported values are GreengrassContainer and NoContainer.

Error: MemorySize configuration for function with arn <function-arn> is not allowed in IsolationMode=NoContainer.

Solution: This error occurs when you specify a MemorySize value and you choose to run without containerization. Lambda functions that are run without containerization cannot have memory

limits. You can either remove the limit or you can change the Lambda function to run in an AWS IoT Greengrass container.

Error: Access Sysfs configuration for function with arn <function-arn> is not allowed in IsolationMode=NoContainer.

Solution: This error occurs when you specify `true` for `AccessSysfs` and you choose to run without containerization. Lambda functions run without containerization must have their code updated to access the file system directly and cannot use `AccessSysfs`. You can either specify a value of `false` for `AccessSysfs` or you can change the Lambda function to run in an AWS IoT Greengrass container.

Error: MemorySize configuration for function with arn <function-arn> is required in IsolationMode=GreengrassContainer.

Solution: This error occurs because you did not specify a `MemorySize` limit for a Lambda function that you are running in an AWS IoT Greengrass container. Specify a `MemorySize` value to resolve the error.

Error: Function <function-arn> refers to resource of type <resource-type> that is not allowed in IsolationMode=NoContainer.

Solution: You cannot access `Local.Device`, `Local.Volume`, `ML_Model.SageMaker.Job`, `ML_Model.S3_Object`, or `S3_Object.Generic_Archive` resource types when you run a Lambda function without containerization. If you need those resource types, you must run in an AWS IoT Greengrass container. You can also access local devices directly when running without containerization by changing the code in your Lambda function.

Error: Execution configuration for function with arn <function-arn> is not allowed.

Solution: This error occurs when you create a system Lambda function with GGIPDetector or GGCloudSpooler and you specified IsolationMode or RunAs configuration. You must omit the Execution parameters for this system Lambda function.

Discovery issues

Use the following information to help troubleshoot issues with the AWS IoT Greengrass Discovery service.

Issues

- [Error: Device is a member of too many groups, devices may not be in more than 10 groups](#)

Error: Device is a member of too many groups, devices may not be in more than 10 groups

Solution: This is a known limitation. A [client device](#) can be a member of up to 10 groups.

Machine learning resource issues

Use the following information to help troubleshoot issues with machine learning resources.

Issues

- [InvalidMLModelOwner - GroupOwnerSetting is provided in ML model resource, but GroupOwner or GroupPermission is not present](#)
- [NoContainer function cannot configure permission when attaching Machine Learning resources. <function-arn> refers to Machine Learnin resource <resource-id> with permission <ro/rw> in resource access policy.](#)

- [Function <function-arn> refers to Machine Learning resource <resource-id> with missing permission in both ResourceAccessPolicy and resource OwnerSetting.](#)
- [Function <function-arn> refers to Machine Learning resource <resource-id> with permission \"rw\", while resource owner setting GroupPermission only allows \"ro\".](#)
- [NoContainer Function <function-arn> refers to resources of nested destination path.](#)
- [Lambda <function-arn> gains access to resource <resource-id> by sharing the same group owner id](#)

InvalidMLModelOwner - GroupOwnerSetting is provided in ML model resource, but GroupOwner or GroupPermission is not present

Solution: You receive this error if a machine learning resource contains the [ResourceDownloadOwnerSetting](#) object but the required GroupOwner or GroupPermission property isn't defined. To resolve this issue, define the missing property.

NoContainer function cannot configure permission when attaching Machine Learning resources. <function-arn> refers to Machine Learning resource <resource-id> with permission <ro/rw> in resource access policy.

Solution: You receive this error if a non-containerized Lambda function specifies function-level permissions to a machine learning resource. Non-containerized functions must inherit permissions from the resource owner permissions defined on the machine learning resource. To resolve this issue, choose to [inherit resource owner permissions](#) (console) or [remove the permissions from the Lambda function's resource access policy](#) (API).

Function <function-arn> refers to Machine Learning resource <resource-id> with missing permission in both ResourceAccessPolicy and resource OwnerSetting.

Solution: You receive this error if permissions to the machine learning resource aren't configured for the attached Lambda function or the resource. To resolve this issue, configure permissions in the [ResourceAccessPolicy](#) property for the Lambda function or the [OwnerSetting](#) property for the resource.

Function <function-arn> refers to Machine Learning resource <resource-id> with permission \"rw\", while resource owner setting GroupPermission only allows \"ro\".

Solution: You receive this error if the access permissions defined for the attached Lambda function exceed the resource owner permissions defined for the machine learning resource. To resolve this issue, set more restrictive permissions for the Lambda function or less restrictive permissions for the resource owner.

NoContainer Function <function-arn> refers to resources of nested destination path.

Solution: You receive this error if multiple machine learning resources attached to a non-containerized Lambda function use the same destination path or a nested destination path. To resolve this issue, specify separate destination paths for the resources.

Lambda <function-arn> gains access to resource <resource-id> by sharing the same group owner id

Solution: You receive this error in `runtime.log` if the same OS group is specified as the Lambda function's [Run as](#) identity and the [resource owner](#) for a machine learning resource, but the resource is not attached to the Lambda function. This configuration gives the Lambda function implicit permissions that it can use to access the resource without AWS IoT Greengrass authorization.

To resolve this issue, use a different OS group for one of the properties or attach the machine learning resource to the Lambda function.

AWS IoT Greengrass core in Docker issues

Use the following information to help troubleshoot issues with running an AWS IoT Greengrass core in a Docker container.

Issues

- [Error: Unknown options: -no-include-email.](#)
- [Warning: IPv4 is disabled. Networking will not work.](#)
- [Error: A firewall is blocking file Sharing between windows and the containers.](#)
- [Error: An error occurred \(AccessDeniedException\) when calling the GetAuthorizationToken operation: User: arn:aws:iam::<account-id>:user/<user-name> is not authorized to perform: ecr:GetAuthorizationToken on resource: *](#)
- [Error: Cannot create container for the service greengrass: Conflict. The container name "/aws-iot-greengrass" is already in use.](#)
- [Error: \[FATAL\]-Failed to reset thread's mount namespace due to an unexpected error: "operation not permitted". To maintain consistency, GGC will crash and need to be manually restarted.](#)

Error: Unknown options: -no-include-email.

Solution: This error can occur when you run the `aws ecr get-login` command. Make sure that you have the latest AWS CLI version installed (for example, run: `pip install awscli --upgrade --user`). If you're using Windows and you installed the CLI using the MSI installer, you must repeat the installation process. For more information, see [Installing the AWS Command Line Interface on Microsoft Windows](#) in the *AWS Command Line Interface User Guide*.

Warning: IPv4 is disabled. Networking will not work.

Solution: You might receive this warning or a similar message when running AWS IoT Greengrass on a Linux computer. Enable IPv4 network forwarding as described in this [step](#). AWS IoT

Greengrass cloud deployment and MQTT communications don't work when IPv4 forwarding isn't enabled. For more information, see [Configure namespaced kernel parameters \(sysctls\) at runtime](#) in the Docker documentation.

Error: A firewall is blocking file Sharing between windows and the containers.

Solution: You might receive this error or a `Firewall Detected` message when running Docker on a Windows computer. This can also occur if you are signed in on a virtual private network (VPN) and your network settings are preventing the shared drive from being mounted. In that situation, turn off VPN and re-run the Docker container.

Error: An error occurred (AccessDeniedException) when calling the GetAuthorizationToken operation: User: arn:aws:iam::<account-id>:user/<user-name> is not authorized to perform: ecr:GetAuthorizationToken on resource: *

You might receive this error when running the `aws ecr get-login-password` command if you don't have sufficient permissions to access an Amazon ECR repository. For more information, see [Amazon ECR Repository Policy Examples](#) and [Accessing One Amazon ECR Repository](#) in the *Amazon ECR User Guide*.

Error: Cannot create container for the service greengrass: Conflict. The container name "/aws-iot-greengrass" is already in use.

Solution: This can occur when the container name is used by an older container. To resolve this issue, run the following command to remove the old Docker container:

```
docker rm -f $(docker ps -a -q -f "name=aws-iot-greengrass")
```

Error: [FATAL]-Failed to reset thread's mount namespace due to an unexpected error: "operation not permitted". To maintain consistency, GGC will crash and need to be manually restarted.

Solution: This error in `runtime.log` can occur when you try to deploy a `GreengrassContainer` Lambda function to an AWS IoT Greengrass core running in a Docker container. Currently, only `NoContainer` Lambda functions can be deployed to a Greengrass Docker container.

To resolve this issue, [make sure that all Lambda functions are in `NoContainer` mode](#) and start a new deployment. Then, when starting the container, don't bind-mount the existing deployment directory onto the AWS IoT Greengrass core Docker container. Instead, create an empty deployment directory in its place and bind-mount that in the Docker container. This allows the new Docker container to receive the latest deployment with Lambda functions running in `NoContainer` mode.

For more information, see [the section called "Run AWS IoT Greengrass in a Docker container"](#).

Troubleshooting with logs

You can configure logging settings for a Greengrass group, such as whether to send logs to CloudWatch Logs, store logs on the local file system, or both. To get detailed information when troubleshooting issues, you can temporarily change the logging level to `DEBUG`. Changes to logging settings take effect when you deploy the group. For more information, see [the section called "Configure logging for AWS IoT Greengrass"](#).

On the local file system, AWS IoT Greengrass stores logs in the following locations. Reading the logs on the file system requires root permissions.

greengrass-root/ggc/var/log/crash.log

Shows messages generated when an AWS IoT Greengrass core crashes.

greengrass-root/ggc/var/log/system/runtime.log

Shows messages about which component failed.

greengrass-root/ggc/var/log/system/

Contains all logs from AWS IoT Greengrass system components, such as the certificate manager and the connection manager. By using the messages in `ggc/var/log/system/` and `ggc/var/`

`log/system/runtime.log`, you should be able to find out which error occurred in AWS IoT Greengrass system components.

greengrass-root/ggc/var/log/system/localwatch/

Contains the logs for the AWS IoT Greengrass component that handles uploading Greengrass logs to CloudWatch Logs. If you cannot view Greengrass logs in CloudWatch, then you can use these logs for troubleshooting.

greengrass-root/ggc/var/log/user/

Contains all logs from user-defined Lambda functions. Check this folder to find error messages from your local Lambda functions.

Note

By default, *greengrass-root* is the `/greengrass` directory. If a [write directory](#) is configured, then the logs are under that directory.

If the logs are configured to be stored on the cloud, use CloudWatch Logs to view log messages. `crash.log` is found only in file system logs on the AWS IoT Greengrass core device.

If AWS IoT is configured to write logs to CloudWatch, check those logs if connection errors occur when system components attempt to connect to AWS IoT.

For more information about AWS IoT Greengrass logging, see [the section called “Monitoring with AWS IoT Greengrass logs”](#).

Note

Logs for AWS IoT Greengrass Core software v1.0 are stored under the *greengrass-root*/`var/log` directory.

Troubleshooting storage issues

When the local file storage is full, some components might start failing:

- Local shadow updates do not occur.

- New AWS IoT Greengrass core MQTT server certificates cannot be downloaded locally.
- Deployments fail.

You should always be aware of the amount of free space available locally. You can calculate free space based on the sizes of deployed Lambda functions, the logging configuration (see [the section called “Troubleshooting with logs”](#)), and the number of shadows stored locally.

Troubleshooting messages

All messages sent locally in AWS IoT Greengrass are sent with QoS 0. By default, AWS IoT Greengrass stores messages in an in-memory queue. Therefore, unprocessed messages are lost when the Greengrass core restarts; for example, after a group deployment or device reboot. However, you can configure AWS IoT Greengrass (v1.6 or later) to cache messages to the file system so they persist across core restarts. You can also configure the queue size. If you configure a queue size, make sure that it's greater than or equal to 262144 bytes (256 KB). Otherwise, AWS IoT Greengrass might not start properly. For more information, see [the section called “MQTT message queue”](#).

Note

When using the default in-memory queue, we recommend that you deploy groups or restart the device when the service disruption is the lowest.

You can also configure the core to establish persistent sessions with AWS IoT. This allows the core to receive messages sent from the AWS Cloud while the core is offline. For more information, see [the section called “MQTT persistent sessions with AWS IoT Core”](#).

Troubleshooting shadow synchronization timeout issues

Significant delays in communication between a Greengrass core device and the cloud might cause shadow synchronization to fail because of a timeout. In this case, you should see log entries similar to the following:

```
[2017-07-20T10:01:58.006Z][ERROR]-cloud_shadow_client.go:57,Cloud shadow
client error: unable to get cloud shadow what_the_thing_is_named for
synchronization. Get https://1234567890abcd.iot.us-west-2.amazonaws.com:8443/things/
```

```
what_the_thing_is_named/shadow: net/http: request canceled (Client.Timeout exceeded
while awaiting headers)
[2017-07-20T10:01:58.006Z][WARN]-sync_manager.go:263,Failed to get cloud
copy: Get https://1234567890abcd.iot.us-west-2.amazonaws.com:8443/things/
what_the_thing_is_named/shadow: net/http: request canceled (Client.Timeout exceeded
while awaiting headers)
[2017-07-20T10:01:58.006Z][ERROR]-sync_manager.go:375,Failed to execute sync operation
{what_the_thing_is_named VersionDiscontinued []}]"
```

A possible fix is to configure the amount of time that the core device waits for a host response. Open the [config.json](#) file in *greengrass-root*/config and add a `system.shadowSyncTimeout` field with a timeout value in seconds. For example:

```
{
  "system": {
    "shadowSyncTimeout": 10
  },
  "coreThing": {
    "caPath": "root-ca.pem",
    "certPath": "cloud.pem.crt",
    "keyPath": "cloud.pem.key",
    ...
  },
  ...
}
```

If no `shadowSyncTimeout` value is specified in `config.json`, the default is 5 seconds.

Note

For AWS IoT Greengrass Core software v1.6 and earlier, the default `shadowSyncTimeout` is 1 second.

Check AWS re:Post

If you're unable to resolve your issue using the troubleshooting information in this topic, you can search the [Troubleshooting](#) or check the [AWS IoT Greengrass tag on AWS re:Post](#) for related issues or post a new question. Members of the AWS IoT Greengrass team actively monitor AWS re:Post.

Document history for AWS IoT Greengrass

The following table describes important changes to the AWS IoT Greengrass Developer Guide after June 2018. For notification about updates to this documentation, you can subscribe to an RSS feed.

Change	Description	Date
Update to end of support for v1.11.x Snap	Updated the end of support information for AWS IoT Greengrass core v 1.11.x Snap on snapcraft.io .	September 22, 2023
End of support for v1.11.x Snap	Added end of support information for AWS IoT Greengrass core v 1.11.x Snap on snapcraft.io .	September 19, 2023
Docker images for AWS IoT Greengrass v1.11.6	The Docker images for AWS IoT Greengrass Core software v1.11.6 are available on Amazon Elastic Container Registry (Amazon ECR) and Docker Hub. We recommend that you always run the latest version.	April 12, 2022
AWS IoT Device Tester (IDT) for AWS IoT Greengrass V1 deprecation	IDT for AWS IoT Greengrass V1 will no longer generate signed qualification reports.	April 4, 2022
Support update for AWS IoT Device Tester for AWS IoT Greengrass	IDT for AWS IoT Greengrass version 4.4.1 now supports using AWS IoT Greengrass core software version v1.11.6 for device qualification.	March 24, 2022
AWS IoT Greengrass version 1.11.6 released	Version 1.11.6 of the AWS IoT Greengrass Core software	March 24, 2022

is available. This version contains performance improvements and bug fixes. We recommend that you always run the latest version.

[IoT SiteWise connector version 12 released](#)

Version 12 of the IoT SiteWise connector is available. This release contains bug fixes.

December 23, 2021

[Docker images for AWS IoT Greengrass v1.11.5 and v1.10.5](#)

The Docker images for AWS IoT Greengrass Core software v1.11.5 and v1.10.5 are available on Amazon Elastic Container Registry (Amazon ECR) and Docker Hub. We recommend that you always run the latest version.

December 22, 2021

[AWS IoT Greengrass V1 maintenance policy](#)

The AWS IoT Greengrass V1 maintenance policy defines the different levels of maintenance and updates for the AWS IoT Greengrass V1 service and the AWS IoT Greengrass core software v1.x.

December 20, 2021

[AWS IoT Device Tester version 4.4.1 released](#)

IDT for AWS IoT Greengrass version 4.4.1 is now available . This release includes the AWS IoT Greengrass qualification suite (GGQ) v1.3.1, and supports using AWS IoT Greengrass core software versions v1.11.5 and v1.10.5 for device qualification.

December 20, 2021

[AWS IoT Greengrass versions 1.11.5 and 1.10.5 released](#)

Versions 1.11.5 and 1.10.5 of the AWS IoT Greengrass Core software are available. These versions contain performance improvements and bug fixes. We recommend that you always run the latest version.

December 12, 2021

[Republished AWS IoT Greengrass v1.11.4 and v1.10.4 Docker images](#)

Docker images for AWS IoT Greengrass Core software versions 1.11.4 and 1.10.4 have been republished on Amazon Elastic Container Registry (Amazon ECR) and Docker Hub to address bug fixes with BusyBox. To use the latest Docker images, use the 1.11.4-1 or 1.10.4-1 tags. For more information about available tags, see [amazon/aws-iot-greengrass](#) in *Docker Hub*.

December 8, 2021

[CloudWatch Metrics connector supports duplicate timestamps in input data](#)

You can now send input data with duplicate timestamps to this connector.

November 19, 2021

[Cross-service confused deputy prevention update](#)

AWS IoT Greengrass supports using the [aws:SourceArn](#) and [aws:SourceAccount](#) global condition context keys in IAM resource policies to prevent the confused deputy problem.

November 1, 2021

[Docker images for AWS IoT Greengrass v1.11.4](#)

The Docker images for AWS IoT Greengrass Core software v1.11.4 are available on Amazon Elastic Container Registry (Amazon ECR) and Docker Hub. We recommend that you always run the latest version.

August 24, 2021

[Published AWS IoT Greengrass v1.11.4 snap](#)

Version 1.11.4 of the AWS IoT Greengrass snap is available . We recommend that you always run the latest version.

August 20, 2021

[Support update for AWS IoT Device Tester for AWS IoT Greengrass](#)

IDT for AWS IoT Greengrass version 4.1.0 now supports using AWS IoT Greengrass core software version v1.11.4 for device qualification.

August 18, 2021

[AWS IoT Greengrass version 1.11.4 released](#)

Version 1.11.4 of the AWS IoT Greengrass Core software is available. This release fixes an issue with stream manager that prevented upgrades to v1.11.3 from an earlier version of the AWS IoT Greengrass Core software. We recommend that you always run the latest version.

August 17, 2021

[VPC endpoints \(AWS PrivateLink\)](#)

AWS IoT Greengrass now supports interface VPC endpoints (AWS PrivateLink) for the AWS IoT Greengrass control plane. You can establish a private connection between your VPC and the AWS IoT Greengrass control plane.

August 16, 2021

[AWS IoT Device Tester version 4.1.0 released](#)

Version 4.1.0 of AWS IoT Device Tester for AWS IoT Greengrass is available. This version supports using AWS IoT Greengrass core software versions 1.11.3 and 1.10.4 for device qualification.

June 23, 2021

[Published AWS IoT Greengrass v1.11.3 snap](#)

Version 1.11.3 of the AWS IoT Greengrass snap contains performance improvements and bug fixes. We recommend that as a best practice you always run the latest version.

June 15, 2021

[Docker images for AWS IoT Greengrass v1.11.3 and v1.10.4 released](#)

The Docker images for AWS IoT Greengrass Core software v1.11.3 and v1.10.4 are available on Amazon Elastic Container Registry (Amazon ECR) and Docker Hub. These versions of AWS IoT Greengrass Core contain performance improvements and bug fixes. We recommend that you always run the latest version.

June 15, 2021

[AWS IoT Greengrass version 1.11.3 released](#)

Version 1.11.3 of the AWS IoT Greengrass Core software is available. This version contains performance improvements and bug fixes. We recommend that you always run the latest version.

June 14, 2021

[AWS IoT Greengrass version 1.10.4 released](#)

Version 1.10.4 of the AWS IoT Greengrass Core software is available. This version contains performance improvements and bug fixes. We recommend that you always run the latest version.

June 14, 2021

[Modbus-TCP Protocol Adapter version 2 released](#)

Version 2 of the Modbus-TCP Protocol Adapter connector is available. This release added support for ASCII, UTF8, and ISO8859 encoded source strings.

May 24, 2021

[Docker application deployment connector version 7 released](#)

Version 7 of the Greengrass Docker application deployment connector is available.

April 5, 2021

[AWS IoT Greengrass version 1.11.1 released](#)

Version 1.11.1 of the AWS IoT Greengrass Core software is available. This version contains performance improvements and bug fixes. We recommend that you always run the latest version.

March 29, 2021

[AWS IoT Device Tester version 4.0.2 released](#)

Version 4.0.2 of AWS IoT Device Tester for AWS IoT Greengrass is available. This version replaces IDT v4.0.0 and adds support for version 1.11.1 of AWS IoT Greengrass Core software. This also fixes an issue that caused IDT to mask Hardware Security Integration (HSI) errors.

March 29, 2021

[IoT SiteWise connector version 11 released](#)

Version 11 of the IoT SiteWise connector is available. This launches support for strings that contain hidden or unprintable characters. This release also includes general performance improvements and bug fixes.

March 24, 2021

[Republished AWS IoT Greengrass v1.11.0 snap](#)

AWS IoT Greengrass snap version 1.11.0 has been republished on Snapcraft to address bug fixes and a possible application crash when using the Python interpreter. AWS IoT Greengrass doesn't provide snaps for software versions 1.10 and 1.9.

March 19, 2021

[Support update for AWS IoT Device Tester for AWS IoT Greengrass](#)

IDT for AWS IoT Greengrass version 4.0.0 now supports using AWS IoT Greengrass core software version v1.10.3 for device qualification.

March 18, 2021

[Republished AWS IoT Greengrass v1.8.4 snap](#)

AWS IoT Greengrass snap version 1.8.4 has been republished on Snapcraft to address bug fixes and a possible application crash when using the Python interpreter.

March 15, 2021

[Republished AWS IoT Greengrass v1.11.0 Docker image for Armv7l](#)

The Docker image for AWS IoT Greengrass Core software version 1.11.0 for the Armv7l platform has been republished on Amazon Elastic Container Registry (Amazon ECR) and Docker Hub to address bug fixes and a possible application crash when using the Python interpreter.

March 8, 2021

[AWS IoT Greengrass v1.10.3 Docker images released](#)

Docker images for AWS IoT Greengrass Core software version 1.10.3 are now available on Amazon Elastic Container Registry (Amazon ECR) and Docker Hub.

March 8, 2021

[Republished AWS IoT Greengrass v1.11.0 and v1.9.4 Docker images](#)

Docker images for AWS IoT Greengrass Core software versions 1.11.0 and 1.9.4 have been republished on Amazon Elastic Container Registry (Amazon ECR) and Docker Hub to address bug fixes and a possible application crash when using the Python interpreter. The Docker images for Armv7l have not been republished at this time.

February 26, 2021

[AWS IoT Greengrass version 1.10.3 released](#)

Version 1.10.3 of the AWS IoT Greengrass Core software is available. This version adds the `systemComponentAuthTimeout` core configuration property and contains performance improvements and bug fixes. We recommend that you always run the latest version.

February 24, 2021

[IoT SiteWise connector version 10 released](#)

Version 10 of the IoT SiteWise connector is available. This release resolves stability issues with the StreamManager client when connection is lost, and improves OPC-UA value handling when a `SourceTimestamp` is absent. Use the IoT SiteWise connector to send local device and equipment data to asset properties in IoT SiteWise.

January 22, 2021

[IoT SiteWise connector
version 9 released](#)

Version 9 of the IoT SiteWise connector is available. This launches support for custom Greengrass StreamManager stream destinations, OPC-UA deadbanding, custom scan mode and custom scan rate. This also includes improved performance during configuration updates made from the IoT SiteWise gateway. Use the IoT SiteWise connector to send local device and equipment data to asset properties in IoT SiteWise.

December 15, 2020

[AWS IoT Device Tester version
4.0.0 released](#)

Version 4.0.0 of AWS IoT Device Tester for AWS IoT Greengrass is available. This version enables you to use IDT to develop and run your custom test suites for device validation. This also includes code signed IDT applications for macOS and Windows.

December 15, 2020

[AWS IoT Greengrass snap
v1.11](#)

Version 1.11.0 of the AWS IoT Greengrass snap supports noncontainerized Lambda functions. We recommend that as a best practice you always run the latest version.

December 6, 2020

IoT SiteWise connector version 8 released	Version 8 of the IoT SiteWise connector is available. This release improves stability when the connector experiences intermittent network connectivity. Use the IoT SiteWise connector to send local device and equipment data to asset properties in IoT SiteWise.	November 19, 2020
Kinesis Firehose connector supports No container mode	You can use the Isolation Mode parameter to configure the containerization mode for the connector.	October 19, 2020
Docker application deployment connector version 6 released	Version 6 of the Greengrass Docker application deployment connector is available.	September 18, 2020
AWS IoT Greengrass version 1.11.0 released	Version 1.11.0 of the AWS IoT Greengrass Core software is available. This version adds the system health telemetry feature and a local health check API. Stream manager can now export data to Amazon Simple Storage Service (Amazon S3) and IoT SiteWise. This version also contains performance improvements and bug fixes. We recommend that you always run the latest version.	September 16, 2020

IoT SiteWise connector version 7 released	Version 7 of the IoT SiteWise connector is available. This release fixes an issue with gateway metrics. Use the IoT SiteWise connector to send local device and equipment data to asset properties in IoT SiteWise.	August 14, 2020
ServiceNow MetricBase Integration, Splunk Integration, and Twilio Notifications connectors support No container mode	You can use the Isolation Mode parameter to configure the containerization mode for the connector.	July 30, 2020
SNS connector supports No container mode	You can use the Isolation Mode parameter to configure the containerization mode for the connector.	July 6, 2020
CloudWatch Metrics connector supports No container mode	You can use the Isolation Mode parameter to configure the containerization mode for the connector.	June 17, 2020
AWS IoT Greengrass version 1.10.2 released	Version 1.10.2 of the AWS IoT Greengrass Core software is available. This version adds the <code>mqttOperationTimeout</code> core configuration property and contains performance improvements and bug fixes. We recommend that you always run the latest version.	June 8, 2020

[Tensorflow machine learning installers deprecated](#)

AWS IoT Greengrass TensorFlow prepackaged machine learning installers have been deprecated. Machine learning samples have been upgraded to Python 3.7.

May 29, 2020

[Chainer framework support and Greengrass machine learning installers deprecated](#)

AWS IoT Greengrass prepackaged machine learning installers and downloads for MXNet and DLR have been deprecated. Chainer framework support and associated downloads have been deprecated.

May 4, 2020

[IoT SiteWise connector version 6 released](#)

Version 6 of the IoT SiteWise connector is available. This release adds support for CloudWatch metrics and automatic discovery of new OPC-UA tags. This means you don't need to restart your gateway when tags change for your OPC-UA sources. This version of the connector requires stream manager and AWS IoT Greengrass Core software v1.10.0 or higher. Use the IoT SiteWise connector to send local device and equipment data to asset properties in IoT SiteWise.

April 29, 2020

Connectors upgraded to Python 3.7	Connectors that support the Python runtime have been upgraded to Python 3.7. We recommend that you upgrade your connector versions from Python 2.7 to Python 3.7.	April 29, 2020
Greengrass device setup can run in silent mode	You can run Greengrass device setup in silent mode so that the script doesn't prompt you for any values.	April 27, 2020
New Docker base images	You can download AWS IoT Greengrass Docker images that are built on Alpine Linux (x86_64, Armv7l, or AArch64) base images.	April 23, 2020
AWS IoT Greengrass version 1.10.1 released	Version 1.10.1 of the AWS IoT Greengrass Core software is available. This version contains performance improvements and bug fixes. We recommend that you always run the latest version.	April 16, 2020
New security chapter	AWS IoT Greengrass security content has been reorganized, with new information added.	March 30, 2020
Use APT package manager to install AWS IoT Greengrass	On supported Debian-based Linux distributions, you can use apt to install the AWS IoT Greengrass Core software on your devices.	February 26, 2020

[IoT SiteWise connector version 5 released](#)

Version 5 of the IoT SiteWise connector is available. This release fixes a compatibility issue with AWS IoT Greengrass Core software v1.9.4. Use the IoT SiteWise connector to send local device and equipment data to asset properties in IoT SiteWise.

February 12, 2020

[New script to quickly set up a core device](#)

You can use Greengrass device setup to configure your core device in minutes. Also, AWS IoT Greengrass now supports Node.js 12.x Lambda functions.

December 20, 2019

[AWS IoT Greengrass version 1.10.0 released](#)

Version 1.10.0 of the AWS IoT Greengrass Core software is available. The new features in this version include Stream manager, container support with the Docker application deployment connector, non-containerized Lambda functions that can access machine learning resources, support for MQTT persistent sessions with AWS IoT, and support for local MQTT traffic over a specified port.

November 25, 2019

[Console support for deployment notifications](#)

Use the Amazon EventBridge console to create event rules that trigger when your Greengrass group deployments change state.

November 14, 2019

[AWS IoT Greengrass version 1.9.4 released](#)

Version 1.9.4 of the AWS IoT Greengrass Core software is available. This version contains performance improvements and bug fixes. As a best practice, we recommend that you always run the latest version.

October 17, 2019

[Console support for managing the Greengrass service role](#)

Use new and improved features in the AWS IoT console to manage your Greengrass service role.

October 4, 2019

[Console support for managing group-level tags](#)

You can create, view, and manage tags for your Greengrass groups in the AWS IoT console.

September 23, 2019

[New machine learning connectors](#)

Use the ML Feedback connector to publish model input and predictions and the ML Object Detection connector to run a local object detection inference service.

September 19, 2019

[AWS IoT Greengrass version 1.9.3 released](#)

Version 1.9.3 of the AWS IoT Greengrass Core software is available. This version allows you to install the AWS IoT Greengrass Core software on Raspbian distributions on Armv6l architectures, supports OTA updates on port 443 with ALPN, and contains a bug fix for binary payloads sent from Python 2.7 Lambda functions to other Lambda functions.

September 12, 2019

[AWS IoT Greengrass version 1.8.4 released](#)

Version 1.8.4 of the AWS IoT Greengrass Core software is available. This version contains performance improvements and bug fixes. If you're running v1.8.x, we recommend that you upgrade to v1.8.4 or v1.9.3. For earlier versions, we recommend that you upgrade to v1.9.3.

August 30, 2019

[AWS IoT Greengrass version 1.9.2 released with support for OpenWrt](#)

Version 1.9.2 of the AWS IoT Greengrass Core software is available. This version allows you to install the AWS IoT Greengrass Core software on OpenWrt distributions with Armv8 (AArch64) and Armv7l architectures.

June 20, 2019

[AWS IoT Greengrass version 1.8.3 released](#)

Version 1.8.3 of the AWS IoT Greengrass Core software is available. This version contains general performance improvements and bug fixes. If you're running v1.8.x, we recommend that you upgrade to v1.8.3 or v1.9.2. For earlier versions, we recommend that you upgrade to v1.9.2.

June 20, 2019

[AWS IoT Greengrass version 1.9.1 released](#)

Version 1.9.1 of the AWS IoT Greengrass Core software is available. This version contains a bug fix for messages from AWS IoT that contain a wildcard character in the topic.

May 10, 2019

[AWS IoT Greengrass version 1.8.2 released](#)

Version 1.8.2 of the AWS IoT Greengrass Core software is available. This version contains general performance improvements and bug fixes. If you're running v1.8.x, we recommend that you upgrade to v1.8.2 or v1.9.0. For earlier versions, we recommend that you upgrade to v1.9.0.

May 2, 2019

[AWS IoT Greengrass version 1.9.0 released](#)

New features: Support for Python 3.7 and Node.js 8.10 Lambda runtimes, optimized MQTT connections, and Elliptic Curve (EC) key support for the local MQTT server.

May 1, 2019

[AWS IoT Greengrass version 1.8.1 released](#)

Version 1.8.1 of the AWS IoT Greengrass Core software is available. This version contains general performance improvements and bug fixes. As a best practice, we recommend that you always run the latest version.

April 18, 2019

[AWS IoT Greengrass snap available on snapcraft](#)

Use the AWS IoT Greengrass Snap Store app to quickly design, test, and deploy software on Linux devices with AWS IoT Greengrass.

April 1, 2019

[Support for more access control using tag-based permissions](#)

You can use tags in AWS Identity and Access Management (IAM) policies to control access to your AWS IoT Greengrass resources.

March 29, 2019

[IoT Analytics connector released](#)

Use the IoT Analytics connector to send local device data to AWS IoT Analytics channels.

March 15, 2019

[Batch support in Kinesis Firehose connector](#)

The Kinesis Firehose connector supports sending batched data records to Amazon Data Firehose at a specified interval.

March 15, 2019

[AWS CloudFormation support for AWS IoT Greengrass resources](#)

Use AWS CloudFormation templates to create and manage AWS IoT Greengrass resources.

March 15, 2019

[AWS IoT Greengrass version 1.8.0 released](#)

New features: Configurable default access identity for Lambda functions, support for HTTPS traffic over port 443, and predictably named client IDs for MQTT connections with AWS IoT.

March 7, 2019

[AWS IoT Greengrass versions 1.7.1 and 1.6.1 released](#)

Versions 1.7.1 and 1.6.1 of the AWS IoT Greengrass Core software are available. These versions require Linux kernel version 3.17 or later. We recommend that customers running any version of the Greengrass core software upgrade to version 1.7.1 immediately.

February 11, 2019

[SageMaker AI Neo deep learning runtime](#)

The SageMaker AI Neo deep learning runtime supports machine learning models that have been optimized by the SageMaker AI Neo deep learning compiler.

November 28, 2018

[Run AWS IoT Greengrass in a Docker container](#)

You can run AWS IoT Greengrass in a Docker container by configuring your Greengrass group to run with no containerization.

November 26, 2018

[AWS IoT Greengrass version 1.7.0 released](#)

New features: Greengrass connectors, local secrets manager, isolation and permission settings for Lambda functions, hardware root of trust security, connection using ALPN or network proxy, and Raspbian Stretch support.

November 26, 2018

[AWS IoT Greengrass software downloads](#)

The AWS IoT Greengrass Core software, AWS IoT Greengrass Core SDK, and AWS IoT Greengrass Machine Learning SDK packages are available for download through Amazon CloudFront.

November 26, 2018

[AWS IoT Device Tester for AWS IoT Greengrass](#)

Use AWS IoT Device Tester for AWS IoT Greengrass to verify that your CPU architecture, kernel configuration, and drivers work with AWS IoT Greengrass.

November 26, 2018

[AWS CloudTrail logging for AWS IoT Greengrass API calls](#)

AWS IoT Greengrass is integrated with AWS CloudTrail, a service that provides a record of actions taken by a user, role, or an AWS service in AWS IoT Greengrass.

October 29, 2018

[Support for TensorFlow v1.10.1 on NVIDIA Jetson TX2](#)

The TensorFlow precompiled library for NVIDIA Jetson TX2 that AWS IoT Greengrass provides now uses TensorFlow v1.10.1. This supports Jetpack 3.3 and CUDA Toolkit 9.0.

October 18, 2018

[Support for MXNet v1.2.1 machine learning resources](#)

AWS IoT Greengrass supports machine learning models that are trained using MXNet v1.2.1.

August 29, 2018

[AWS IoT Greengrass version 1.6.0 released](#)

New features: Lambda executables, configurable message queue, configurable reconnect retry interval, volume resources under /proc, and configurable write directory.

July 26, 2018

Earlier updates

The following table describes important changes to the AWS IoT Greengrass Developer Guide before July 2018.

Change	Description	Date
AWS IoT Greengrass Version 1.5.0 Released	New features: - Local machine learning inference using cloud-trained models. For more information, see Perform machine learning inference. - Greengrass Lambda functions support binary input data, in addition to JSON. For more information, see AWS IoT Greengrass Core versions.	March 29, 2018

Change	Description	Date
AWS IoT Greengrass Version 1.3.0 Released	New features: • Over-the-air (OTA) update agent capable of handling cloud-deployed, Greengrass update jobs. For more information, see OTA updates of AWS IoT Greengrass Core software. • Access local peripherals and resources from Greengrass Lambda functions. For more information, see Access local resources with Lambda functions and connectors.	November 27, 2017
AWS IoT Greengrass Version 1.1.0 Released	New features: • Reset deployed AWS IoT Greengrass groups. For more information, see Reset deployments. • Support for Node.js 6.10 and Java 8 Lambda runtimes, in addition to Python 2.7.	September 20, 2017
AWS IoT Greengrass Version 1.0.0 Released	AWS IoT Greengrass is generally available.	June 7, 2017