



User Guide

# AWS Fault Injection Service



# AWS Fault Injection Service: User Guide

Copyright © 2026 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

<b>What is AWS FIS?</b> .....	<b>1</b>
Concepts .....	1
Actions .....	2
Targets .....	2
Stop conditions .....	2
Supported AWS services .....	3
Access AWS FIS .....	3
Pricing .....	3
<b>Planning your experiments</b> .....	<b>4</b>
Basic principles and guidelines .....	4
Experiment planning guidelines .....	5
<b>Experiment template components</b> .....	<b>7</b>
Template syntax .....	7
Get started .....	8
Actions .....	8
Action syntax .....	8
Action identifiers .....	10
Action parameters .....	10
Action targets .....	10
Action duration .....	11
Example actions .....	12
Targets .....	15
Target syntax .....	16
Resource types .....	17
Identify target resources .....	18
Selection mode .....	23
Example targets .....	23
Example filters .....	24
Stop conditions .....	29
Stop condition syntax .....	30
Learn more .....	30
Experiment role .....	31
Prerequisites .....	31
Option 1: Create an experiment role and attach an AWS managed policy .....	33

Option 2: Create an experiment role and add an inline policy document .....	33
Experiment report configuration .....	35
Experiment report configuration syntax .....	38
Experiment report permissions .....	40
Experiment report best practices .....	42
Experiment options .....	42
Account targeting .....	43
Empty target resolution mode .....	44
Actions mode .....	45
<b>Actions reference .....</b>	<b>46</b>
Fault injection actions .....	47
aws:fis:inject-api-internal-error .....	47
aws:fis:inject-api-throttle-error .....	48
aws:fis:inject-api-unavailable-error .....	49
Recovery action .....	49
aws:arc:start-zonal-autoshift .....	49
Wait action .....	51
aws:fis:wait .....	51
Amazon CloudWatch actions .....	51
aws:cloudwatch:assert-alarm-state .....	51
Amazon DynamoDB actions .....	52
aws:dynamodb:global-table-pause-replication .....	52
Amazon Aurora DSQL actions .....	56
aws:dsql:cluster-connection-failure .....	56
Amazon EBS actions .....	57
aws:ebs:pause-volume-io .....	57
aws:ebs:volume-io-latency .....	58
Amazon EC2 actions .....	59
aws:ec2:api-insufficient-instance-capacity-error .....	59
aws:ec2:asg-insufficient-instance-capacity-error .....	60
aws:ec2:reboot-instances .....	61
aws:ec2:send-spot-instance-interruptions .....	61
aws:ec2:stop-instances .....	62
aws:ec2:terminate-instances .....	63
Amazon ECS actions .....	64
aws:ecs:drain-container-instances .....	64

aws:ecs:stop-task .....	65
aws:ecs:task-cpu-stress .....	66
aws:ecs:task-io-stress .....	66
aws:ecs:task-kill-process .....	67
aws:ecs:task-network-blackhole-port .....	68
aws:ecs:task-network-latency .....	69
aws:ecs:task-network-packet-loss .....	70
Amazon EKS actions .....	72
aws:eks:inject-kubernetes-custom-resource .....	72
aws:eks:pod-cpu-stress .....	73
aws:eks:pod-delete .....	74
aws:eks:pod-io-stress .....	75
aws:eks:pod-memory-stress .....	77
aws:eks:pod-network-blackhole-port .....	78
aws:eks:pod-network-latency .....	79
aws:eks:pod-network-packet-loss .....	80
aws:eks:terminate-nodegroup-instances .....	82
Amazon ElastiCache actions .....	82
aws:elasticache:replicationgroup-interrupt-az-power .....	82
Amazon Kinesis Data Streams actions .....	83
aws:kinesis:stream-provisioned-throughput-exception .....	84
aws:kinesis:stream-expired-iterator-exception .....	84
AWS Lambda actions .....	85
aws:lambda:invocation-add-delay .....	85
aws:lambda:invocation-error .....	86
aws:lambda:invocation-http-integration-response .....	87
Amazon MemoryDB action .....	88
aws:memorydb:multi-region-cluster-pause-replication .....	88
Network actions .....	89
aws:network:disrupt-connectivity .....	89
aws:network:route-table-disrupt-cross-region-connectivity .....	90
aws:network:transit-gateway-disrupt-cross-region-connectivity .....	92
aws:network:disrupt-vpc-endpoint .....	93
Amazon RDS actions .....	94
aws:rds:failover-db-cluster .....	94
aws:rds:reboot-db-instances .....	94

Amazon S3 actions .....	95
aws:s3:bucket-pause-replication .....	95
Systems Manager actions .....	96
aws:ssm:send-command .....	97
aws:ssm:start-automation-execution .....	97
AWS Direct Connect actions .....	98
aws:directconnect:virtual-interface-disconnect .....	99
SSM document actions .....	99
Use the aws:ssm:send-command action .....	100
Pre-configured AWS FIS SSM documents .....	101
Examples .....	111
Limitations .....	111
Rollback scripts .....	111
Troubleshooting .....	113
ECS task actions .....	113
Actions .....	114
Limitations .....	114
Requirements .....	115
Reference version of the script .....	118
Example experiment template .....	120
EKS Pod actions .....	121
Actions .....	122
Limitations .....	123
Requirements .....	124
Create an experiment role .....	124
Configure the Kubernetes service account .....	124
Grant IAM users and roles access to Kubernetes APIs .....	126
Pod container images .....	127
Example experiment template .....	129
AWS Lambda actions .....	130
Actions .....	131
Limitations .....	131
Prerequisites .....	131
Configure Lambda functions .....	133
Configure an AWS FIS experiment .....	133
Logging .....	134

Advanced topics .....	135
AWS FIS Lambda extension versions .....	141
<b>Managing experiment templates .....</b>	<b>145</b>
Create an experiment template .....	145
View experiment templates .....	148
Generate a target preview .....	148
Start an experiment from a template .....	149
Update an experiment template .....	150
Tag experiment templates .....	150
Delete an experiment template .....	151
Example templates .....	151
Stop EC2 instances based on filters .....	152
Stop a specified number of EC2 instances .....	153
Run a pre-configured AWS FIS SSM document .....	154
Run a predefined Automation runbook .....	155
Throttle API actions on EC2 instances with the target IAM role .....	156
Stress test CPU of pods in a Kubernetes cluster .....	157
Provisioned throughput exception for specified number of Kinesis Data Streams .....	159
Experiment role permissions example .....	160
<b>Managing experiments .....</b>	<b>162</b>
Start an experiment .....	162
View your experiments .....	163
Experiment states .....	163
Action states .....	164
Tag an experiment .....	164
Stop an experiment .....	165
List resolved targets .....	165
<b>Tutorials .....</b>	<b>167</b>
Test instance stop and start .....	167
Prerequisites .....	167
Step 1: Create an experiment template .....	167
Step 2: Start the experiment .....	171
Step 3: Track the experiment progress .....	171
Step 4: Verify the experiment result .....	171
Step 5: Clean up .....	172
Run CPU stress on an instance .....	172

Prerequisites .....	173
Step 1: Create a CloudWatch alarm for a stop condition .....	174
Step 2: Create an experiment template .....	174
Step 3: Start the experiment .....	176
Step 4: Track the experiment progress .....	177
Step 5: Verify the experiment results .....	177
Step 6: Clean up .....	172
Test Spot Instance interruptions .....	179
Prerequisites .....	180
Step 1: Create an experiment template .....	181
Step 2: Start the experiment .....	184
Step 3: Track the experiment progress .....	184
Step 4: Verify the experiment result .....	184
Step 5: Clean up .....	185
Simulate a connectivity event .....	186
Prerequisites .....	187
Step 1: Create an AWS FIS experiment template .....	188
Step 2: Ping an Amazon S3 endpoint .....	189
Step 3: Start your AWS FIS experiment .....	190
Step 4: Track your AWS FIS experiment progress .....	190
Step 5: Verify Amazon S3 network disruption .....	191
Step 5: Clean up .....	191
Schedule a recurring experiment .....	192
Prerequisites .....	192
Step 1: Create an IAM role and policy .....	192
Step 2: Create an Amazon EventBridge Scheduler .....	195
Step 3: Verify your experiment .....	196
Step 4: Clean up .....	196
<b>Working with the scenario library .....</b>	<b>197</b>
Viewing a scenario .....	197
Using a scenario .....	198
Exporting a scenario .....	199
Scenarios reference .....	199
AZ Availability: Power Interruption .....	202
AZ: Application Slowdown .....	217
Cross-AZ: Traffic Slowdown .....	223

Cross-Region: Connectivity .....	230
<b>Working with multi-account experiments .....</b>	<b>245</b>
Concepts .....	246
Best practices .....	246
Prerequisites .....	246
Permissions .....	247
Stop conditions (optional) .....	250
Safety levers for multi-account experiments (optional) .....	250
Create a multi-account experiment template .....	250
Update a target account configuration .....	252
Delete a target account configuration .....	252
<b>Scheduling experiments .....</b>	<b>254</b>
Create a scheduler role .....	254
Create an experiment schedule .....	258
To update schedule using the console .....	259
Update an experiment schedule .....	259
Disable or delete an experiment schedule .....	260
<b>Safety levers .....</b>	<b>261</b>
Concepts for safety levers .....	261
Safety lever resource .....	261
Working with safety levers .....	262
Viewing a safety lever .....	262
Engaging a safety lever .....	263
Disengaging a safety lever .....	263
<b>Monitoring experiments .....</b>	<b>264</b>
Monitor using CloudWatch .....	265
Monitor AWS FIS experiments .....	265
AWS FIS usage metrics .....	266
Monitor using EventBridge .....	267
Experiment logging .....	268
Permissions .....	269
Log schema .....	269
Log destinations .....	270
Example log records .....	271
Enable experiment logging .....	276
Disable experiment logging .....	276

Log API calls with AWS CloudTrail .....	277
Use CloudTrail .....	277
Understand AWS FIS log file entries .....	278
<b>Troubleshooting .....</b>	<b>283</b>
Error codes .....	283
<b>Security .....</b>	<b>286</b>
Data protection .....	286
Encryption at rest .....	287
Encryption in transit .....	288
Identity and access management .....	288
Audience .....	288
Authenticating with identities .....	289
Managing access using policies .....	290
How AWS Fault Injection Service works with IAM .....	292
Policy examples .....	297
Use service-linked roles .....	307
AWS managed policies .....	309
Infrastructure security .....	314
AWS PrivateLink .....	314
Considerations .....	315
Create an interface VPC endpoint .....	315
Create a VPC endpoint policy .....	315
<b>Tagging your resources .....</b>	<b>317</b>
Tagging restrictions .....	317
Work with tags .....	317
<b>Quotas and limitations .....</b>	<b>319</b>
<b>Document history .....</b>	<b>333</b>

# What is AWS Fault Injection Service?

AWS Fault Injection Service (AWS FIS) is a managed service that enables you to perform fault injection experiments on your AWS workloads. Fault injection is based on the principles of chaos engineering. These experiments stress an application by creating disruptive events so that you can observe how your application responds. You can then use this information to improve the performance and resiliency of your applications so that they behave as expected.

To use AWS FIS, you set up and run experiments that help you create the real-world conditions needed to uncover application issues that can be difficult to find otherwise. AWS FIS provides templates that generate disruptions, and the controls and guardrails that you need to run experiments in production, such as automatically rolling back or stopping the experiment if specific conditions are met.

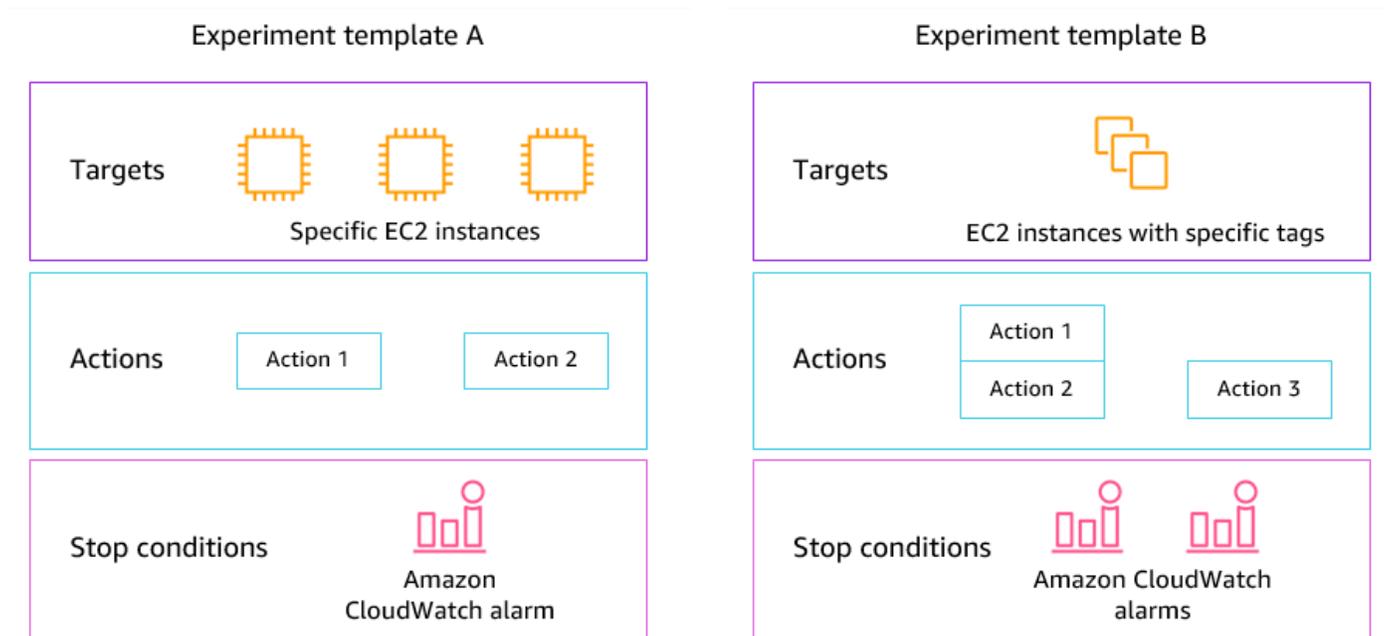
## Important

AWS FIS carries out real actions on real AWS resources in your system. Therefore, before you use AWS FIS to run experiments in production, we strongly recommend that you complete a planning phase and run the experiments in a pre-production environment.

For more information about planning your experiment, see [Test Reliability](#) and [Planning your AWS FIS experiments](#). For more information about AWS FIS, see [AWS Fault Injection Service](#).

## AWS FIS concepts

To use AWS FIS, you run *experiments* on your AWS resources to test your theory of how an application or system will perform under fault conditions. To run experiments, you first create an *experiment template*. An experiment template is the blueprint of your experiment. It contains the *actions*, *targets*, and *stop conditions* for the experiment. After you create an experiment template, you can use it to run an experiment. While your experiment is running, you can track its progress and view its status. An experiment is complete when all of the actions in the experiment have run.



## Actions

An *action* is an activity that AWS FIS performs on an AWS resource during an experiment. AWS FIS provides a set of preconfigured actions based on the type of AWS resource. Each action runs for a specified duration during an experiment, or until you stop the experiment. Actions can run sequentially or simultaneously (in parallel).

## Targets

A *target* is one or more AWS resources on which AWS FIS performs an action during an experiment. You can choose specific resources, or you can select a group of resources based on specific criteria, such as tags or state.

## Stop conditions

AWS FIS provides the controls and guardrails that you need to run experiments safely on your AWS workloads. A *stop condition* is a mechanism to stop an experiment if it reaches a threshold that you define as an Amazon CloudWatch alarm. If a stop condition is triggered while the experiment is running, AWS FIS stops the experiment.

## Supported AWS services

AWS FIS provides preconfigured actions for specific types of targets across AWS services. For the list of supported services and their actions, see [AWS FIS Actions reference](#).

For single-account experiments, the target resources must be in the same AWS account as the experiment. You can run AWS FIS experiments that target resources in a different AWS account using AWS FIS multi-account experiments.

For more information, see [Actions for AWS FIS](#).

## Access AWS FIS

You can work with AWS FIS in any of the following ways:

- **AWS Management Console** — Provides a web interface that you can use to access AWS FIS. For more information, see [Working with the AWS Management Console](#).
- **AWS Command Line Interface (AWS CLI)** — Provides commands for a broad set of AWS services, including AWS FIS, and is supported on Windows, macOS, and Linux. For more information, see [AWS Command Line Interface](#). For more information about the commands for AWS FIS, see [fis](#) in the *AWS CLI Command Reference*.
- **AWS CloudFormation** — Create templates that describe your AWS resources. You use the templates to provision and manage these resources as a single unit. For more information, see the [AWS Fault Injection Service resource type reference](#).
- **AWS SDKs** — Provides language-specific APIs and takes care of many of the connection details, such as calculating signatures, handling request retries, and handling errors. For more information, see [AWS SDKs](#).
- **HTTPS API** — Provides low-level API actions that you can call using HTTPS requests. For more information, see the [AWS Fault Injection Service API Reference](#).

## Pricing for AWS FIS

You are charged per minute that an action runs, from start to finish, based on the number of target accounts for your experiment. For more information, see [AWS FIS Pricing](#).

# Planning your AWS FIS experiments

Fault injection is the process of stressing an application in testing or production environments by creating disruptive events, such as server outages or API throttling. From observing how the system responds, you can then implement improvements. When you run experiments on your system, it can help you to identify systemic weaknesses in a controlled manner, before those weaknesses affect the customers who depend on your system. Then you can proactively address the issues to help prevent unpredictable outcomes.

Before you get started running fault injection experiments using AWS FIS, we recommend that you familiarize yourself with the following principles and guidelines.

## Important

AWS FIS carries out real actions on real AWS resources in your system. Therefore, before you get started using AWS FIS to run experiments, we strongly recommend that you first complete a planning phase and a test in a pre-production or test environment.

## Contents

- [Basic principles and guidelines](#)
- [Experiment planning guidelines](#)

## Basic principles and guidelines

Before starting experiments with AWS FIS, take the following steps:

1. **Identify the target deployment for the experiment** — Start by identifying the target deployment. If this is your first experiment, we recommend starting in a pre-production or test environment.
2. **Review the application architecture** — You must ensure that you have identified all of the application components, dependencies, and recovery procedures for each component. Begin with reviewing the application architecture. Depending on the application, refer to the [AWS Well-Architected Framework](#).

3. **Define steady-state behavior** — Define the steady-state behavior of your system in terms of important technical and business metrics, such as latency, CPU load, failed sign-ins per minute, number of retries, or page load speed.
4. **Form a hypothesis** — Form a hypothesis of how you expect the system behavior to change during the experiment. A hypothesis definition follows this format:

If *fault injection action* is performed, the *business or technical metric impact* should not exceed *value*.

For example, a hypothesis for an authentication service might read as follows: "If network latency increases by 10%, there is less than a 1% increase in sign-in failures." After the experiment is completed, you evaluate whether the application resiliency aligns with your business and technical expectations.

We also recommend following these guidelines when working with AWS FIS:

- Always start experimenting with AWS FIS in a test environment. Never start with a production environment. As you progress in your fault injection experiments, you can experiment in other controlled environments beyond the test environment.
- Build your team's confidence in your application resilience by starting with small, simple experiments, such as running the **aws:ec2:stop-instances** action on one target.
- Fault injection can cause real issues. Proceed with caution and make sure that your first fault injections are on test instances so that no customers are affected.
- Test, test, and test some more. Fault injection is meant to be implemented in a controlled environment with well-planned experiments. This allows you to build confidence in the abilities of your application and your tools to withstand turbulent conditions.
- We strongly recommend that you have an excellent monitoring and alerting program in place before you begin. Without it, you won't be able to understand or measure the impact of your experiments, which is critical to sustainable fault injection practices.

## Experiment planning guidelines

With AWS FIS, you run experiments on your AWS resources to test your theory of how an application or system will perform under fault conditions.

The following are recommended guidelines for planning your AWS FIS experiments.

- **Review outage history** — Review the previous outages and events for your system. This can help you to build up a picture of the overall health and resiliency of your system. Before you start running experiments on your system, you should address known issues and weaknesses in your system.
- **Identify services with the largest impact** — Review your services and identify the ones that have the biggest impact on your end users or customers if they go down or do not function correctly.
- **Identify the target system** — The target system is the system on which you will run experiments. If you are new to AWS FIS or you have never run fault injection experiments before, we recommend that you start by running experiments on a pre-production or test system.
- **Consult with your team** — Ask what they are worried about. You can form a hypothesis to prove or disprove their concerns. You can also ask your team what they are not worried about. This question can reveal two common fallacies: the sunk cost fallacy and the confirmation bias fallacy. Forming a hypothesis based on your team's answers can help provide more information about the reality of your system's state.
- **Review your application architecture** — Conduct a review of your system or application and ensure that you have identified all of the application components, dependencies, and recovery procedures for every component.

We recommend that you review the AWS Well-Architected Framework. The framework can help you build secure, high-performing, resilient, and efficient infrastructure for your applications and workloads. For more information, see [AWS Well-Architected](#).

- **Identify the applicable metrics** — You can monitor the impact of an experiment on your AWS resources using Amazon CloudWatch metrics. You can use these metrics to determine the baseline or "steady state" when your application is performing optimally. Then, you can monitor these metrics during or after the experiment to determine the impact. For more information, see [Monitor AWS FIS usage metrics using Amazon CloudWatch](#).
- **Define an acceptable performance threshold for your system** — Identify the metric that represents an acceptable, steady state for your system. You will use this metric to create one or more CloudWatch alarms that represent a stop condition for your experiment. If the alarm is triggered, the experiment is automatically stopped. For more information, see [Stop conditions for AWS FIS](#).

# AWS FIS experiment template components

You use the following components to construct experiment templates:

## Actions

The [AWS FIS actions](#) that you want to run. Actions can be run in a set order that you specify, or they can be run simultaneously. For more information, see [Actions](#).

## Targets

The AWS resources on which a specific action is carried out. For more information, see [Targets](#).

## Stop conditions

The CloudWatch alarms that define a threshold at which your application performance is not acceptable. If a stop condition is triggered while an experiment is running, AWS FIS stops the experiment. For more information, see [Stop conditions](#).

## Experiment role

An IAM role that grants AWS FIS the permissions required so that it can run experiments on your behalf. For more information, see [Experiment role](#).

## Experiment report configuration

The configuration to enable experiment reports. For more information, see [Experiment report configurations for AWS FIS](#).

## Experiment options

Options for the experiment template. For more information, see [Experiment options for AWS FIS](#).

Your account has quotas related to AWS FIS. For example, there is a quota on the number of actions per experiment template. For more information, see [Quotas and limitations](#).

## Template syntax

The following is the syntax for an experiment template.

```
{
```

```
    "description": "string",
    "targets": {},
    "actions": {},
    "stopConditions": [],
    "roleArn": "arn:aws:iam::123456789012:role/AllowFISActions",
    "experimentReportConfiguration": {},
    "experimentOptions": {},
    "tags": {}
}
```

For examples, see [Example templates](#).

## Get started

To create an experiment template using the AWS Management Console, see [Create an experiment template](#).

To create an experiment template using the AWS CLI, see [Example AWS FIS experiment templates](#).

## Actions for AWS FIS

To create an experiment template, you must define one or more actions. For a list of predefined actions provided by AWS FIS, see [Actions reference](#).

You can run an action only once during an experiment. To run the same AWS FIS action more than once in the same experiment, add it to the template multiple times using different names.

### Contents

- [Action syntax](#)
- [Action identifiers](#)
- [Action parameters](#)
- [Action targets](#)
- [Action duration](#)
- [Example actions](#)

## Action syntax

The following is the syntax for an action.

```
{
  "actions": {
    "action_name": {
      "actionId": "aws:service:action-type",
      "description": "string",
      "parameters": {
        "name": "value"
      },
      "startAfter": ["action_name", ...],
      "targets": {
        "ResourceType": "target_name"
      }
    }
  }
}
```

When you define an action, you provide the following:

### ***action\_name***

A name for the action.

### **actionId**

The [action identifier](#).

### **description**

An optional description.

### **parameters**

Any [action parameters](#).

### **startAfter**

Any actions that must complete before this action can start. Otherwise, the action runs at the start of the experiment.

### **targets**

Any [action targets](#).

For examples, see [the section called "Example actions"](#).

## Action identifiers

Each AWS FIS action has an identifier with the following format:

```
aws:service-name:action-type
```

For example, the following action stops the target Amazon EC2 instances:

```
aws:ec2:stop-instances
```

For a complete list of actions, see the [AWS FIS Actions reference](#).

## Action parameters

Some AWS FIS actions have additional parameters that are specific to the action. These parameters are used to pass information to AWS FIS when the action is run.

AWS FIS supports custom fault types using the `aws:ssm:send-command` action, which uses the SSM Agent and an SSM command document to create the fault condition on the targeted instances. The `aws:ssm:send-command` action includes a `documentArn` parameter that takes the Amazon Resource Name (ARN) of an SSM document as a value. You specify values for parameters when you add the action to your experiment template.

For more information about specifying parameters for the `aws:ssm:send-command` action, see [Use the aws:ssm:send-command action](#).

Where possible, you can input a rollback configuration (also referred to as a *post action*) within the action parameters. A post action returns the target to the state that it was in before the action was run. The post action runs after the time specified in the action duration. Not all actions can support post actions. For example, if the action terminates an Amazon EC2 instance, you cannot recover the instance after it has been terminated.

## Action targets

An action runs on the target resources that you specify. After you define a target, you can specify its name when you define an action.

```
"targets": {  
  "ResourceType": "resource_name"  
}
```

AWS FIS actions support the following resource types for action targets:

- **AutoScalingGroups** – Amazon EC2 Auto Scaling groups
- **Buckets** – Amazon S3 buckets
- **Cluster** – Amazon EKS clusters
- **Clusters** – Amazon ECS, Aurora DSQL, or Amazon Aurora DB clusters
- **DBInstances** – Amazon RDS DB instances
- **Functions** – AWS Lambda functions
- **Instances** – Amazon EC2 instances
- **KinesisStreams** – Kinesis data streams
- **ManagedResources** – Amazon EKS clusters, Amazon EC2 Application and Network Load Balancers, and Amazon EC2 Auto Scaling groups that are enabled for ARC zonal shift.
- **MultiRegionClusters** – Amazon MemoryDB multi-Region clusters
- **Nodegroups** – Amazon EKS node groups
- **Pods** – Kubernetes pods on Amazon EKS
- **ReplicationGroups** – ElastiCache Replication Groups
- **Roles** – IAM roles
- **SpotInstances** – Amazon EC2 Spot Instances
- **Subnets** – VPC subnets
- **Tables** – Amazon DynamoDB multi-Region strongly and eventually consistent global tables
- **Tasks** – Amazon ECS tasks
- **TransitGateways** – Transit gateways
- **VirtualInterfaces** – Direct Connect Virtual Interfaces
- **Volumes** – Amazon EBS volumes
- **VPCEndpoints** – Amazon VPC endpoints

For examples, see [the section called “Example actions”](#).

## Action duration

If an action includes a parameter that you can use to specify the duration of the action, by default, the action is considered complete only after the specified duration has elapsed. If you have set the `emptyTargetResolutionMode` experiment option to `skip`, then the action will complete

immediately with status 'skipped' when no targets were resolved. For example, if you specify a duration of 5 minutes, AWS FIS considers the action complete after 5 minutes. It then starts the next action, until all actions are complete.

Duration can be either the length of time that an action condition is maintained or the length of time for which metrics are monitored. For example, latency is injected for the duration of time specified. For near instantaneous action types, such as terminating an instance, stop conditions are monitored for the duration of time specified.

If an action includes a post action within the action parameters, the post action runs after the action completes. The time it takes to complete the post action might cause a delay between the specified action duration and the beginning of the next action (or the end of the experiment, if all other actions are complete).

## Example actions

The following are example actions.

### Examples

- [Stop EC2 instances](#)
- [Interrupt Spot Instances](#)
- [Disrupt network traffic](#)
- [Terminate EKS workers](#)
- [Start ARC zonal autoshift](#)

### Example: Stop EC2 instances

The following action stops the EC2 instances identified using the target named *targetInstances*. After two minutes, it restarts the target instances.

```
"actions": {
  "stopInstances": {
    "actionId": "aws:ec2:stop-instances",
    "parameters": {
      "startInstancesAfterDuration": "PT2M"
    },
    "targets": {
      "Instances": "targetInstances"
    }
  }
}
```

```

    }
  }
}

```

### Example: Interrupt Spot Instances

The following action stops the Spot Instances identified using the target named *targetSpotInstances*. It waits two minutes before interrupting the Spot Instance.

```

"actions": {
  "interruptSpotInstances": {
    "actionId": "aws:ec2:send-spot-instance-interruptions",
    "parameters": {
      "durationBeforeInterruption": "PT2M"
    },
    "targets": {
      "SpotInstances": "targetSpotInstances"
    }
  }
}

```

### Example: Disrupt network traffic

The following action denies traffic between the target subnets and subnets in other Availability Zones.

```

"actions": {
  "disruptAZConnectivity": {
    "actionId": "aws:network:disrupt-connectivity",
    "parameters": {
      "scope": "availability-zone",
      "duration": "PT5M"
    },
    "targets": {
      "Subnets": "targetSubnets"
    }
  }
}

```

### Example: Terminate EKS workers

The following action terminates 50% of the EC2 instances in the EKS cluster identified using the target named *targetNodeGroups*.

```
"actions": {
  "terminateWorkers": {
    "actionId": "aws:eks:terminate-nodegroup-instances",
    "parameters": {
      "instanceTerminationPercentage": "50"
    },
    "targets": {
      "Nodegroups": "targetNodeGroups"
    }
  }
}
```

### Example: Start ARC zonal autoshift

The following action starts an ARC Zonal autoshift that shifts managed resources away from *az-in-parameters* for *duration-in-parameters*. The resource type ManagedResources is used as a key for the target name in AWS FIS experiment template.

```
{
  "description": "aaa",
  "targets": {
    "ManagedResources-Target-1": {
      "resourceType": "aws:arc:zonal-shift-managed-resource",
      "resourceArns": [
        "arn:aws:elasticloadbalancing:us-east-1:0124567890:loadbalancer/app/application/11223312312516",
      ],
      "selectionMode": "ALL"
    }
  },
  "actions": {
    "arc": {
      "actionId": "aws:arc:start-zonal-autoshift",
      "parameters": {
        "availabilityZoneIdentifier": "us-east-1a",
        "duration": "PT1M"
      },
      "targets": {
        "ManagedResources": "ManagedResources-Target-1"
      }
    }
  }
}
```

```
    }
  }
},
"stopConditions": [
  {
    "source": "none"
  }
],
"roleArn": "arn:aws:iam::718579638765:role/fis",
"tags": {},
"experimentOptions": {
  "accountTargeting": "single-account",
  "emptyTargetResolutionMode": "fail"
}
}
```

## Targets for AWS FIS

A target is one or more AWS resources on which an action is performed by AWS Fault Injection Service (AWS FIS) during an experiment. Targets can be in the same AWS account as the experiment, or in a different account using a multi-account experiment. To learn more about targeting resources in a different account, see [Working with multi-account experiments](#).

You define targets when you [create an experiment template](#). You can use the same target for multiple actions in your experiment template.

AWS FIS identifies all targets at the start of the experiment, before starting any of the actions in the actions set. AWS FIS uses the target resources that it selects for the entire experiment. If no targets are found, the experiment fails.

### Contents

- [Target syntax](#)
- [Resource types](#)
- [Identify target resources](#)
  - [Resource filters](#)
  - [Resource parameters](#)
- [Selection mode](#)
- [Example targets](#)

- [Example filters](#)

## Target syntax

The following is the syntax for a target.

```
{
  "targets": {
    "target_name": {
      "resourceType": "resource-type",
      "resourceArns": [
        "resource-arn"
      ],
      "resourceTags": {
        "tag-key": "tag-value"
      },
      "parameters": {
        "parameter-name": "parameter-value"
      },
      "filters": [
        {
          "path": "path-string",
          "values": ["value-string"]
        }
      ],
      "selectionMode": "value"
    }
  }
}
```

When you define a target, you provide the following:

### target\_name

A name for the target.

### resourceType

The [resource type](#).

### resourceArns

The Amazon Resource Names (ARN) of specific resources.

## resourceTags

The tags applied to specific resources.

## parameters

The [parameters](#) that identify targets using specific attributes.

## filters

The [resource filters](#) scopes the identified target resources using specific attributes.

## selectionMode

The [selection mode](#) for the identified resources.

For examples, see [the section called "Example targets"](#).

## Resource types

Each AWS FIS action is performed on a specific AWS resource type. When you define a target, you must specify exactly one resource type. When you specify a target for an action, the target must be the resource type supported by the action.

The following resource types are supported by AWS FIS:

- **aws:arc:zonal-shift-managed-resource** – An AWS resource that is registered with ARC zonal shift
- **aws:directconnect:virtual-interface** – A Direct Connect Virtual Interface
- **aws:dsq:cluster** – An Amazon Aurora DSQL cluster
- **aws:dynamodb:global-table** – An Amazon DynamoDB multi-Region global table
- **aws:ec2:autoscaling-group** – An Amazon EC2 Auto Scaling group
- **aws:ec2:ebs-volume** – An Amazon EBS volume
- **aws:ec2:instance** – An Amazon EC2 instance
- **aws:ec2:spot-instance** – An Amazon EC2 Spot Instance
- **aws:ec2:subnet** – An Amazon VPC subnet
- **aws:ec2:transit-gateway** – A transit gateway
- **aws:ec2:vpc-endpoint** – An Amazon VPC Endpoint

- **aws:ecs:cluster** – An Amazon ECS cluster
- **aws:ecs:task** – An Amazon ECS task
- **aws:eks:cluster** – An Amazon EKS cluster
- **aws:eks:nodegroup** – An Amazon EKS node group
- **aws:eks:pod** – A Kubernetes pod
- **aws:elasticache:replicationgroup** – An ElastiCache Replication Group
- **aws:iam:role** – An IAM role
- **aws:kinesis:stream** – An Amazon Kinesis data stream
- **aws:lambda:function** – An AWS Lambda function
- **aws:memorydb:multi-region-cluster** – An Amazon MemoryDB multi-Region cluster
- **aws:rds:cluster** – An Amazon Aurora DB cluster
- **aws:rds:db** – An Amazon RDS DB instance
- **aws:s3:bucket** – An Amazon S3 bucket

## Identify target resources

When you define a target in the AWS FIS console, you can choose specific AWS resources (of a specific resource type) to target. Or, you can let AWS FIS identify a group of resources based on the criteria that you provide.

To identify your target resources, you can specify the following:

- **Resource IDs** – The resource IDs of specific AWS resources. All resource IDs must represent the same type of resource.
- **Resource tags** – The tags applied to specific AWS resources.
- **Resource filters** – The path and values that represent resources with specific attributes. For more information, see [Resource filters](#).
- **Resource parameters** – The parameters that represent resources that meet specific criteria. For more information, see [Resource parameters](#).

## Considerations

- You can't specify both a resource ID and a resource tag for the same target.
- You can't specify both a resource ID and a resource filter for the same target.

- If you specify a resource tag with an empty tag value, it is not equivalent to a wildcard. It matches resources that have a tag with the specified tag key and an empty tag value.
- If you specify more than one tag, all specified tags have to be present on the target resource for it to be selected (AND).

## Resource filters

Resource filters are queries that identify target resources according to specific attributes. AWS FIS applies the query to the output of an API action that contains the canonical description of the AWS resource, according to the resource type that you specify. Resources that have attributes that match the query are included in the target definition.

Each filter is expressed as an attribute path and possible values. A path is a sequence of elements, separated by periods, that describe the path to reach an attribute in the output of the **Describe** action for a resource. Each period stands for the expansion of an element. Each element must be expressed in Pascal case, even if the output of the **Describe** action for a resource is in camel case. For example, you should use `AvailabilityZone`, not `availablityZone` as an attribute element.

```
"filters": [  
  {  
    "path": "Component.Component.Component",  
    "values": [  
      "string"  
    ]  
  }  
],
```

The following logic applies to all resource filters:

- If multiple filters are provided, including filters with the same path, all filters have to be matched for a resource to be selected – AND
- If multiple values are provided for a single filter, any one value needs to be matched for a resource to be selected – OR
- If multiple values are found at the path location of the describe API call, any one value needs to be matched for a resource to be selected – OR
- To match on tag key/value pairs you should select target resources by tags instead (see above).

The following table includes the API actions and AWS CLI commands that you can use to get the canonical descriptions for each resource type. AWS FIS runs these actions on your behalf to apply the filters that you specify. The corresponding documentation describes the resources that are included in the results by default. For example, the documentation for **DescribeInstances** states that recently terminated instances might appear in the results.

Resource type	API action	AWS CLI command
<b>aws:arc:zonal-shift-managed-resource</b>	ListManagedResources	list-managed-resources
<b>aws:directconnect:virtual-interface</b>	<a href="#">DescribeVirtualInterfaces</a>	<a href="#">describe-virtual-interfaces</a>
<b>aws:ec2:autoscaling-group</b>	<a href="#">DescribeAutoScalingGroups</a>	<a href="#">describe-auto-scaling-groups</a>
<b>aws:ec2:ebs-volume</b>	<a href="#">DescribeVolumes</a>	<a href="#">describe-volumes</a>
<b>aws:ec2:instance</b>	<a href="#">DescribeInstances</a>	<a href="#">describe-instances</a>
<b>aws:ec2:subnet</b>	<a href="#">DescribeSubnets</a>	<a href="#">describe-subnets</a>
<b>aws:ec2:transit-gateway</b>	<a href="#">DescribeTransitGateways</a>	<a href="#">describe-transit-gateways</a>
<b>aws:ec2:vpc-endpoint</b>	<a href="#">DescribeVpcEndpoints</a>	<a href="#">describe-vpc-endpoints</a>
<b>aws:ecs:cluster</b>	<a href="#">DescribeClusters</a>	<a href="#">describe-clusters</a>
<b>aws:ecs:task</b>	<a href="#">DescribeTasks</a>	<a href="#">describe-tasks</a>
<b>aws:eks:cluster</b>	<a href="#">DescribeClusters</a>	<a href="#">describe-clusters</a>
<b>aws:eks:nodegroup</b>	<a href="#">DescribeNodegroup</a>	<a href="#">describe-nodegroup</a>
<b>aws:elasticache:replication group</b>	<a href="#">DescribeReplicationGroups</a>	<a href="#">describe-replication-groups</a>
<b>aws:iam:role</b>	<a href="#">ListRoles</a>	<a href="#">list-roles</a>
<b>aws:kinesis:stream</b>	<a href="#">DescribeStreamSummary</a>	<a href="#">describe-stream-summary</a>

Resource type	API action	AWS CLI command
<b>aws:lambda:function</b>	<a href="#">ListFunctions</a>	<a href="#">list-functions</a>
<b>aws:memorydb:multi-region-cluster</b>	<a href="#">DescribeMultiRegionClusters</a>	<a href="#">describe-multi-region-clusters</a>
<b>aws:rds:cluster</b>	<a href="#">DescribeDBClusters</a>	<a href="#">describe-db-clusters</a>
<b>aws:rds:db</b>	<a href="#">DescribeDBInstances</a>	<a href="#">describe-db-instances</a>
<b>aws:s3:bucket</b>	<a href="#">ListBuckets</a>	<a href="#">list-buckets</a>
<b>aws:dynamodb:global-table</b>	<a href="#">DescribeTable</a>	<a href="#">describe-table</a>
<b>aws:dsq:cluster</b>	<a href="#">GetCluster</a>	<a href="#">get-cluster</a>

For examples, see [the section called “Example filters”](#).

## Resource parameters

Resource parameters identify target resources according to specific criteria.

The following resource type supports parameters.

### **aws:ec2:ebs-volume**

- `availabilityZoneIdentifier` – The code (for example, `us-east-1a`) of the Availability Zone that contains the target volumes.

### **aws:ec2:subnet**

- `availabilityZoneIdentifier` – The code (for example, `us-east-1a`) or AZ ID (for example, `use1-az1`) of the Availability Zone that contains the target subnets.
- `vpc` – The VPC that contains the target subnets. Does not support more than one VPC per account.

### **aws:ecs:task**

- `cluster` – The cluster that contains the target tasks.
- `service` – The service that contains the target tasks.

## aws:eks:pod

- `availabilityZoneIdentifier` – Optional. The Availability Zone that contains the target pods. For example, `us-east-1d`. We determine the Availability Zone of a pod by comparing its hostIP and the CIDR of the cluster subnet.
- `clusterIdentifier` – Required. The name or ARN of the target EKS cluster.
- `namespace` – Required. The Kubernetes namespace of the target pods.
- `selectorType` – Required. The selector type. The possible values are `labelSelector`, `deploymentName`, and `podName`.
- `selectorValue` – Required. The selector value. This value depends on the value of `selectorType`.
- `targetContainerName` – Optional. The name of the target container as defined in the pod spec. The default is the first container defined in each target pod spec.

## aws:lambda:function

- `functionQualifier` – Optional. The version or alias of the function to target. If no qualifier is specified all invocations will be considered for targeting. If an alias with multiple versions is specified, all versions included in the alias will be considered for targeting as long as they are invoked using an ARN containing the alias. If the special alias `$LATEST` is used, invocations to the base function ARN and invocations including `$LATEST` in the ARN will be considered for fault injection. For more information on Lambda versions, see [Manage Lambda function versions](#) in the *AWS Lambda user guide*.

## aws:rds:cluster

- `writerAvailabilityZoneIdentifiers` – Optional. The Availability Zones of the writer of the DB cluster. Possible values are: a comma separated list of Availability Zone identifiers, `all`.

## aws:rds:db

- `availabilityZoneIdentifiers` – Optional. The Availability Zones of the DB instance to be affected. Possible values are: a comma separated list of Availability Zone identifiers, `all`.

## aws:elasticache:replicationgroup

- `availabilityZoneIdentifier` – Required. The code (for example, `us-east-1a`) or AZ ID (for example, `use1-az1`) of the Availability Zone that contains the target nodes.

## Selection mode

You scope the identified resources by specifying a selection mode. AWS FIS supports the following selection modes:

- ALL – Run the action on all targets.
- COUNT(*n*) – Run the action on the specified number of targets, chosen from the identified targets at random. For example, COUNT(1) selects one of the identified targets.
- PERCENT(*n*) – Run the action on the specified percentage of targets, chosen from the identified targets at random. For example, PERCENT(25) selects 25% of the identified targets.

If you have an odd number of resources and specify 50%, AWS FIS rounds down. For example, if you add five Amazon EC2 instances as targets and scope to 50%, AWS FIS rounds down to two instances. You can't specify a percentage that is less than one resource. For example, if you add four Amazon EC2 instances and scope to 5%, AWS FIS can't select an instance.

If you define multiple targets using the same target resource type, AWS FIS can select the same resource multiple times.

Regardless of which selection mode you use, if the scope that you specify identifies no resources, the experiment fails.

## Example targets

The following are example targets.

### Examples

- [Instances in the specified VPC with the specified tags](#)
- [Tasks with the specified parameters](#)

### Example: Instances in the specified VPC with the specified tags

The possible targets for this example are Amazon EC2 instances in the specified VPC with the tag `env=prod`. The selection mode specifies that AWS FIS chooses one of these targets at random.

```
{
```

```

"targets": {
  "randomInstance": {
    "resourceType": "aws:ec2:instance",
    "resourceTags": {
      "env": "prod"
    },
    "filters": [
      {
        "path": "VpcId",
        "values": [
          "vpc-aabbcc11223344556"
        ]
      }
    ],
    "selectionMode": "COUNT(1)"
  }
}

```

### Example: Tasks with the specified parameters

The possible targets for this example are Amazon ECS tasks with the specified cluster and service. The selection mode specifies that AWS FIS choose one of these targets at random.

```

{
  "targets": {
    "randomTask": {
      "resourceType": "aws:ecs:task",
      "parameters": {
        "cluster": "myCluster",
        "service": "myService"
      },
      "selectionMode": "COUNT(1)"
    }
  }
}

```

### Example filters

The following are example filters.

## Examples

- [EC2 instances](#)
- [DB clusters](#)

### Example: EC2 instances

When you specify a filter for an action that supports the **aws:ec2:instance** resource type, AWS FIS uses the Amazon EC2 **describe-instances** command and applies the filter to identify the targets.

The **describe-instances** command returns JSON output where each instance is a structure under `Instances`. The following is partial output that includes fields marked with *italics*. We'll provide examples that use these fields to specify an attribute path from the structure of the JSON output.

```
{
  "Reservations": [
    {
      "Groups": [],
      "Instances": [
        {
          "ImageId": "ami-0011111111111111",
          "InstanceId": "i-00aaaaaaaaaaaaaaaa",
          "InstanceType": "t2.micro",
          "KeyName": "virginia-kp",
          "LaunchTime": "2020-09-30T11:38:17.000Z",
          "Monitoring": {
            "State": "disabled"
          },
          "Placement": {
            "AvailabilityZone": "us-east-1a",
            "GroupName": "",
            "Tenancy": "default"
          },
          "PrivateDnsName": "ip-10-0-1-240.ec2.internal",
          "PrivateIpAddress": "10.0.1.240",
          "ProductCodes": [],
          "PublicDnsName": "ec2-203-0-113-17.compute-1.amazonaws.com",
          "PublicIpAddress": "203.0.113.17",
          "State": {
            "Code": 16,
```

```

        "Name": "running"
    },
    "StateTransitionReason": "",
    "SubnetId": "subnet-aabbcc11223344556",
    "VpcId": "vpc-00bbbbbbbbbbbbbbbb",
    ...
    "NetworkInterfaces": [
    {
        ...
        "Groups": [
            {
                "GroupName": "sec-group-1",
                "GroupId": "sg-a0011223344556677"
            },
            {
                "GroupName": "sec-group-1",
                "GroupId": "sg-b9988776655443322"
            }
        ],
        ...
    },
    ...
],
...
{
    ...
}
],
"OwnerId": "123456789012",
"ReservationId": "r-aaaaabbbbb111111"
},
...
]
}

```

To select instances in a specific Availability Zone using a resource filter, specify the attribute path for `AvailabilityZone` and the code for the Availability Zone as the value. For example:

```

"filters": [
  {
    "path": "Placement.AvailabilityZone",
    "values": [ "us-east-1a" ]
  }
]

```

```
],
```

To select instances in a specific subnet using a resource filter, specify the attribute path for `SubnetId` and the ID of the subnet as the value. For example:

```
"filters": [  
  {  
    "path": "SubnetId",  
    "values": [ "subnet-aabbcc11223344556" ]  
  }  
],
```

To select instances that are in a specific instance state, specify the attribute path for `Name` and one of the following state names as the value: `pending` | `running` | `shutting-down` | `terminated` | `stopping` | `stopped`. For example:

```
"filters": [  
  {  
    "path": "State.Name",  
    "values": [ "running" ]  
  }  
],
```

To select instances that have *any* of a number of security groups attached, specify a single filter with the attribute path for `GroupId` and multiple security group IDs. For example:

```
"filters": [  
  {  
    "path": "NetworkInterfaces.Groups.GroupId",  
    "values": [  
      "sg-a0011223344556677",  
      "sg-f1100110011001100"  
    ]  
  }  
],
```

To select instances that have *all* of a number of security groups attached, specify multiple filters with the attribute path for `GroupId` and a single security group ID for each filter. For example:

```
"filters": [  

```

```

    {
      "path": "NetworkInterfaces.Groups.GroupId",
      "values": [
        "sg-a0011223344556677"
      ]
    },
    {
      "path": "NetworkInterfaces.Groups.GroupId",
      "values": [
        "sg-b9988776655443322"
      ]
    }
  ],
],

```

### Example: Amazon RDS cluster (DB cluster)

When you specify a filter for an action that supports the **aws:rds:cluster** resource type, AWS FIS runs the Amazon RDS **describe-db-clusters** command and applies the filter to identify the targets.

The **describe-db-clusters** command returns JSON output similar to the following for each DB cluster. The following is partial output that includes fields marked with *italics*. We'll provide examples that use these fields to specify an attribute path from the structure of the JSON output.

```

[
  {
    "AllocatedStorage": 1,
    "AvailabilityZones": [
      "us-east-2a",
      "us-east-2b",
      "us-east-2c"
    ],
    "BackupRetentionPeriod": 7,
    "DatabaseName": "",
    "DBClusterIdentifier": "database-1",
    "DBClusterParameterGroup": "default.aurora-postgresql11",
    "DBSubnetGroup": "default-vpc-01234567abc123456",
    "Status": "available",
    "EarliestRestorableTime": "2020-11-13T15:08:32.211Z",
    "Endpoint": "database-1.cluster-example.us-east-2.rds.amazonaws.com",
    "ReaderEndpoint": "database-1.cluster-ro-example.us-east-2.rds.amazonaws.com",
    "MultiAZ": false,
    "Engine": "aurora-postgresql",
  }
]

```

```
        "EngineVersion": "11.7",
        ...
    }
]
```

To apply a resource filter that returns only the DB clusters that use a specific DB engine, specify the attribute path as `Engine` and the value as `aurora-postgresql` as shown in the following example.

```
"filters": [
  {
    "path": "Engine",
    "values": [ "aurora-postgresql" ]
  }
],
```

To apply a resource filter that returns only the DB clusters in a specific Availability Zone, specify the attribute path and value as shown in the following example.

```
"filters": [
  {
    "path": "AvailabilityZones",
    "values": [ "us-east-2a" ]
  }
],
```

## Stop conditions for AWS FIS

AWS Fault Injection Service (AWS FIS) provides controls and guardrails for you to run experiments safely on AWS workloads. A *stop condition* is a mechanism to stop an experiment if it reaches a threshold that you define as an Amazon CloudWatch alarm. If a stop condition is triggered during an experiment, AWS FIS stops the experiment. You cannot resume a stopped experiment.

To create a stop condition, first define the steady state for your application or service. The steady state is when your application is performing optimally, defined in terms of business or technical metrics. For example, latency, CPU load, or number of retries. You can use the steady state to create a CloudWatch alarm that you can use to stop an experiment if your application or service reaches a state where its performance is not acceptable. For more information, see [Using Amazon CloudWatch alarms](#) in the *Amazon CloudWatch User Guide*.

Your account has a quota on the number of stop conditions that you can specify in an experiment template. For more information, see [Quotas and limitations for AWS Fault Injection Service](#).

## Stop condition syntax

When you create an experiment template, you specify one or more stop conditions by specifying the CloudWatch alarms that you created.

```
{
  "stopConditions": [
    {
      "source": "aws:cloudwatch:alarm",
      "value": "arn:aws:cloudwatch:region:123456789012:alarm:alarm-name"
    }
  ]
}
```

The following example indicates that the experiment template does not specify a stop condition.

```
{
  "stopConditions": [
    {
      "source": "none"
    }
  ]
}
```

## Learn more

For a tutorial that demonstrates how to create a CloudWatch alarm and add a stop condition to an experiment template, see [Run CPU stress on an instance](#).

For more information about the CloudWatch metrics that are available for the resource types supported by AWS FIS, see the following:

- [Monitor your instances using CloudWatch](#)
- [Amazon ECS CloudWatch metrics](#)
- [Monitoring Amazon RDS metrics using CloudWatch](#)
- [Monitoring Run Command metrics using CloudWatch](#)

# IAM roles for AWS FIS experiments

AWS Identity and Access Management (IAM) is an AWS service that helps an administrator securely control access to AWS resources. To use AWS FIS, you must create an IAM role that grants AWS FIS the permissions required so that AWS FIS can run experiments on your behalf. You specify this experiment role when you create an experiment template. For a single-account experiment, the IAM policy for the experiment role must grant permission to modify the resources that you specify as targets in your experiment template. For a multi-account experiment, the experiment role must grant the orchestrator role permission to assume the IAM role for each target account. For more information, see [Permissions for multi-account experiments](#).

We recommend that you follow the standard security practice of granting least privilege. You can do so by specifying specific resource ARNs or tags in your policies.

To help you get started with AWS FIS quickly, we provide AWS managed policies that you can specify when you create an experiment role. Alternatively, you can also use these policies as a model as you create your own inline policy documents.

## Contents

- [Prerequisites](#)
- [Option 1: Create an experiment role and attach an AWS managed policy](#)
- [Option 2: Create an experiment role and add an inline policy document](#)

## Prerequisites

Before you begin, install the AWS CLI and create the required trust policy.

### Install the AWS CLI

Before you begin, install and configure the AWS CLI. When you configure the AWS CLI, you are prompted for AWS credentials. The examples in this procedure assume that you also configured a default Region. Otherwise, add the `--region` option to each command. For more information, see [Installing or updating the AWS CLI](#) and [Configuring the AWS CLI](#).

### Create a trust relationship policy

An experiment role must have a trust relationship that allows the AWS FIS service to assume the role. Create a text file named `fis-role-trust-policy.json` and add the following trust relationship policy.

## JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Principal": {
        "Service": [
          "fis.amazonaws.com"
        ]
      },
      "Action": "sts:AssumeRole"
    }
  ]
}
```

We recommend that you use the `aws:SourceAccount` and `aws:SourceArn` condition keys to protect yourself against [the confused deputy problem](#). The source account is the owner of the experiment and the source ARN is the ARN of the experiment. For example, you should add the following condition block to your trust policy.

```
"Condition": {
  "StringEquals": {
    "aws:SourceAccount": "account_id"
  },
  "ArnLike": {
    "aws:SourceArn": "arn:aws:fis:region:account_id:experiment/*"
  }
}
```

### Add permissions to assume target account roles (multi-account experiments only)

For multi-account experiments, you need permissions that allows orchestrator account to assume target account roles. You can modify the following example and add as an inline policy document to assume target account roles:

```
{
  "Effect": "Allow",
  "Action": "sts:AssumeRole",
  "Resource": [
    "arn:aws:iam::target_account_id:role/role_name"
  ]
}
```

```
]
}
```

## Option 1: Create an experiment role and attach an AWS managed policy

Use one of the AWS managed policies from AWS FIS to get started quickly.

### To create an experiment role and attach an AWS managed policy

1. Verify that there is a managed policy for the AWS FIS actions in your experiment. Otherwise, you'll need to create your own inline policy document instead. For more information, see [the section called "AWS managed policies"](#).
2. Use the following [create-role](#) command to create a role and add the trust policy that you created in the prerequisites.

```
aws iam create-role --role-name my-fis-role --assume-role-policy-document
file://fis-role-trust-policy.json
```

3. Use the following [attach-role-policy](#) command to attach the AWS managed policy.

```
aws iam attach-role-policy --role-name my-fis-role --policy-arn fis-policy-arn
```

Where *fis-policy-arn* is one of the following:

- `arn:aws:iam::aws:policy/service-role/AWSFaultInjectionSimulatorEC2Access`
- `arn:aws:iam::aws:policy/service-role/AWSFaultInjectionSimulatorECSAccess`
- `arn:aws:iam::aws:policy/service-role/AWSFaultInjectionSimulatorEKSAccess`
- `arn:aws:iam::aws:policy/service-role/AWSFaultInjectionSimulatorNetworkAccess`
- `arn:aws:iam::aws:policy/service-role/AWSFaultInjectionSimulatorRDSAccess`
- `arn:aws:iam::aws:policy/service-role/AWSFaultInjectionSimulatorSSMAccess`

## Option 2: Create an experiment role and add an inline policy document

Use this option for actions that don't have a managed policy, or to include only the permissions that are required for your specific experiment.

## To create an experiment and add an inline policy document

1. Use the following [create-role](#) command to create a role and add the trust policy that you created in the prerequisites.

```
aws iam create-role --role-name my-fis-role --assume-role-policy-document
file://fis-role-trust-policy.json
```

2. Create a text file named `fis-role-permissions-policy.json` and add a permissions policy. For an example that you can use as a starting point, see the following.

- **Fault injection actions** – Start from the following policy.

JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "AllowFISExperimentRoleFaultInjectionActions",
      "Effect": "Allow",
      "Action": [
        "fis:InjectApiInternalError",
        "fis:InjectApiThrottleError",
        "fis:InjectApiUnavailableError"
      ],
      "Resource": "arn:*:fis:*:*:experiment/*"
    }
  ]
}
```

- **Amazon EBS actions** – Start from the following policy.

JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "ec2:DescribeVolumes"
      ],
    }
  ]
}
```

```

        "Resource": "*"
    },
    {
        "Effect": "Allow",
        "Action": [
            "ec2:PauseVolumeIO"
        ],
        "Resource": "arn:aws:ec2:*:*:volume/*"
    }
]
}

```

- **Amazon EC2 actions** – Start from the [AWSFaultInjectionSimulatorEC2Access](#) policy.
  - **Amazon ECS actions** – Start from the [AWSFaultInjectionSimulatorECSAccess](#) policy.
  - **Amazon EKS actions** – Start from the [AWSFaultInjectionSimulatorEKSAccess](#) policy.
  - **Network actions** – Start from the [AWSFaultInjectionSimulatorNetworkAccess](#) policy.
  - **Amazon RDS actions** – Start from the [AWSFaultInjectionSimulatorRDSAccess](#) policy.
  - **Systems Manager actions** – Start from the [AWSFaultInjectionSimulatorSSMAccess](#) policy.
3. Use the following [put-role-policy](#) command to add the permissions policy that you created in the previous step.

```
aws iam put-role-policy --role-name my-fis-role --policy-name my-fis-policy --policy-document file://fis-role-permissions-policy.json
```

## Experiment report configurations for AWS FIS

You can enable AWS Fault Injection Service (FIS) to generate reports for experiments, making it easier to produce evidence of resilience testing. The experiment report is a PDF document that summarizes experiment actions and optionally captures application response from a CloudWatch dashboard that you specify. To see an example experiment report, download the zip file [here](#).

To enable and configure the contents of the report generated for the experiment, you define the experiment report configuration for the experiment template. When you specify a CloudWatch dashboard, AWS FIS includes a snapshot graph of all the widgets in the given dashboard annotated with experiment start and end time over a duration that you specify, as shown in the example below.

This example demonstrates the impact of a packet loss experiment in an Availability Zone (AZ). When packet loss is introduced in AZ use1-az6, traffic shifts away from use1-az6 and towards use1-az4, so that the number of bytes processed by the load balancer in that AZ declines.

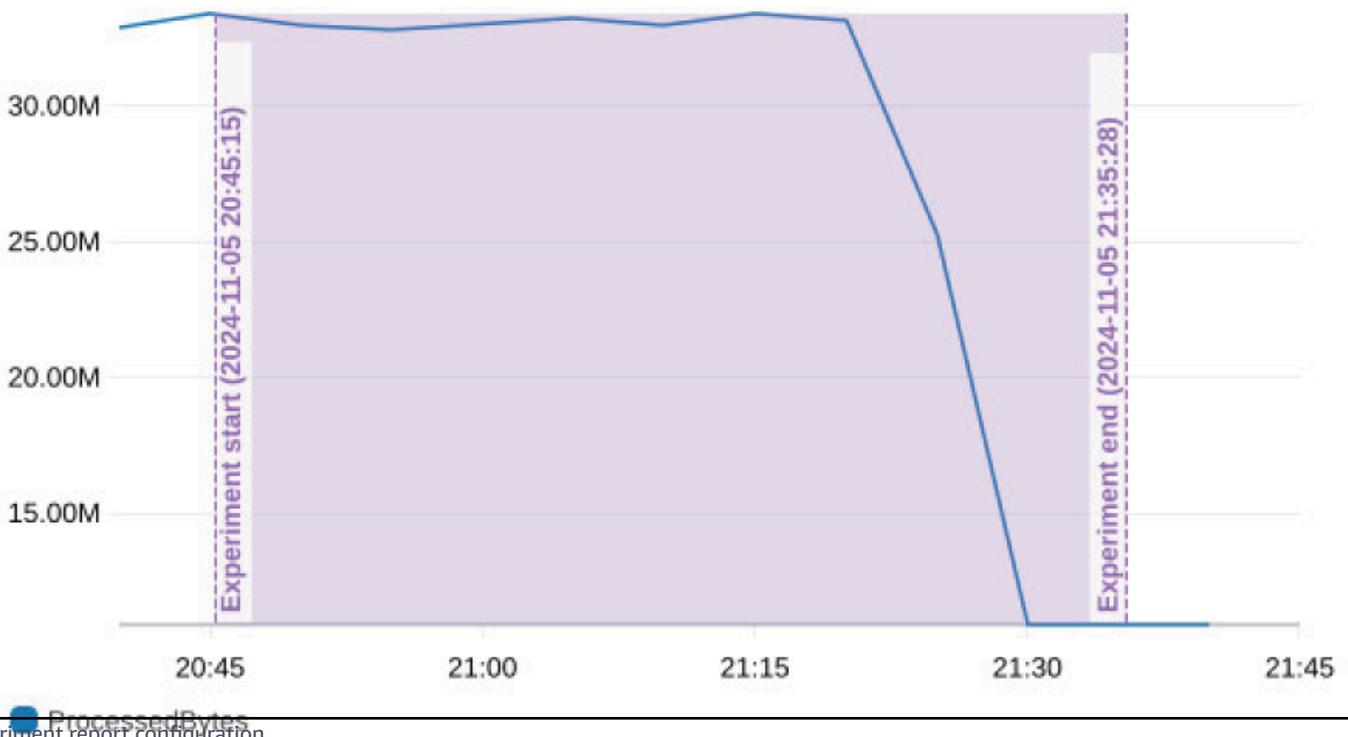
### NLB ProcessedBytes use1-az4

Bytes



### NLB ProcessedBytes use1-az6

Bytes



When the experiment ends, the report can be downloaded from the AWS FIS console and is also stored in an Amazon S3 bucket. If you include a CloudWatch dashboard in your report configuration, images of each widget are also delivered. Reports are not generated for experiments that are cancelled or run as part of target preview (with **actionsMode** set to `skip-all`). Once the experiment exceeds the experiment data retention limit, the report will only be available from the Amazon S3 bucket. AWS FIS charges apply for each delivered report, except for those that fail with internal errors. For more information, see [AWS Fault Injection Service pricing](#) and [Quotas and limitations for AWS Fault Injection Service](#). Ingest and storage charges for Amazon S3 and CloudWatch API charges for **GetMetricWidgetImage** and **GetDashboard** requests may apply. For more information, see [Amazon S3 pricing](#) and [CloudWatch pricing](#).

## Contents

- [Experiment report configuration syntax](#)
- [Experiment report permissions](#)
- [Experiment report best practices](#)

## Experiment report configuration syntax

The following is the syntax for the experiment report configuration, an optional section of the experiment template.

```
{
  "experimentReportConfiguration": {
    "outputs": {
      "s3Configuration": {
        "bucketName": "my-bucket-name",
        "prefix": "report-storage-prefix"
      }
    },
    "dataSources": {
      "cloudWatchDashboards": [
        {
          "dashboardIdentifier": "arn:aws:cloudwatch::123456789012:dashboard/MyDashboard"
        }
      ]
    },
    "preExperimentDuration": "PT20M",
    "postExperimentDuration": "PT20M"
  }
}
```

```
}  
}
```

Using the `experimentReportConfiguration`, you can customize the output destination, input data, and time windows for the data to include in the experiment report, which can help you better understand the impact and results of your AWS FIS experiments. When you define the experiment report configuration, you provide the following:

## outputs

Section of the `experimentReportConfiguration` that specifies where the experiment report will be delivered. In `outputs`, you specify the `s3Configuration` by providing the following:

- `bucketName` - The name of the Amazon S3 bucket where the report will be stored. The bucket must be in the same region as the experiment.
- `prefix` (Optional) - A prefix within the Amazon S3 bucket where the report will be stored. This field is strongly recommended so that you can limit access to the prefix only.

## dataSources

Optional section of the `experimentReportConfiguration` that specifies the additional data sources that will be included in the experiment report.

- `cloudWatchDashboards` - An array of the CloudWatch dashboards that will be included in the report. Limited to one CloudWatch dashboard.
- `dashboardIdentifier` - The ARN of the CloudWatch dashboard. Snapshot graphs of every widget with the type `metric` in this dashboard will be included in the report, with the exception of cross-region metrics.

## preExperimentDuration

Optional section of the `experimentReportConfiguration` that defines the pre-experiment duration for the CloudWatch dashboard metrics to include in the report, up to 30 minutes. This should be a period that represents your application steady state. For example, a pre-experiment duration of 5 minutes means the snapshot graphs will include metrics 5 minutes before the experiment starts. The format for the duration is ISO 8601 and the default is 20 minutes.

## postExperimentDuration

Optional section of the `experimentReportConfiguration` that defines the post-experiment duration for the CloudWatch dashboard metrics to include in the report, up to 2 hours. This

should be a duration that represents your application steady state or recovery period. For example, if you specify a post-experiment duration of 5 minutes, the snapshot graphs will include metrics until 5 minutes after the experiment ends. The format for the duration is ISO 8601 and the default is 20 minutes.

## Experiment report permissions

To enable AWS FIS to generate and store the experiment report, you need to allow the following operations from your AWS FIS experiment IAM role:

- `cloudwatch:GetDashboard`
- `cloudwatch:GetMetricWidgetImage`
- `s3:GetObject`
- `s3:PutObject`

We recommend you follow AWS security best practices and restrict the experiment role to the bucket and prefix. The following is an example policy statement that restricts the experiment role access.

JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Action": [
        "s3:PutObject",
        "s3:GetObject"
      ],
      "Resource": "arn:aws:s3:::my-experiment-report-bucket/my-prefix/*",
      "Effect": "Allow"
    },
    {
      "Action": [
        "cloudwatch:GetDashboard"
      ],

```

```

        "Resource": "arn:aws:cloudwatch::012345678912:dashboard/my-
experiment-report-dashboard",
        "Effect": "Allow"
    },
    {
        "Action": [
            "cloudwatch:GetMetricWidgetImage"
        ],
        "Resource": "*",
        "Effect": "Allow"
    }
]
}

```

## Additional permissions for reports delivered to Amazon S3 buckets encrypted with customer managed keys (CMK)

If the Amazon S3 bucket you specify in the `S3Configuration` is encrypted with CMK, you need to grant the following additional permissions to the FIS experiment role on your KMS key policy:

- `kms:GenerateDataKey`
- `kms:Decrypt`

The following is an example KMS key policy statement that allows the FIS experiment role to write reports to encrypted buckets:

```

{
  "Sid": "Allow FIS experiment report",
  "Effect": "Allow",
  "Principal":
  {
    "AWS": [
      "arn:aws:iam::012345678912:role/FISExperimentRole",
    ]
  },
  "Action": [
    "kms:Decrypt",
    "kms:GenerateDataKey"
  ],
  "Resource": "*"
}

```

```
}
```

## Experiment report best practices

The following are best practices for using the AWS FIS experiment report configuration:

- Before you start an experiment, generate a target preview to verify that your experiment template is configured as you expect. The target preview will give you information about the expected targets of your experiment. To learn more, see [Generate a target preview from an experiment template](#).
- The report should not be used for troubleshooting failed experiments. Instead, use experiment logs to troubleshoot experiment errors. We recommend that you rely on the report only for experiments that you have previously run and successfully completed.
- Restrict the experiment IAM role put and get object access to the S3 destination bucket and prefix. We recommend that you dedicate the bucket / prefix to AWS FIS experiment reports only, and do not grant other AWS services access to this bucket and prefix.
- Use Amazon S3 Object Lock to prevent the report from getting being deleted or overwritten for a fixed amount of time or indefinitely. To learn more, see [Locking objects with Object Lock](#).
- If your CloudWatch dashboard is in a separate account within the same region, you can use CloudWatch cross-account observability to enable your AWS FIS orchestrator account as the monitoring account and the separate account as the source account from the CloudWatch console or Observability Access Manager commands in the AWS CLI and API. To learn more, see [CloudWatch cross-account observability](#).

## Experiment options for AWS FIS

Experiment options are optional settings for an experiment. You can define certain experiment options on the experiment template. Additional experiment options are set when you begin the experiment.

The following is the syntax for experiment options that you define on the experiment template.

```
{
  "experimentOptions": {
    "accountTargeting": "single-account | multi-account",
    "emptyTargetResolutionMode": "fail | skip"
```

```
}  
}
```

If you do not specify any experiment options when you create the experiment template, the default for each option is used.

The following is the syntax for experiment options that you set when you begin the experiment.

```
{  
  "experimentOptions": {  
    "actionsMode": "run-all | skip-all"  
  }  
}
```

If you do not specify any experiment options when you begin the experiment, the default `run-all` is used.

## Contents

- [Account targeting](#)
- [Empty target resolution mode](#)
- [Actions mode](#)

## Account targeting

If you have multiple AWS accounts with resources that you want to target in an experiment, you can define a multi-account experiment using the account targeting experiment option. You run multi-account experiments from an orchestrator account that impacts resources in multiple target accounts. The orchestrator account owns the AWS FIS experiment template and experiment. A target account is an individual AWS account with resources that can be affected by an AWS FIS experiment. For more information, see [Working with multi-account experiments for AWS FIS](#).

You use account targeting to indicate the location of your target resources. You can provide two values for account targeting:

- **single-account** – Default. The experiment will only target resources in the AWS account where AWS FIS experiment runs.
- **multi-account** – The experiment can target resources in multiple AWS accounts.

## Target account configurations

To run a multi-account experiment, you must define one or more target account configurations. A target account configuration specifies the `accountId`, `roleArn`, and `description` for each account with resources targeted in the experiment. The account IDs of the target account configurations for an experiment template must be unique.

When you create a multi-account experiment template, the experiment template will return a read-only field, `targetAccountConfigurationsCount`, that is a count of all the target account configurations for the experiment template.

The following is the syntax for a target account configuration.

```
{
  accountId: "123456789012",
  roleArn: "arn:aws:iam::123456789012:role/AllowFISActions",
  description: "fis-ec2-test"
}
```

When you create a target account configuration, you provide the following:

### **accountId**

12-digit AWS account ID of the target account.

### **roleArn**

An IAM Role granting AWS FIS permissions to take actions in target account.

### **description**

An optional description.

To learn more about how to work with target account configurations, see [Working with multi-account experiments for AWS FIS](#).

## Empty target resolution mode

This mode gives you the option to allow experiments to complete even when a target resource is not resolved.

- **fail** – Default. If no resources are resolved for the target, the experiment is terminated immediately with a status of `failed`.
- **skip** – If no resources are resolved for the target, the experiment will continue and any actions with no resolved targets are skipped. Actions with targets that are defined using unique identifiers, such as ARNs, cannot be skipped. If a target defined using a unique identifier is not found the experiment is terminated immediately with a status of `failed`

## Actions mode

Actions mode is an optional parameter that you can specify when you start an experiment. You can set actions mode to `skip-all` to generate a target preview before injecting faults into your target resources. The target preview allows you to verify the following:

- That you have configured your experiment template to target the resources you expect. The actual resources that are targeted when you start this experiment may be different from the preview because resources may be removed, updated, or sampled randomly.
- That your logging configurations are set up correctly.
- That for multi-account experiments you have correctly set up an IAM role for each of your target account configurations.

### Note

The `skip-all` mode does not allow you to verify that you have the necessary permissions to run the AWS FIS experiment and take actions on your resources.

The actions mode parameter accepts the following values:

- `run-all` - (Default) The experiment will take actions on target resources.
- `skip-all` - The experiment will skip all actions on target resources.

To learn more about how to set the actions mode parameter when you start an experiment, see [Generate a target preview from an experiment template](#).

# AWS FIS Actions reference

An action is the fault injection activity that you run on a target using AWS Fault Injection Service (AWS FIS). AWS FIS provides preconfigured actions for specific types of targets across AWS services. You add actions to an experiment template, which you then use to run experiments.

This reference describes the common actions in AWS FIS, including information about the action parameters and the required IAM permissions. You can also list the supported AWS FIS actions using the AWS FIS console or the [list-actions](#) command from the AWS Command Line Interface (AWS CLI). Once you have the name of a specific action, you can view detailed information about the action by using the [get-action](#) command. For more information on using AWS FIS commands with the AWS CLI, see the [AWS Command Line Interface User Guide](#) and [fis](#) in the *AWS CLI Command Reference*.

For more information on how AWS FIS actions work, see [Actions for AWS FIS](#) and [How AWS Fault Injection Service works with IAM](#).

## Actions

- [Fault injection actions](#)
- [Recovery action](#)
- [Wait action](#)
- [Amazon CloudWatch actions](#)
- [Amazon DynamoDB actions](#)
- [Amazon Aurora DSQL actions](#)
- [Amazon EBS actions](#)
- [Amazon EC2 actions](#)
- [Amazon ECS actions](#)
- [Amazon EKS actions](#)
- [Amazon ElastiCache actions](#)
- [Amazon Kinesis Data Streams actions](#)
- [AWS Lambda actions](#)
- [Amazon MemoryDB action](#)
- [Network actions](#)
- [Amazon RDS actions](#)

- [Amazon S3 actions](#)
- [Systems Manager actions](#)
- [AWS Direct Connect actions](#)
- [Use Systems Manager SSM documents with AWS FIS](#)
- [Use the AWS FIS aws:ecs:task actions](#)
- [Use the AWS FIS aws:eks:pod actions](#)
- [Use the AWS FIS aws:lambda:function actions](#)

## Fault injection actions

AWS FIS supports the following fault injection actions.

### Actions

- [aws:fis:inject-api-internal-error](#)
- [aws:fis:inject-api-throttle-error](#)
- [aws:fis:inject-api-unavailable-error](#)

### aws:fis:inject-api-internal-error

Injects Internal Errors into requests made by the the target IAM role. The specific response depends on each service and API. For more information, please review the SDK and API documentation of your service.

### Resource type

- **aws:iam:role**

### Parameters

- **duration** – The duration, from one minute to 12 hours. In the AWS FIS API, the value is a string in ISO 8601 format. For example, PT1M represents one minute. In the AWS FIS console, you enter the number of seconds, minutes, or hours.
- **service** – The target AWS API namespace. The supported value is `ec2` and `kinesis`.
- **percentage** – The percentage (1-100) of calls to inject the fault into.

- **operations** – The operations to inject the fault into, separated using commas. For a list of the API actions for the ec2 namespace, see [Amazon EC2 API Reference](#) and [Amazon Kinesis Data Streams API Reference](#).

## Permissions

- `fis:InjectApiInternalError`

## aws:fis:inject-api-throttle-error

Injects throttling errors into requests made by the target IAM role. The specific response depends on each service and API. For more information, please review the SDK and API documentation of your service.

## Resource type

- `aws:iam:role`

## Parameters

- **duration** – The duration, from one minute to 12 hours. In the AWS FIS API, the value is a string in ISO 8601 format. For example, PT1M represents one minute. In the AWS FIS console, you enter the number of seconds, minutes, or hours.
- **service** – The target AWS API namespace. The supported value is `ec2` and `kinesis`.
- **percentage** – The percentage (1-100) of calls to inject the fault into.
- **operations** – The operations to inject the fault into, separated using commas. For a list of the API actions for the ec2 namespace, see [Amazon EC2 API Reference](#) and [Amazon Kinesis Data Streams API Reference](#).

## Permissions

- `fis:InjectApiThrottleError`

## aws:fis:inject-api-unavailable-error

Injects Unavailable errors into requests made by the target IAM role. The specific response depends on each service and API. For more information, please review the SDK and API documentation of your service.

### Resource type

- `aws:iam:role`

### Parameters

- **duration** – The duration, from one minute to 12 hours. In the AWS FIS API, the value is a string in ISO 8601 format. For example, PT1M represents one minute. In the AWS FIS console, you enter the number of seconds, minutes, or hours.
- **service** – The target AWS API namespace. The supported value is `ec2` and `kinesis`.
- **percentage** – The percentage (1-100) of calls to inject the fault into.
- **operations** – The operations to inject the fault into, separated using commas. For a list of the API actions for the `ec2` namespace, see [Amazon EC2 API Reference](#) and [Amazon Kinesis Data Streams API Reference](#).

### Permissions

- `fis:InjectApiUnavailableError`

## Recovery action

Recovery actions are performed to mitigate risk or protect applications after impairment.

AWS FIS supports the following recovery actions.

### aws:arc:start-zonal-autoshift

Automatically shifts traffic for supported resources away from a potentially impaired Availability Zone (AZ) and reroutes them to healthy AZs in the same AWS Region. This allows for experiencing zonal autoshift through FIS. Zonal autoshift is a capability in Amazon Application Recovery Controller (ARC) that allows AWS to shift traffic for a resource away from an AZ, on your behalf,

when AWS determines that there is an impairment that could potentially affect customers in the AZ.

When you run the `aws:arc:start-zonal-autoshift` action, AWS FIS manages the zonal shift using the `StartZonalShift`, `UpdateZonalShift`, and `CancelZonalShift` APIs with the `expiresIn` field for these requests set to 1 minute as a safety mechanism. This enables AWS FIS to quickly rollback the zonal shift in the case of any unexpected events such as network outages or system issues. In the ARC console, the expiration time field will display AWS FIS-managed, and the actual expected expiration is determined by the duration specified in the zonal shift action.

## Resource type

- **aws:arc:zonal-shift-managed-resource**

Zonal shift managed resources are resource types including Amazon EKS clusters, Amazon EC2 Application and Network Load Balancers, and Amazon EC2 Auto Scaling groups that can be enabled for ARC zonal autoshift. For more information, see [supported resources](#) and [enabling zonal autoshift resources](#) in the *ARC Developer Guide*.

## Parameters

- **duration** – The length of time for which traffic will be shifted. In the AWS FIS API, the value is a string in ISO 8601 format. For example, PT1M represents one minute. In the AWS FIS console, you enter the number of seconds, minutes, or hours.
- **availabilityZoneIdentifier** – Traffic moves away from this AZ. This can be an AZ name (us-east-1a) or AZ ID (use1-az1).
- **managedResourceTypes** – The resource types from which traffic will be shifted, separated by commas. Possible options are ASG (Auto Scaling Group), ALB (Application Load Balancer), NLB (Network Load Balancer), and EKS (Amazon EKS).
- **zonalAutoshiftStatus** – The `zonalAutoshiftStatus` status of the resources that you want to target. Possible options are ENABLED DISABLED, and ANY. The default is ENABLED.

## Permissions

- `arc-zonal-shift:StartZonalShift`
- `arc-zonal-shift:GetManagedResource`
- `arc-zonal-shift:UpdateZonalShift`

- `arc-zonal-shift:CancelZonalShift`
- `arc-zonal-shift:ListManagedResources`
- `autoscaling:DescribeTags`
- `tag:GetResources`

## Wait action

AWS FIS supports the following wait action.

### **aws:fis:wait**

Runs the AWS FIS wait action.

#### Parameters

- **duration** – The duration, from one minute to 12 hours. In the AWS FIS API, the value is a string in ISO 8601 format. For example, PT1M represents one minute. In the AWS FIS console, you enter the number of seconds, minutes, or hours.

#### Permissions

- None

## Amazon CloudWatch actions

AWS FIS supports the following Amazon CloudWatch action.

### **aws:cloudwatch:assert-alarm-state**

Verifies that the specified alarms are in one of the specified alarm states.

#### Resource type

- None

#### Parameters

- **alarmArns** – The ARNs of the alarms, separated by commas. You can specify up to five alarms.

- **alarmStates** – The alarm states, separated by commas. The possible alarm states are OK, ALARM, and INSUFFICIENT\_DATA.

## Permissions

- `cloudwatch:DescribeAlarms`

## Amazon DynamoDB actions

AWS FIS supports the following Amazon DynamoDB action.

### **aws:dynamodb:global-table-pause-replication**

Pauses Amazon DynamoDB multi-Region global table replication to any replica table. Tables may continue to be replicated for up to 5 minutes after action begins.

### **Multi-Region strongly consistent (MRSC) global tables**

The following statements will be dynamically appended to the policy for the target DynamoDB MRSC global table:

```
{
  "Statement": [
    {
      "Sid": "DoNotModifyFisDynamoDbPauseReplicationEXPxxxxxxxxxxxxxxxx",
      "Effect": "Deny",
      "Principal": {
        "AWS": "*"
      },
      "Action": [
        "dynamodb:UpdateTable"
      ],
      "Resource": "arn:aws:dynamodb:us-east-1:123456789012:table/ExampleGlobalTable",
      "Condition": {
        "DateLessThan": {
          "aws:CurrentTime": "2024-04-10T09:51:41.511Z"
        },
        "ArnEquals": {
          "aws:PrincipalArn": "arn:aws:iam::123456789012:role/aws-service-role/replication.dynamodb.amazonaws.com/AWSServiceRoleForDynamoDBReplication"
        }
      }
    }
  ]
}
```

```

    }
  },
  {
    "Sid":
    "DoNotModifyFisDynamoDbPauseReplicationEXPxxxxxxxxxxxxxxxxxxxxForApplicationAutoScaling",
    "Effect": "Deny",
    "Principal": {
      "AWS": "*"
    },
    "Action": [
      "dynamodb:DescribeTable",
      "dynamodb:UpdateTable"
    ],
    "Resource": "arn:aws:dynamodb:us-east-1:123456789012:table/ExampleGlobalTable",
    "Condition": {
      "DateLessThan": {
        "aws:CurrentTime": "2024-04-10T09:51:41.511Z"
      },
      "ArnEquals": {
        "aws:PrincipalArn": "arn:aws:iam::123456789012:role/
aws-service-role/dynamodb.application-autoscaling.amazonaws.com/
AWSServiceRoleForApplicationAutoScaling_DynamoDBTable"
      }
    }
  }
]
}

```

If a target table does not have any attached resource policies, a resource policy is created for the duration of the experiment, and automatically deleted when the experiment ends. Otherwise, the fault statement is inserted into an existing policy, without any additional modifications to the existing policy statements. The fault statement is then removed from the policy at the end of the experiment.

Target Amazon DynamoDB MRSC global tables are subject to an additional quota. This quota enforces that no single table may be subject to more than 5,040 minutes of impairment in a 7-day rolling window.

### Multi-Region eventually consistent (MREC) global tables

The following statement will be dynamically appended to the policy for the target DynamoDB MREC global table:

```

{
  "Statement": [
    {
      "Sid": "DoNotModifyFisDynamoDbPauseReplicationEXPxxxxxxxxxxxxxxxx",
      "Effect": "Deny",
      "Principal": {
        "AWS": "*"
      },
      "Action": [
        "dynamodb:GetItem",
        "dynamodb:PutItem",
        "dynamodb:UpdateItem",
        "dynamodb>DeleteItem",
        "dynamodb:DescribeTable",
        "dynamodb:UpdateTable",
        "dynamodb:Scan",
        "dynamodb:DescribeTimeToLive",
        "dynamodb:UpdateTimeToLive"
      ],
      "Resource": "arn:aws:dynamodb:us-east-1:123456789012:table/ExampleGlobalTable",
      "Condition": {
        "DateLessThan": {
          "aws:CurrentTime": "2024-04-10T09:51:41.511Z"
        },
        "ArnEquals": {
          "aws:PrincipalArn": "arn:aws:iam::123456789012:role/aws-service-role/replication.dynamodb.amazonaws.com/AWSServiceRoleForDynamoDBReplication"
        }
      }
    }
  ]
}

```

The following statement will be dynamically appended to the stream policy for the target DynamoDB MREC global table:

```

{
  "Statement": [
    {
      "Sid": "DoNotModifyFisDynamoDbPauseReplicationEXPxxxxxxxxxxxxxxxx",
      "Effect": "Deny",
      "Principal": {
        "AWS": "*"
      }
    }
  ]
}

```

```

    },
    "Action": [
      "dynamodb:GetRecords",
      "dynamodb:DescribeStream",
      "dynamodb:GetShardIterator"
    ],
    "Resource": "arn:aws:dynamodb:us-east-1:123456789012:table/ExampleGlobalTable/
stream/2023-08-31T09:50:24.025",
    "Condition": {
      "DateLessThan": {
        "aws:CurrentTime": "2024-04-10T09:51:41.511Z"
      },
      "ArnEquals": {
        "aws:PrincipalArn": "arn:aws:iam::123456789012:role/aws-service-role/
replication.dynamodb.amazonaws.com/AWSServiceRoleForDynamoDBReplication"
      }
    }
  }
]
}

```

If a target table or stream does not have any attached resource policies, a resource policy is created for the duration of the experiment, and automatically deleted when the experiment ends. Otherwise, the fault statement is inserted into an existing policy, without any additional modifications to the existing policy statements. The fault statement is then removed from the policy at the end of the experiment.

## Resource type

- **aws:dynamodb:global-table**

## Parameters

- **duration** – In the AWS FIS API, the value is a string in ISO 8601 format. For example, PT1M represents one minute. In the AWS FIS console, you enter the number of seconds, minutes, or hours.

## Permissions

- **dynamodb:PutResourcePolicy**

- dynamodb:DeleteResourcePolicy
- dynamodb:GetResourcePolicy
- dynamodb:DescribeTable
- tag:GetResources
- dynamodb:InjectError \*

\* The permission is only required if you are targeting MRSC global tables

## Amazon Aurora DSQL actions

AWS FIS supports the following Amazon Aurora DSQL actions.

### aws:dsql:cluster-connection-failure

Creates controlled connection failures in an Aurora DSQL cluster for a specified duration to test application resilience.

#### Resource type

- **aws:dsql:cluster**

#### Parameters

- **duration** – The duration, from one minute to 12 hours. In the AWS FIS API, the value is a string in ISO 8601 format. For example, PT1M represents one minute. In the AWS FIS console, you enter the number of seconds, minutes, or hours.
- **percentage** – The percentage (1-100) of calls to inject the fault into.

#### Permissions

- dsql:InjectError
- dsql:GetCluster
- tag:GetResources

To initiate the experiment with Aurora DSQL, see [Fault injection testing](#) in the *Aurora DSQL User Guide*.

# Amazon EBS actions

AWS FIS supports the following Amazon EBS action.

## Actions

- [aws:ebs:pause-volume-io](#)
- [aws:ebs:volume-io-latency](#)

## aws:ebs:pause-volume-io

Pauses I/O operations on target EBS volumes. The target volumes must be in the same Availability Zone and must be attached to instances built on the Nitro System. The volumes can't be attached to instances on an Outpost.

To initiate the experiment using the Amazon EC2 console, see [Fault testing on Amazon EBS](#) in the *Amazon EC2 User Guide*.

## Resource type

- **aws:ec2:ebs-volume**

## Parameters

- **duration** – The duration, from one second to 12 hours. In the AWS FIS API, the value is a string in ISO 8601 format. For example, PT1M represents one minute, PT5S represents five seconds, and PT6H represents six hours. In the AWS FIS console, you enter the number of seconds, minutes, or hours. If the duration is small, such as PT5S, the I/O is paused for the specified duration, but it might take longer for the experiment to complete due to the time it takes to initialize the experiment.

## Permissions

- `ec2:DescribeVolumes`
- `ec2:PauseVolumeIO`
- `tag:GetResources`

## aws:ebs:volume-io-latency

Injects latency on I/O operations of target EBS volumes. The target volumes must be in the same Availability Zone and must be attached to [Nitro-based instances](#). The volumes cannot be attached to instances in an Outpost.

To initiate the experiment using the Amazon EC2 console, see [Fault testing on Amazon EBS](#) in the *Amazon EBS User Guide*.

### Resource type

- **aws:ec2:ebs-volume**

### Parameters

- **readIOPercentage** – The percentage of read I/O operations that latency will be injected on, from 0.1% to 100.%. This is the percentage of all read I/O operations on the volume that will be impacted during the experiment. The default is 100.
- **readIOLatencyMilliseconds** – The amount of latency injected on read I/O operations in milliseconds, from 1ms (io2 volumes) or 10ms (non-io2 volumes) to 60 seconds. This is the latency value that will be observed on the specified percentage of the read I/O during the experiment. The default is 100.
- **writeIOPercentage** – The percentage of write I/O operations that latency will be injected on, from 0.1% to 100.%. This is the percentage of all write I/O operations on the volume that will be impacted during the experiment. The default is 100.
- **writeIOLatencyMilliseconds** – The amount of latency injected on write I/O operations in milliseconds, from 1ms (io2 volumes) or 10ms (non-io2 volumes) to 60 seconds. This is the latency value that will be observed on the specified percentage of the read I/O during the experiment. The default is 100.
- **duration** – The duration for which the latency will be injected, from 1 second to 12 hours.

### Permissions

- `ec2:DescribeVolumes`
- `ec2:InjectVolumeIOLatency`
- `tag:GetResources`

# Amazon EC2 actions

AWS FIS supports the following Amazon EC2 actions.

## Actions

- [aws:ec2:api-insufficient-instance-capacity-error](#)
- [aws:ec2:asg-insufficient-instance-capacity-error](#)
- [aws:ec2:reboot-instances](#)
- [aws:ec2:send-spot-instance-interruptions](#)
- [aws:ec2:stop-instances](#)
- [aws:ec2:terminate-instances](#)

AWS FIS also supports fault injection actions through the AWS Systems Manager SSM Agent. Systems Manager uses an SSM document that defines actions to perform on EC2 instances. You can use your own document to inject custom faults, or you can use pre-configured SSM documents. For more information, see [the section called "SSM document actions"](#).

## **aws:ec2:api-insufficient-instance-capacity-error**

Injects `InsufficientInstanceCapacity` error responses on requests made by the target IAM roles. Supported operations are `RunInstances`, `CreateCapacityReservation`, `StartInstances`, `CreateFleet` calls. Requests that include capacity asks in multiple Availability Zones are not supported. This action doesn't support defining targets using resource tags, filters, or parameters.

For the Auto Scaling `LaunchInstances` operation, `InsufficientInstanceCapacity` errors will be returned in the response's `errors` field, but the Auto Scaling group's desired capacity will still be updated, allowing the asynchronous scaling process to potentially launch instances. For broader testing of insufficient capacity handling with `LaunchInstances`, consider using this action together with [the section called "aws:ec2:asg-insufficient-instance-capacity-error"](#).

## Resource type

- **aws:iam:role**

## Parameters

- **duration** – In the AWS FIS API, the value is a string in ISO 8601 format. For example, PT1M represents one minute. In the AWS FIS console, you enter the number of seconds, minutes, or hours.
- **availabilityZoneIdentifiers** – The comma separated list of Availability Zones. Supports Zone IDs (e.g. "use1-az1, use1-az2") and Zone names (e.g. "us-east-1a").
- **percentage** – The percentage (1-100) of calls to inject the fault into.

## Permissions

- `ec2:InjectApiError` with condition key `ec2:FisActionId` value set to `aws:ec2:api-insufficient-instance-capacity-error` and `ec2:FisTargetArns` condition key set to target IAM roles.

For an example policy, see [Example: Use condition keys for `ec2:InjectApiError`](#).

## **aws:ec2:asg-insufficient-instance-capacity-error**

Injects `InsufficientInstanceCapacity` error responses on requests made by the target Auto Scaling groups. This action only supports Auto Scaling groups using launch templates. To learn more about insufficient instance capacity errors, see the [Amazon EC2 user guide](#).

## Resource type

- **aws:ec2:autoscaling-group**

## Parameters

- **duration** – In the AWS FIS API, the value is a string in ISO 8601 format. For example, PT1M represents one minute. In the AWS FIS console, you enter the number of seconds, minutes, or hours.
- **availabilityZoneIdentifiers** – The comma separated list of Availability Zones. Supports Zone IDs (e.g. "use1-az1, use1-az2") and Zone names (e.g. "us-east-1a").
- **percentage** – Optional. The percentage (1-100) of the target Auto Scaling group's launch requests to inject the fault. The default is 100.

## Permissions

- `ec2:InjectApiError` with condition key `ec2:FisActionId` value set to `aws:ec2:asg-insufficient-instance-capacity-error` and `ec2:FisTargetArns` condition key set to target Auto Scaling groups.
- `autoscaling:DescribeAutoScalingGroups`

For an example policy, see [Example: Use condition keys for ec2:InjectApiError](#).

## aws:ec2:reboot-instances

Runs the Amazon EC2 API action [RebootInstances](#) on the target EC2 instances.

### Resource type

- `aws:ec2:instance`

### Parameters

- None

### Permissions

- `ec2:RebootInstances`
- `ec2:DescribeInstances`

### AWS managed policy

- [AWSFaultInjectionSimulatorEC2Access](#)

## aws:ec2:send-spot-instance-interruptions

Interrupts the target Spot Instances. Sends a [Spot Instance interruption notice](#) to target Spot Instances two minutes before interrupting them. The interruption time is determined by the specified **durationBeforeInterruption** parameter. Two minutes after the interruption time, the Spot Instances are terminated or stopped, depending on their interruption behavior. A Spot Instance that was stopped by AWS FIS remains stopped until you restart it.

Immediately after the action is initiated, the target instance receives an [EC2 instance rebalance recommendation](#). If you specified **durationBeforeInterruption**, there could be a delay between the rebalance recommendation and the interruption notice.

For more information, see [the section called “Test Spot Instance interruptions”](#). Alternatively, to initiate the experiment by using the Amazon EC2 console, see [Initiate a Spot Instance interruption](#) in the *Amazon EC2 User Guide*.

### Resource type

- **aws:ec2:spot-instance**

### Parameters

- **durationBeforeInterruption** – The time to wait before interrupting the instance, from 2 to 15 minutes. In the AWS FIS API, the value is a string in ISO 8601 format. For example, PT2M represents two minutes. In the AWS FIS console, you enter the number of minutes.

### Permissions

- `ec2:SendSpotInstanceInterruptions`
- `ec2:DescribeInstances`

### AWS managed policy

- [AWSFaultInjectionSimulatorEC2Access](#)

## aws:ec2:stop-instances

Runs the Amazon EC2 API action [StopInstances](#) on the target EC2 instances.

### Resource type

- **aws:ec2:instance**

## Parameters

- **startInstancesAfterDuration** – Optional. The time to wait before starting the instance, from one minute to 12 hours. In the AWS FIS API, the value is a string in ISO 8601 format. For example, PT1M represents one minute. In the AWS FIS console, you enter the number of seconds, minutes, or hours. If the instance has an encrypted EBS volume, you must grant AWS FIS permission to the KMS key used to encrypt the volume, or add the experiment role to the KMS key policy.
- **completeIfInstancesTerminated** – Optional. If true, and if `startInstancesAfterDuration` is also true, this action will not fail when targeted EC2 instances have been terminated by a separate request outside of FIS and cannot be restarted. For example, Auto Scaling groups may terminate stopped EC2 instances under their control before this action completes. The default is false.

## Permissions

- `ec2:StopInstances`
- `ec2:StartInstances`
- `ec2:DescribeInstances` – Optional. Required with **completeIfInstancesTerminated** to validate instance state at end of action.
- `kms:CreateGrant` – Optional. Required with **startInstancesAfterDuration** to restart instances with encrypted volumes.

## AWS managed policy

- [AWSFaultInjectionSimulatorEC2Access](#)

## aws:ec2:terminate-instances

Runs the Amazon EC2 API action [TerminateInstances](#) on the target EC2 instances.

## Resource type

- `aws:ec2:instance`

## Parameters

- None

## Permissions

- `ec2:TerminateInstances`
- `ec2:DescribeInstances`

## AWS managed policy

- [AWSFaultInjectionSimulatorEC2Access](#)

## Amazon ECS actions

AWS FIS supports the following Amazon ECS actions.

### Actions

- [aws:ecs:drain-container-instances](#)
- [aws:ecs:stop-task](#)
- [aws:ecs:task-cpu-stress](#)
- [aws:ecs:task-io-stress](#)
- [aws:ecs:task-kill-process](#)
- [aws:ecs:task-network-blackhole-port](#)
- [aws:ecs:task-network-latency](#)
- [aws:ecs:task-network-packet-loss](#)

## aws:ecs:drain-container-instances

Runs the Amazon ECS API action [UpdateContainerInstancesState](#) to drain the specified percentage of underlying Amazon EC2 instances on the target clusters.

### Resource type

- `aws:ecs:cluster`

### Parameters

- **drainagePercentage** – The percentage (1-100).

- **duration** – The duration, from one minute to 12 hours. In the AWS FIS API, the value is a string in ISO 8601 format. For example, PT1M represents one minute. In the AWS FIS console, you enter the number of seconds, minutes, or hours.

### Permissions

- `ecs:DescribeClusters`
- `ecs:UpdateContainerInstancesState`
- `ecs:ListContainerInstances`
- `tag:GetResources`

### AWS managed policy

- [AWSFaultInjectionSimulatorECSAccess](#)

## aws:ecs:stop-task

Runs the Amazon ECS API action [StopTask](#) to stop the target task.

### Resource type

- `aws:ecs:task`

### Parameters

- None

### Permissions

- `ecs:DescribeTasks`
- `ecs:ListTasks`
- `ecs:StopTask`
- `tag:GetResources`

## AWS managed policy

- [AWSFaultInjectionSimulatorECSAccess](#)

## aws:ecs:task-cpu-stress

Runs CPU stress on the target tasks. Uses the [AWSFIS-Run-CPU-Stress](#) SSM document. The tasks must be managed by AWS Systems Manager. For more information, see [ECS task actions](#).

### Resource type

- **aws:ecs:task**

### Parameters

- **duration** – The duration of the stress test, in ISO 8601 format.
- **percent** – Optional. The target load percentage, from 0 (no load) to 100 (full load). The default is 100.
- **workers** – Optional. The number of stressors to use. The default is 0, which uses all stressors.
- **installDependencies** – Optional. If this value is `True`, Systems Manager installs the required dependencies on the sidecar container for the SSM agent, if they are not already installed. The default is `True`. The dependency is **stress-ng**.

### Permissions

- `ecs:DescribeTasks`
- `ssm:SendCommand`
- `ssm:ListCommands`
- `ssm:CancelCommand`

## aws:ecs:task-io-stress

Runs I/O stress on the target tasks. Uses the [AWSFIS-Run-IO-Stress](#) SSM document. The tasks must be managed by AWS Systems Manager. For more information, see [ECS task actions](#).

## Resource type

- **aws:ecs:task**

## Parameters

- **duration** – The duration of the stress test, in ISO 8601 format.
- **percent** – Optional. The percentage of free space on the file system to use during the stress test. The default is 80%.
- **workers** – Optional. The number of workers. Workers perform a mix of sequential, random, and memory-mapped read/write operations, forced synchronizing, and cache dropping. Multiple child processes perform different I/O operations on the same file. The default is 1.
- **installDependencies** – Optional. If this value is `True`, Systems Manager installs the required dependencies on the sidecar container for the SSM agent, if they are not already installed. The default is `True`. The dependency is **stress-ng**.

## Permissions

- `ecs:DescribeTasks`
- `ssm:SendCommand`
- `ssm:ListCommands`
- `ssm:CancelCommand`

## aws:ecs:task-kill-process

Stops the specified process in the tasks, using the **killall** command. Uses the [AWSFIS-Run-Kill-Process](#) SSM document. The task definition must have `pidMode` set to `task`. The tasks must be managed by AWS Systems Manager. For more information, see [ECS task actions](#).

## Resource type

- **aws:ecs:task**

## Parameters

- **processName** – The name of the process to stop.

- **signal** – Optional. The signal to send along with the command. The possible values are SIGTERM (which the receiver can choose to ignore) and SIGKILL (which cannot be ignored). The default is SIGTERM.
- **installDependencies** – Optional. If this value is `True`, Systems Manager installs the required dependencies on the sidecar container for the SSM agent, if they are not already installed. The default is `True`. The dependency is **killall**.

## Permissions

- `ecs:DescribeTasks`
- `ssm:SendCommand`
- `ssm:ListCommands`
- `ssm:CancelCommand`

## aws:ecs:task-network-blackhole-port

Drops inbound or outbound traffic for the specified protocol and port, using the [Amazon ECS Fault Injection endpoints](#). Uses the [AWSFIS-Run-Network-Blackhole-Port-ECS](#) SSM document. The task definition must have `pidMode` set to `task`. The tasks must be managed by AWS Systems Manager. You can't set `networkMode` to `bridge` in the task definition. For more information, see [ECS task actions](#).

When `useEcsFaultInjectionEndpoints` is set to `false`, the fault uses the `iptables` tool, and uses the [AWSFIS-Run-Network-Blackhole-Port](#) SSM document.

## Resource type

- **aws:ecs:task**

## Parameters

- **duration** – The duration of the test, in ISO 8601 format.
- **port** – The port number.
- **trafficType** – The type of traffic. The possible values are `ingress` and `egress`.
- **protocol** – Optional. The protocol. The possible values are `tcp` and `udp`. The default is `tcp`.

- **installDependencies** – Optional. If this value is `True`, Systems Manager installs the required dependencies on the sidecar container for the SSM agent, if they are not already installed. The default is `True`. The dependencies are **atd**, **curl-minimal**, **dig** and **jq**.
- **useEcsFaultInjectionEndpoints** – Optional. If set to `true`, the Amazon ECS Fault Injection APIs will be used. The default is `false`.

## Permissions

- `ecs:DescribeTasks`
- `ssm:SendCommand`
- `ssm:ListCommands`
- `ssm:CancelCommand`

## aws:ecs:task-network-latency

Adds latency and jitter to the network interface for egress traffic to specific sources, using the [Amazon ECS Fault Injection endpoints](#). Uses the [AWSFIS-Run-Network-Latency-ECS](#) SSM document. The task definition must have `pidMode` set to `task`. The tasks must be managed by AWS Systems Manager. You can't set `networkMode` to `bridge` in the task definition. For more information, see [ECS task actions](#).

When `useEcsFaultInjectionEndpoints` is set to `false`, the fault uses the `tc` tool, and uses the [AWSFIS-Run-Network-Latency-Sources](#) SSM document.

Use the `flowsPercent` parameter to add latency on a percentage of the connections. To use the `flowsPercent` parameter, the ECS Agent version should be `1.100.0` or higher.

To use AZ names or AZ IDs in the `sources` parameter, all targets of the action must be on the same VPC.

## Resource type

- **aws:ecs:task**

## Parameters

- **duration** – The duration of the test, in ISO 8601 format.

- **delayMilliseconds** – Optional. The delay, in milliseconds. The default is 200.
- **jitterMilliseconds** – Optional. The jitter, in milliseconds. The default is 10.
- **flowsPercent** – Optional. The percentage of network flows that will be affected by the action. The default is 100%.
- **sources** – Optional. The sources, separated by commas, without spaces. The possible values are: an IPv4 address, an IPv4 CIDR block, a domain name, an AZ name (us-east-1a), an AZ ID (use1-az1), ALL, DYNAMODB, and S3. If you specify DYNAMODB or S3, this applies only to the Regional endpoint in the current Region. The default is ALL, which matches all IPv4 traffic.
- **installDependencies** – Optional. If this value is `True`, Systems Manager installs the required dependencies on the sidecar container for the SSM agent, if they are not already installed. The default is `True`. The dependencies are **atd**, **curl-minimal**, **dig**, **jq** and **lsof**.
- **useEcsFaultInjectionEndpoints** – Optional. If set to true, the Amazon ECS Fault Injection APIs will be used. The default is false.

## Permissions

- `ecs:DescribeTasks`
- `ecs:DescribeContainerInstances`
- `ec2:DescribeInstances`
- `ec2:DescribeSubnets`
- `ssm:SendCommand`
- `ssm:ListCommands`
- `ssm:CancelCommand`

## aws:ecs:task-network-packet-loss

Adds packet loss to the network interface for egress traffic to specific sources, using the [Amazon ECS Fault Injection endpoints](#). Uses the [AWSFIS-Run-Network-Packet-Loss-ECS](#) SSM document. The task definition must have `pidMode` set to `task`. The tasks must be managed by AWS Systems Manager. You can't set `networkMode` to `bridge` in the task definition. For more information, see [ECS task actions](#).

When `useEcsFaultInjectionEndpoints` is set to `false`, the fault uses the `tc` tool, and uses the [AWSFIS-Run-Network-Packet-Loss-Sources](#) SSM document.

Use the `flowsPercent` parameter to inject packet loss on a percentage of the connections. To use the `flowsPercent` parameter, the ECS Agent version should be 1.100.0 or higher.

To use AZ names or AZ IDs in the `sources` parameter, all targets of the action must be on the same VPC.

## Resource type

- `aws:ecs:task`

## Parameters

- **duration** – The duration of the test, in ISO 8601 format.
- **lossPercent** – Optional. The percentage of packet loss. The default is 7%.
- **flowsPercent** – Optional. The percentage of network flows that will be affected by the action. The default is 100%.
- **sources** – Optional. The sources, separated by commas, without spaces. The possible values are: an IPv4 address, an IPv4 CIDR block, a domain name, an AZ name (us-east-1a), an AZ ID (use1-az1), ALL, DYNAMODB, and S3. If you specify DYNAMODB or S3, this applies only to the Regional endpoint in the current Region. The default is ALL, which matches all IPv4 traffic.
- **installDependencies** – Optional. If this value is `True`, Systems Manager installs the required dependencies on the sidecar container for the SSM agent, if they are not already installed. The default is `True`. The dependencies are **atd**, **curl-minimal**, **dig**, **jq** and **lsof**.
- **useEcsFaultInjectionEndpoints** – Optional. If set to `true`, the Amazon ECS Fault Injection APIs will be used. The default is `false`.

## Permissions

- `ecs:DescribeTasks`
- `ecs:DescribeContainerInstances`
- `ec2:DescribeInstances`
- `ec2:DescribeSubnets`
- `ssm:SendCommand`
- `ssm:ListCommands`
- `ssm:CancelCommand`

# Amazon EKS actions

AWS FIS supports the following Amazon EKS actions.

## Actions

- [aws:eks:inject-kubernetes-custom-resource](#)
- [aws:eks:pod-cpu-stress](#)
- [aws:eks:pod-delete](#)
- [aws:eks:pod-io-stress](#)
- [aws:eks:pod-memory-stress](#)
- [aws:eks:pod-network-blackhole-port](#)
- [aws:eks:pod-network-latency](#)
- [aws:eks:pod-network-packet-loss](#)
- [aws:eks:terminate-nodegroup-instances](#)

## aws:eks:inject-kubernetes-custom-resource

Runs a ChaosMesh or Litmus experiment on a single target cluster. You must install ChaosMesh or Litmus on the target cluster.

When you create an experiment template and define a target of type `aws:eks:cluster`, you must target this action to a single Amazon Resource Name (ARN). This action doesn't support defining targets using resource tags, filters, or parameters.

When you install ChaosMesh, you must specify the appropriate container runtime. Starting with Amazon EKS version 1.23, the default runtime changed from Docker to **containerd**. Starting with version 1.24, Docker was removed.

## Resource type

- **aws:eks:cluster**

## Parameters

- **kubernetesApiVersion** – The API version of the [Kubernetes custom resource](#). The possible values are `chaos-mesh.org/v1alpha1` | `litmuschaos.io/v1alpha1`.

- **kubernetesKind** – The Kubernetes custom resource kind. The value depends on the API version.
  - `chaos-mesh.org/v1alpha1` – The possible values are `AWSChaos` | `DNSChaos` | `GCPChaos` | `HTTPChaos` | `IOChaos` | `JVMChaos` | `KernelChaos` | `NetworkChaos` | `PhysicalMachineChaos` | `PodChaos` | `PodHttpChaos` | `PodIOChaos` | `PodNetworkChaos` | `Schedule` | `StressChaos` | `TimeChaos` |
  - `litmuschaos.io/v1alpha1` – The possible value is `ChaosEngine`.
- **kubernetesNamespace** – The [Kubernetes namespace](#).
- **kubernetesSpec** – The spec section of the Kubernetes custom resource, in JSON format.
- **maxDuration** – The maximum time allowed for the automation execution to complete, from one minute to 12 hours. In the AWS FIS API, the value is a string in ISO 8601 format. For example, `PT1M` represents one minute. In the AWS FIS console, you enter the number of seconds, minutes, or hours.

## Permissions

No AWS Identity and Access Management (IAM) permissions are required for this action. The permissions required to use this action are controlled by Kubernetes using RBAC authorization. For more information, see [Using RBAC Authorization](#) in the official Kubernetes documentation. For more information about Chaos Mesh, see the [official Chaos Mesh documentation](#). For more information about Litmus, see the [official Litmus documentation](#).

## aws:eks:pod-cpu-stress

Runs CPU stress on the target pods. For more information, see [EKS Pod actions](#).

### Resource type

- **aws:eks:pod**

### Parameters

- **duration** – The duration of the stress test, in ISO 8601 format.
- **percent** – Optional. The target load percentage, from 0 (no load) to 100 (full load). The default is 100.
- **workers** – Optional. The number of stressors to use. The default is 0, which uses all stressors.

- **kubernetesServiceAccount** – The Kubernetes service account. For information about the required permissions, see [the section called “Configure the Kubernetes service account”](#).
- **fisPodContainerImage** – Optional. The container image used to create the fault injector pod. The default is to use the images provided by AWS FIS. For more information, see [the section called “Pod container images”](#).
- **maxErrorsPercent** – Optional. The percentage of targets that can fail before the fault injection fails. The default is 0.
- **fisPodLabels** – Optional. The Kubernetes labels that are attached to the fault orchestration pod created by FIS.
- **fisPodAnnotations** – Optional. The Kubernetes annotations that are attached to the fault orchestration pod created by FIS.
- **fisPodSecurityPolicy** – Optional. The [Kubernetes Security Standards](#) policy to use for the fault orchestration pod created by FIS and the ephemeral containers. Possible values are `privileged`, `baseline` and `restricted`. This action is compatible with all policy levels.

## Permissions

- `eks:DescribeCluster`
- `ec2:DescribeSubnets`
- `tag:GetResources`

## AWS managed policy

- [AWSFaultInjectionSimulatorEKSAccess](#)

## **aws:eks:pod-delete**

Deletes the target pods. For more information, see [EKS Pod actions](#).

## Resource type

- **aws:eks:pod**

## Parameters

- **gracePeriodSeconds** – Optional. The duration, in seconds, to wait for the pod to terminate gracefully. If the value is 0, we perform the action immediately. If the value is nil, we use the default grace period for the pod.
- **kubernetesServiceAccount** – The Kubernetes service account. For information about the required permissions, see [the section called “Configure the Kubernetes service account”](#).
- **fisPodContainerImage** – Optional. The container image used to create the fault injector pod. The default is to use the images provided by AWS FIS. For more information, see [the section called “Pod container images”](#).
- **maxErrorsPercent** – Optional. The percentage of targets that can fail before the fault injection fails. The default is 0.
- **fisPodLabels** – Optional. The Kubernetes labels that are attached to the fault orchestration pod created by FIS.
- **fisPodAnnotations** – Optional. The Kubernetes annotations that are attached to the fault orchestration pod created by FIS.
- **fisPodSecurityPolicy** – Optional. The [Kubernetes Security Standards](#) policy to use for the fault orchestration pod created by FIS and the ephemeral containers. Possible values are `privileged`, `baseline` and `restricted`. This action is compatible with all policy levels.

## Permissions

- `eks:DescribeCluster`
- `ec2:DescribeSubnets`
- `tag:GetResources`

## AWS managed policy

- [AWSFaultInjectionSimulatorEKSAccess](#)

## aws:eks:pod-io-stress

Runs I/O stress on the target pods. For more information, see [EKS Pod actions](#).

## Resource type

- **aws:eks:pod**

## Parameters

- **duration** – The duration of the stress test, in ISO 8601 format.
- **workers** – Optional. The number of workers. Workers perform a mix of sequential, random, and memory-mapped read/write operations, forced synchronizing, and cache dropping. Multiple child processes perform different I/O operations on the same file. The default is 1.
- **percent** – Optional. The percentage of free space on the file system to use during the stress test. The default is 80%.
- **kubernetesServiceAccount** – The Kubernetes service account. For information about the required permissions, see [the section called “Configure the Kubernetes service account”](#).
- **fisPodContainerImage** – Optional. The container image used to create the fault injector pod. The default is to use the images provided by AWS FIS. For more information, see [the section called “Pod container images”](#).
- **maxErrorsPercent** – Optional. The percentage of targets that can fail before the fault injection fails. The default is 0.
- **fisPodLabels** – Optional. The Kubernetes labels that are attached to the fault orchestration pod created by FIS.
- **fisPodAnnotations** – Optional. The Kubernetes annotations that are attached to the fault orchestration pod created by FIS.
- **fisPodSecurityPolicy** – Optional. The [Kubernetes Security Standards](#) policy to use for the fault orchestration pod created by FIS and the ephemeral containers. Possible values are `privileged`, `baseline` and `restricted`. This action is compatible with all policy levels.

## Permissions

- `eks:DescribeCluster`
- `ec2:DescribeSubnets`
- `tag:GetResources`

## AWS managed policy

- [AWSFaultInjectionSimulatorEKSAccess](#)

## aws:eks:pod-memory-stress

Runs memory stress on the target pods. For more information, see [EKS Pod actions](#).

### Resource type

- **aws:eks:pod**

### Parameters

- **duration** – The duration of the stress test, in ISO 8601 format.
- **workers** – Optional. The number of stressors to use. The default is 1.
- **percent** – Optional. The percentage of virtual memory to use during the stress test. The default is 80%.
- **kubernetesServiceAccount** – The Kubernetes service account. For information about the required permissions, see [the section called “Configure the Kubernetes service account”](#).
- **fisPodContainerImage** – Optional. The container image used to create the fault injector pod. The default is to use the images provided by AWS FIS. For more information, see [the section called “Pod container images”](#).
- **maxErrorsPercent** – Optional. The percentage of targets that can fail before the fault injection fails. The default is 0.
- **fisPodLabels** – Optional. The Kubernetes labels that are attached to the fault orchestration pod created by FIS.
- **fisPodAnnotations** – Optional. The Kubernetes annotations that are attached to the fault orchestration pod created by FIS.
- **fisPodSecurityPolicy** – Optional. The [Kubernetes Security Standards](#) policy to use for the fault orchestration pod created by FIS and the ephemeral containers. Possible values are `privileged`, `baseline` and `restricted`. This action is compatible with all policy levels.

### Permissions

- `eks:DescribeCluster`

- `ec2:DescribeSubnets`
- `tag:GetResources`

## AWS managed policy

- [AWSFaultInjectionSimulatorEKSAccess](#)

## aws:eks:pod-network-blackhole-port

Drops inbound or outbound traffic for the specified protocol and port. Only compatible with the [Kubernetes Security Standards](#) privilegedpolicy. For more information, see [EKS Pod actions](#).

### Resource type

- `aws:eks:pod`

### Parameters

- **duration** – The duration of the test, in ISO 8601 format.
- **protocol** – The protocol. The possible values are `tcp` and `udp`.
- **trafficType** – The type of traffic. The possible values are `ingress` and `egress`.
- **port** – The port number.
- **kubernetesServiceAccount** – The Kubernetes service account. For information about the required permissions, see [the section called “Configure the Kubernetes service account”](#).
- **fisPodContainerImage** – Optional. The container image used to create the fault injector pod. The default is to use the images provided by AWS FIS. For more information, see [the section called “Pod container images”](#).
- **maxErrorsPercent** – Optional. The percentage of targets that can fail before the fault injection fails. The default is 0.
- **fisPodLabels** – Optional. The Kubernetes labels that are attached to the fault orchestration pod created by FIS.
- **fisPodAnnotations** – Optional. The Kubernetes annotations that are attached to the fault orchestration pod created by FIS.

## Permissions

- `eks:DescribeCluster`
- `ec2:DescribeSubnets`
- `tag:GetResources`

## AWS managed policy

- [AWSFaultInjectionSimulatorEKSAccess](#)

## aws:eks:pod-network-latency

Adds latency and jitter to the network interface using the `tc` tool for traffic to or from specific sources. Only compatible with the [Kubernetes Security Standards](#) privileged policy. For more information, see [EKS Pod actions](#).

Use the `flowsPercent` parameter to add latency on a percentage of the connections.

## Resource type

- `aws:eks:pod`

## Parameters

- **duration** – The duration of the test, in ISO 8601 format.
- **interface** – Optional. The network interfaces, separated by commas. ALL and DEFAULT values are supported. The default is DEFAULT, which will target the primary network interface for the Operating System.
- **delayMilliseconds** – Optional. The delay, in milliseconds. The default is 200.
- **jitterMilliseconds** – Optional. The jitter, in milliseconds. The default is 10.
- **flowsPercent** – Optional. The percentage of network flows that will be affected by the action. The default is 100%.
- **sources** – Optional. The sources, separated by commas, without spaces. The possible values are: an IPv4 address, an IPv4 CIDR block, a domain name, an AZ name (`us-east-1a`), an AZ ID (`use1-az1`), ALL, DYNAMODB, and S3. If you specify DYNAMODB or S3, this applies only to the Regional endpoint in the current Region. For domain names, 10 DNS resolution attempts are made to

collect IP addresses. Due to DNS load balancing and rotation, this action may not impair all possible IP addresses the domain could resolve to. The default is ALL, which matches all IPv4 traffic.

- **kubernetesServiceAccount** – The Kubernetes service account. For information about the required permissions, see [the section called “Configure the Kubernetes service account”](#).
- **fisPodContainerImage** – Optional. The container image used to create the fault injector pod. The default is to use the images provided by AWS FIS. For more information, see [the section called “Pod container images”](#).
- **maxErrorsPercent** – Optional. The percentage of targets that can fail before the fault injection fails. The default is 0.
- **fisPodLabels** – Optional. The Kubernetes labels that are attached to the fault orchestration pod created by FIS.
- **fisPodAnnotations** – Optional. The Kubernetes annotations that are attached to the fault orchestration pod created by FIS.

## Permissions

- `eks:DescribeCluster`
- `ec2:DescribeSubnets`
- `tag:GetResources`

## AWS managed policy

- [AWSFaultInjectionSimulatorEKSAccess](#)

## aws:eks:pod-network-packet-loss

Adds packet loss to the network interface using the `tc` tool. Only compatible with the [Kubernetes Security Standards](#) privileged policy. For more information, see [EKS Pod actions](#).

Use the `flowsPercent` parameter to inject packet loss on a percentage of the connections.

## Resource type

- `aws:eks:pod`

## Parameters

- **duration** – The duration of the test, in ISO 8601 format.
- **interface** – Optional. The network interfaces, separated by commas. ALL and DEFAULT values are supported. The default is DEFAULT, which will target the primary network interface for the Operating System.
- **lossPercent** – Optional. The percentage of packet loss. The default is 7%.
- **flowsPercent** – Optional. The percentage of network flows that will be affected by the action. The default is 100%.
- **sources** – Optional. The sources, separated by commas, without spaces. The possible values are: an IPv4 address, an IPv4 CIDR block, a domain name, an AZ name (us-east-1a), an AZ ID (use1-az1), ALL, DYNAMODB, and S3. If you specify DYNAMODB or S3, this applies only to the Regional endpoint in the current Region. For domain names, 10 DNS resolution attempts are made to collect IP addresses. Due to DNS load balancing and rotation, this action may not impair all possible IP addresses the domain could resolve to. The default is ALL, which matches all IPv4 traffic.
- **kubernetesServiceAccount** – The Kubernetes service account. For information about the required permissions, see [the section called “Configure the Kubernetes service account”](#).
- **fisPodContainerImage** – Optional. The container image used to create the fault injector pod. The default is to use the images provided by AWS FIS. For more information, see [the section called “Pod container images”](#).
- **maxErrorsPercent** – Optional. The percentage of targets that can fail before the fault injection fails. The default is 0.
- **fisPodLabels** – Optional. The Kubernetes labels that are attached to the fault orchestration pod created by FIS.
- **fisPodAnnotations** – Optional. The Kubernetes annotations that are attached to the fault orchestration pod created by FIS.

## Permissions

- `eks:DescribeCluster`
- `ec2:DescribeSubnets`
- `tag:GetResources`

## AWS managed policy

- [AWSFaultInjectionSimulatorEKSAccess](#)

## aws:eks:terminate-nodegroup-instances

Runs the Amazon EC2 API action [TerminateInstances](#) on the target node group. Only compatible with Amazon EKS managed node groups. Self-managed node groups are not supported. For more information, see [EKS manage compute](#).

### Resource type

- `aws:eks:nodegroup`

### Parameters

- `instanceTerminationPercentage` – The percentage (1-100) of instances to terminate.

### Permissions

- `ec2:DescribeInstances`
- `ec2:TerminateInstances`
- `eks:DescribeNodegroup`
- `tag:GetResources`

## AWS managed policy

- [AWSFaultInjectionSimulatorEKSAccess](#)

## Amazon ElastiCache actions

AWS FIS supports the following ElastiCache action.

### aws:elasticache:replicationgroup-interrupt-az-power

Interrupts power to nodes in the specified Availability Zone for target ElastiCache replication groups with Multi-AZ enabled. Only one Availability Zone per replication group can be impacted at

a time. When a primary node is targeted, the corresponding read replica with the least replication lag is promoted to primary. Read replica replacements in the specified Availability Zone are blocked for the duration of this action, which means that target Replication Groups operate with reduced capacity. The target for this action supports both Redis and Valkey engines. The action does not support the "serverless" deployment option.

### Resource type

- `aws:elasticache:replicationgroup`

### Parameters

- **duration** – The duration, from one minute to 12 hours. In the AWS FIS API, the value is a string in ISO 8601 format. For example, PT1M represents one minute. In the AWS FIS console, you enter the number of seconds, minutes, or hours.

### Permissions

- `elasticache:InterruptClusterAzPower`
- `elasticache:DescribeReplicationGroups`
- `tag:GetResources`

#### Note

The ElastiCache interrupt AZ power action now supports all replication group types, including Valkey and Redis. To better represent this functionality, the action has been renamed. If you are currently using `aws:elasticache:interrupt-cluster-az-power`, we recommend that you migrate to the new action `aws:elasticache:replicationgroup-interrupt-az-power` to take advantage of the latest features.

## Amazon Kinesis Data Streams actions

Amazon Kinesis Data Streams supports the following Kinesis actions.

### Actions

- [aws:kinesis:stream-provisioned-throughput-exception](#)
- [aws:kinesis:stream-expired-iterator-exception](#)

## aws:kinesis:stream-provisioned-throughput-exception

Injects ProvisionedThroughputExceededException error responses on requests for the targeted Kinesis Data Streams. Supported operations include: GetRecords, GetShardIterator, PutRecord, and PutRecords.

### Resource type

- **aws:kinesis:stream**

### Parameters

- **duration** – The duration, which ranges from one minute to 12 hours. In the AWS FISAPI, the value is a string in ISO 8601 format. For example, PT1M represents one minute. In the AWS FIS console, you enter the number of seconds, minutes, or hours.
- **percentage** – The percentage (1-100) of calls to inject the fault into.

### Permissions

- `kinesis:InjectApiError`

## aws:kinesis:stream-expired-iterator-exception

Injects ExpiredIteratorException error responses for GetRecords calls targeting specified Kinesis Data Streams.

### Resource type

- **aws:kinesis:stream**

## Parameters

- **duration** – The duration, which ranges from one minute to 12 hours. In the AWS FISAPI, the value is a string in ISO 8601 format. For example, PT1M represents one minute. In the AWS FIS console, you enter the number of seconds, minutes, or hours.
- **percentage** – The percentage (1-100) of calls to inject the fault into.

## Permissions

- `kinesis:InjectApiError`

## AWS Lambda actions

AWS Lambda supports the following Lambda actions

### Actions

- [aws:lambda:invocation-add-delay](#)
- [aws:lambda:invocation-error](#)
- [aws:lambda:invocation-http-integration-response](#)

## aws:lambda:invocation-add-delay

Delays starting a function for a number of milliseconds that you specify. The effect of this action is similar to Lambda cold starts, but the additional time is spent as part of the billed duration and is applied to all execution environments rather than only affecting new execution environments. This means that you may experience both a Lambda cold start and this delay. By setting a latency value higher than the timeout configured on the Lambda function, this action will also provide access to a high fidelity timeout event.

### Resource type

- `aws:lambda:function`

## Parameters

- **duration** – The length of time that the action lasts. In the AWS FIS API, the value is a string in ISO 8601 format. For example, PT1M represents one minute. In the AWS FIS console, you enter the number of seconds, minutes, or hours.
- **invocationPercentage** – Optional. The percentage (1-100) of function invocations to inject the fault into. The default is 100.
- **startupDelayMilliseconds** – Optional. The amount of time in milliseconds (0-900,000) to wait between invocation and execution of function code. The default is 1000.

## Permissions

- s3:PutObject
- s3>DeleteObject
- lambda:GetFunction
- tag:GetResources

## aws:lambda:invocation-error

Marks Lambda function invocations as failed. This action is useful for testing error handling mechanisms, such as alarms and retry configurations. While using this action, you select whether or not to run the function code before returning an error.

## Resource type

- **aws:lambda:function**

## Parameters

- **duration** – The length of time that the action lasts. In the AWS FIS API, the value is a string in ISO 8601 format. For example, PT1M represents one minute. In the AWS FIS console, you enter the number of seconds, minutes, or hours.
- **invocationPercentage** – Optional. The percentage (1-100) of function invocations to inject the fault into. The default is 100.

- **preventExecution** – If the value is true, the action will return the error without executing the function.

## Permissions

- `s3:PutObject`
- `s3:DeleteObject`
- `lambda:GetFunction`
- `tag:GetResources`

## aws:lambda:invocation-http-integration-response

Modifies the behavior of the function. You select a content type and HTTP response code to support integrations with ALB, API-GW and VPC Lattice. To enable selectively impacting upstream or downstream integrations, you can choose whether to directly return the modified response or whether to run the function and replace the response after the function finishes execution.

### Resource type

- **aws:lambda:function**

### Parameters

- **contentTypeHeader** – String value of HTTP content type header to return from Lambda function.
- **duration** – The length of time that the action lasts. In the AWS FIS API, the value is a string in ISO 8601 format. For example, PT1M represents one minute. In the AWS FIS console, you enter the number of seconds, minutes, or hours.
- **invocationPercentage** – Optional. The percentage (1-100) of function invocations to inject the fault into. The default is 100.
- **preventExecution** – If the value is true, the action will return the response without executing the function.
- **statusCode** – Value of HTTP status code (000-999) to return from Lambda function.

## Permissions

- `s3:PutObject`
- `s3>DeleteObject`
- `lambda:GetFunction`
- `tag:GetResources`

## Amazon MemoryDB action

AWS FIS supports the following MemoryDB action.

### **`aws:memorydb:multi-region-cluster-pause-replication`**

Pauses the replication between one regional cluster and all other regional clusters within the multi-Region cluster. The regional cluster targeted is the cluster in the Region where the FIS experiment is running. While the replication is paused, the multi-Region cluster cannot be updated. Once the action completes, it may take a few minutes for the multi-Region cluster to return to an available state. To learn more about Amazon MemoryDB Multi-Region, see the [Amazon MemoryDB Multi-Region Developer Guide](#). For Region availability, see [MemoryDB Multi-Region Prerequisites and limitations](#).

### Resource type

- `aws:memorydb:multi-region-cluster`

### Parameters

- **duration** – The duration, from one minute to 12 hours. In the AWS FIS API, the value is a string in ISO 8601 format. For example, PT1M represents one minute. In the AWS FIS console, you enter the number of seconds, minutes, or hours.

## Permissions

- `memorydb:DescribeMultiRegionClusters`
- `memorydb:PauseMultiRegionClusterReplication`
- `tag:GetResources`

# Network actions

AWS FIS supports the following network actions.

## Actions

- [aws:network:disrupt-connectivity](#)
- [aws:network:route-table-disrupt-cross-region-connectivity](#)
- [aws:network:transit-gateway-disrupt-cross-region-connectivity](#)
- [aws:network:disrupt-vpc-endpoint](#)

## aws:network:disrupt-connectivity

Denies the specified traffic to the target subnets by temporarily cloning the original network access control list (network ACL) associated with the targeted subnet. FIS adds deny rules to the cloned network ACL, which has a tag `managedbyFIS=true`, and associates it with the subnet for the duration of the action. At action completion, FIS deletes the cloned network ACL and restores the original network ACL association.

## Resource type

- **aws:ec2:subnet**

## Parameters

- **scope** – The type of traffic to deny. When the scope is not `all`, the maximum number of entries in network ACLs is 20. The possible values are:
  - `all` – Denies all traffic entering and leaving the subnet. Note that this option allows intra-subnet traffic, including traffic to and from network interfaces in the subnet.
  - `availability-zone` – Denies intra-VPC traffic to and from subnets in other Availability Zones. The maximum number of subnets that can be targeted in a VPC is 30.
  - `dynamodb` – Denies traffic to and from the Regional endpoint for DynamoDB in the current Region.
  - `prefix-list` – Denies traffic to and from the specified prefix list.
  - `s3` – Denies traffic to and from the Regional endpoint for Amazon S3 in the current Region.

- `s3express` – Denies traffic to and from the zonal endpoint for Amazon S3 Express One Zone in the target subnets' AZ. Target subnets must reside in AZs where S3 Express One Zone is currently available. For more information, see [S3 Express One Zone Availability Zones and Regions](#).
- `vpc` – Denies traffic entering and leaving the VPC.
- **duration** – The duration, from one minute to 12 hours. In the AWS FIS API, the value is a string in ISO 8601 format. For example, PT1M represents one minute. In the AWS FIS console, you enter the number of seconds, minutes, or hours.
- **prefixListIdentifier** – If the scope is `prefix-list`, this is the identifier of the customer managed prefix list. You can specify a name, an ID, or an ARN. The prefix list can have at most 10 entries.

## Permissions

- `ec2:CreateNetworkAcl` – Creates the network ACL with the tag `managedByFIS=true`.
- `ec2:CreateNetworkAclEntry` – The network ACL must have the tag `managedByFIS=true`.
- `ec2:CreateTags`
- `ec2>DeleteNetworkAcl` – The network ACL must have the tag `managedByFIS=true`.
- `ec2:DescribeManagedPrefixLists`
- `ec2:DescribeNetworkAcls`
- `ec2:DescribeSubnets`
- `ec2:DescribeVpcs`
- `ec2:GetManagedPrefixListEntries`
- `ec2:ReplaceNetworkAclAssociation`

## AWS managed policy

- [AWSFaultInjectionSimulatorNetworkAccess](#)

## **aws:network:route-table-disrupt-cross-region-connectivity**

Blocks traffic that originates in the target subnets and is destined for the specified Region. Creates route tables that include all routes for the Region to isolate. To allow FIS to create these route

tables, raise the Amazon VPC quota for routes per route table to 250 (or 350 if the `region` parameter is `us-east-1`) plus the number of routes in your existing route tables.

## Resource type

- `aws:ec2:subnet`

## Parameters

- `region` – The code of the Region to isolate (for example, `eu-west-1`).
- `duration` – The length of time the action lasts. In the AWS FIS API, the value is a string in ISO 8601 format. For example, `PT1M` represents one minute. In the AWS FIS console, you enter the number of seconds, minutes, or hours.

## Permissions

- `ec2:AssociateRouteTable`
- `ec2:CreateManagedPrefixList` †
- `ec2:CreateNetworkInterface` †
- `ec2:CreateRoute` †
- `ec2:CreateRouteTable` †
- `ec2:CreateTags` †
- `ec2>DeleteManagedPrefixList` †
- `ec2>DeleteNetworkInterface` †
- `ec2>DeleteRouteTable` †
- `ec2:DescribeManagedPrefixLists`
- `ec2:DescribeNetworkInterfaces`
- `ec2:DescribeRouteTables`
- `ec2:DescribeSubnets`
- `ec2:DescribeVpcPeeringConnections`
- `ec2:DescribeVpcs`
- `ec2:DisassociateRouteTable`

- `ec2:GetManagedPrefixListEntries`
- `ec2:ModifyManagedPrefixList` †
- `ec2:ModifyVpcEndpoint`
- `ec2:ReplaceRouteTableAssociation`

† Scoped using the tag `managedByFIS=true`. You do not need to manage this tag. AWS FIS adds and removes this tag during the experiment.

### AWS managed policy

- [AWSFaultInjectionSimulatorNetworkAccess](#)

## **aws:network:transit-gateway-disrupt-cross-region-connectivity**

Blocks traffic from the target transit gateway peering attachments that is destined for the specified Region.

### Resource type

- **aws:ec2:transit-gateway**

### Parameters

- `region` – The code of the Region to isolate (for example, `eu-west-1`).
- `duration` – The length of time the action lasts. In the AWS FIS API, the value is a string in ISO 8601 format. For example, `PT1M` represents one minute. In the AWS FIS console, you enter the number of seconds, minutes, or hours.

### Permissions

- `ec2:AssociateTransitGatewayRouteTable`
- `ec2:DescribeTransitGatewayAttachments`
- `ec2:DescribeTransitGatewayPeeringAttachments`
- `ec2:DescribeTransitGateways`
- `ec2:DisassociateTransitGatewayRouteTable`

## AWS managed policy

- [AWSFaultInjectionSimulatorNetworkAccess](#)

## aws:network:disrupt-vpc-endpoint

Blocks inbound and outbound traffic of the target interface VPC endpoints. FIS creates a managed security group with empty rules and temporarily replaces security groups of the target VPC endpoints with this managed security group. If modifications are made to the target resources during action execution, the action will fail and the resources will not be restored to their pre-experiment state. Additionally, if a FIS-managed security group is modified during action execution, it will not be deleted by FIS. .

### Resource type

- **aws:ec2:vpc-endpoint**

### Parameters

- **duration** – The length of time the action lasts. In the AWS FIS API, the value is a string in ISO 8601 format. For example, PT1M represents one minute. In the AWS FIS console, you enter the number of seconds, minutes, or hours.

### Permissions

- ec2:DescribeVpcEndpoints
- ec2:DescribeSecurityGroups
- ec2:ModifyVpcEndpoint
- ec2:CreateSecurityGroup
- ec2>DeleteSecurityGroup
- ec2:RevokeSecurityGroupEgress
- ec2:CreateTags
- vpce:AllowMultiRegion \*

\* The permission is only required if you are targeting cross-region VPC endpoints

# Amazon RDS actions

AWS FIS supports the following Amazon RDS actions.

## Actions

- [aws:rds:failover-db-cluster](#)
- [aws:rds:reboot-db-instances](#)

## aws:rds:failover-db-cluster

Runs the Amazon RDS API action [FailoverDBCluster](#) on the target Aurora DB cluster. RDS clusters and DocumentDB clusters are supported.

### Resource type

- **aws:rds:cluster**

### Parameters

- None

### Permissions

- `rds:FailoverDBCluster`
- `rds:DescribeDBClusters`
- `tag:GetResources`

### AWS managed policy

- [AWSFaultInjectionSimulatorRDSAccess](#)

## aws:rds:reboot-db-instances

Runs the Amazon RDS API action [RebootDBInstance](#) on the target DB instance. RDS clusters and DocumentDB clusters are supported.

## Resource type

- **aws:rds:db**

## Parameters

- **forceFailover** – Optional. If the value is true, and if instances are Multi-AZ, forces failover from one Availability Zone to another. The default is false.

## Permissions

- `rds:RebootDBInstance`
- `rds:DescribeDBInstances`
- `tag:GetResources`

## AWS managed policy

- [AWSFaultInjectionSimulatorRDSAccess](#)

# Amazon S3 actions

AWS FIS supports the following Amazon S3 action.

## Actions

- [aws:s3:bucket-pause-replication](#)

## aws:s3:bucket-pause-replication

Pauses replication from target source buckets to destination buckets. Destination buckets can be in different AWS Regions or within the same Region as the source bucket. Existing objects may continue to be replicated for up to one hour after action begins. This action only supports targeting by tags. To learn more about Amazon S3 Replication, see the [Amazon S3 user guide](#).

## Resource type

- **aws:s3:bucket**

## Parameters

- **duration** – The duration, from one minute to 12 hours. In the AWS FIS API, the value is a string in ISO 8601 format. For example, PT1M represents one minute. In the AWS FIS console, you enter the number of seconds, minutes, or hours.
- **region** – The AWS region where destination buckets are located.
- **destinationBuckets** – Optional. Comma separated list of destination S3 bucket(s).
- **prefixes** – Optional. Comma separated list of S3 object key prefixes from replication rule filters. Replication rules of target buckets with a filter based on the prefix(es) will be paused.

## Permissions

- S3:PutReplicationConfiguration with condition key S3:IsReplicationPauseRequest set to True
- S3:GetReplicationConfiguration with condition key S3:IsReplicationPauseRequest set to True
- S3:PauseReplication
- S3:ListAllMyBuckets
- tag:GetResources

For an example policy, see [Example: Use condition keys for aws:s3:bucket-pause-replication](#).

## Systems Manager actions

AWS FIS supports the following Systems Manager actions.

### Actions

- [aws:ssm:send-command](#)
- [aws:ssm:start-automation-execution](#)

## aws:ssm:send-command

Runs the Systems Manager API action [SendCommand](#) on the target EC2 instances. The Systems Manager document (SSM document) defines the actions that Systems Manager performs on your instances. For more information, see [Use the aws:ssm:send-command action](#).

### Resource type

- **aws:ec2:instance**

### Parameters

- **documentArn** – The Amazon Resource Name (ARN) of the document. In the console, this parameter is completed for you if you choose a value from **Action type** that corresponds to one of the [pre-configured AWS FIS SSM documents](#).
- **documentVersion** – Optional. The version of the document. If empty, the default version runs.
- **documentParameters** – Conditional. The required and optional parameters that the document accepts. The format is a JSON object with keys that are strings and values that are either strings or arrays of strings.
- **duration** – The duration, from one minute to 12 hours. In the AWS FIS API, the value is a string in ISO 8601 format. For example, PT1M represents one minute. In the AWS FIS console, you enter the number of seconds, minutes, or hours.

### Permissions

- `ssm:SendCommand`
- `ssm:ListCommands`
- `ssm:CancelCommand`

### AWS managed policy

- [AWSFaultInjectionSimulatorEC2Access](#)

## aws:ssm:start-automation-execution

Runs the Systems Manager API action [StartAutomationExecution](#).

## Resource type

- None

## Parameters

- **documentArn** – The Amazon Resource Name (ARN) of the automation document.
- **documentVersion** – Optional. The version of the document. If empty, the default version runs.
- **documentParameters** – Conditional. The required and optional parameters that the document accepts. The format is a JSON object with keys that are strings and values that are either strings or arrays of strings.
- **maxDuration** – The maximum time allowed for the automation execution to complete, from one minute to 12 hours. In the AWS FIS API, the value is a string in ISO 8601 format. For example, PT1M represents one minute. In the AWS FIS console, you enter the number of seconds, minutes, or hours.

## Permissions

- `ssm:GetAutomationExecution`
- `ssm:StartAutomationExecution`
- `ssm:StopAutomationExecution`
- `iam:PassRole` – Optional. Required if the automation document assumes a role.

## AWS managed policy

- [AWSFaultInjectionSimulatorSSMAccess](#)

## AWS Direct Connect actions

AWS FIS supports the following AWS Direct Connect action.

### Actions

- [aws:directconnect:virtual-interface-disconnect](#)

## aws:directconnect:virtual-interface-disconnect

Tests the resilience of the AWS Direct Connect connection by temporarily disrupting the Border Gateway Protocol (BGP) sessions between the on-premises networks and peers associated with target Virtual Interfaces (VIFs). Before initiating the experiment, FIS verifies that all VIFs targeted in the experiment are in an 'available' state and each VIF has all BGP peers with 'available' state and 'up' BGP status. During the experiment, BGP peering sessions for the targeted Virtual Interfaces will be placed in the down state. For the detailed information about Direct Connect failover testing, please refer to the [AWS Direct Connect documentation](#).

### Resource type

- aws:directconnect:virtual-interface

### Parameters

- `duration` – The duration, from 10 minutes to 12 hours. In the AWS FIS API, the value is a string in ISO 8601 format. For example, PT10M represents ten minutes. In the AWS FIS console, you enter the number of seconds, minutes, or hours.

### Permissions

- `directconnect:DescribeVirtualInterfaces`
- `directconnect:StartBgpFailoverTest`
- `directconnect:ListVirtualInterfaceTestHistory`
- `directconnect:StopBgpFailoverTest`
- `tag:GetResources`

## Use Systems Manager SSM documents with AWS FIS

AWS FIS supports custom fault types through the AWS Systems Manager SSM Agent and the AWS FIS action [aws:ssm:send-command](#). Pre-configured Systems Manager SSM documents (SSM documents) that can be used to create common fault injection actions are available as public AWS documents that begin with the AWSFIS- prefix.

SSM Agent is Amazon software that can be installed and configured on Amazon EC2 instances, on-premises servers, or virtual machines (VMs). This makes it possible for Systems Manager to

manage these resources. The agent processes requests from Systems Manager, and then runs them as specified in the request. You can include your own SSM document to inject custom faults, or reference one of the public Amazon-owned documents.

## Requirements

For actions that require SSM Agent to run the action on the target, you must ensure the following:

- The agent is installed on the target. SSM Agent is installed by default on some Amazon Machine Images (AMIs). Otherwise, you can install the SSM Agent on your instances. For more information, see [Manually install SSM Agent for EC2 instances](#) in the *AWS Systems Manager User Guide*.
- Systems Manager has permission to perform actions on your instances. You grant access using an IAM instance profile. For more information, see [Create an IAM instance profile for Systems Manager](#) and [Attach an IAM instance profile to an EC2 instance](#) in the *AWS Systems Manager User Guide*.

## Use the `aws:ssm:send-command` action

An SSM document defines the actions that Systems Manager performs on your managed instances. Systems Manager includes a number of pre-configured documents, or you can create your own. For more information about creating your own SSM document, see [Creating Systems Manager documents](#) in the *AWS Systems Manager User Guide*. For more information about SSM documents in general, see [AWS Systems Manager documents](#) in the *AWS Systems Manager User Guide*.

AWS FIS provides pre-configured SSM documents. You can view the pre-configured SSM documents under **Documents** in the AWS Systems Manager console: <https://console.aws.amazon.com/systems-manager/documents>. You can also choose from a selection of pre-configured documents in the AWS FIS console. For more information, see [Pre-configured AWS FIS SSM documents](#).

To use an SSM document in your AWS FIS experiments, you can use the [aws:ssm:send-command](#) action. This action fetches and runs the specified SSM document on your target instances.

When you use the `aws:ssm:send-command` action in your experiment template, you must specify additional parameters for the action, including the following:

- **documentArn** – Required. The Amazon Resource Name (ARN) of the SSM document.

- **documentParameters** – Conditional. The required and optional parameters that the SSM document accepts. The format is a JSON object with keys that are strings and values that are either strings or arrays of strings.
- **documentVersion** – Optional. The version of the SSM document to run.

You can view the information for an SSM document (including the parameters for the document) by using the Systems Manager console or the command line.

### To view information about an SSM document using the console

1. Open the AWS Systems Manager console at <https://console.aws.amazon.com/systems-manager/>.
2. In the navigation pane, choose **Documents**.
3. Select the document, and choose the **Details** tab.

### To view information about an SSM document using the command line

Use the SSM [describe-document](#) command.

### Learn more about action state

SSM action state is determined by [SSM command statuses](#).

## Pre-configured AWS FIS SSM documents

You can use pre-configured AWS FIS SSM documents with the `aws:ssm:send-command` action in your experiment templates.

### Requirements

- The pre-configured SSM documents provided by AWS FIS are supported only on the following operating systems:
  - Amazon Linux 2023, Amazon Linux 2
  - Ubuntu
  - RHEL 8, 9
  - CentOS 9
- The pre-configured SSM documents provided by AWS FIS are supported only on EC2 instances. They are not supported on other types of managed nodes, such as on-premises servers.

To use these SSM documents in experiments on ECS tasks, use the corresponding [the section called “Amazon ECS actions”](#). For example, the `aws:ecs:task-cpu-stress` action uses the AWSFIS-Run-CPU-Stress document.

## Documents

- [AWSFIS-Run-CPU-Stress](#)
- [AWSFIS-Run-Disk-Fill](#)
- [AWSFIS-Run-IO-Stress](#)
- [AWSFIS-Run-Kill-Process](#)
- [AWSFIS-Run-Memory-Stress](#)
- [AWSFIS-Run-Network-Blackhole-Port](#)
- [AWSFIS-Run-Network-Latency](#)
- [AWSFIS-Run-Network-Latency-Sources](#)
- [AWSFIS-Run-Network-Packet-Loss](#)
- [AWSFIS-Run-Network-Packet-Loss-Sources](#)

## Difference between action duration and DurationSeconds in AWS FIS SSM documents

Some SSM documents limit their own execution time, for example the `DurationSeconds` parameter is used by some of the pre-configured AWS FIS SSM documents. As a result you need to specify two independent durations in the AWS FIS action definition:

- **Action duration:** For experiments with a single action, the action duration is equivalent to the experiment duration. With multiple actions, the experiment duration depends on individual action durations and the order in which they are run. AWS FIS monitors each action until its action duration elapsed.
- Document parameter **DurationSeconds:** The duration, specified in seconds, for which the SSM document will execute.

You can choose different values for the two types of duration:

- **Action duration exceeds DurationSeconds:** The SSM document execution finishes before the action is complete. AWS FIS waits until the action duration elapsed before subsequent actions are started.

- **Action duration is shorter than DurationSeconds:** The SSM document continues the execution after the action is complete. If the SSM document execution is still in progress and the action duration elapsed then the action status is set to Completed. AWS FIS only monitors the execution until the action duration elapsed.

Note that some SSM documents have variable durations. For example AWS FIS SSM documents have the option to install prerequisites, which can extend the overall execution duration beyond the specified DurationSeconds parameter. Thus, if you set the action duration and DurationSeconds to the same value, it is possible that the SSM script may run longer than the action duration.

## AWSFIS-Run-CPU-Stress

Runs CPU stress on an instance using the **stress-ng** tool. Uses the [AWSFIS-Run-CPU-Stress](#) SSM document.

### Action type (console only)

aws:ssm:send-command/AWSFIS-Run-CPU-Stress

### ARN

arn:aws:ssm:region::document/AWSFIS-Run-CPU-Stress

### Document parameters

- **DurationSeconds** – Required. The duration of the CPU stress test, in seconds.
- **CPU** – Optional. The number of CPU stressors to use. The default is 0, which uses all CPU stressors.
- **LoadPercent** – Optional. The target CPU load percentage, from 0 (no load) to 100 (full load). The default is 100.
- **InstallDependencies** – Optional. If the value is `True`, Systems Manager installs the required dependencies on the target instances if they are not already installed. The default is `True`. The dependency is **stress-ng**.

The following is an example of the string you can enter in the console.

```
{"DurationSeconds":"60", "InstallDependencies":"True"}
```

## AWSFIS-Run-Disk-Fill

Allocates disk space on the root volume of an instance to simulate a disk full fault. Uses the [AWSFIS-Run-Disk-Fill](#) SSM document.

If the experiment injecting this fault is stopped, either manually or through a stop condition, AWS FIS attempts to roll back by canceling the running SSM document. However, if the disk is 100% full, either due to the fault or the fault plus application activity, Systems Manager might be unable to complete the cancel operation. Therefore, if you might need to stop the experiment, ensure that the disk will not become 100% full.

### Action type (console only)

aws:ssm:send-command/AWSFIS-Run-Disk-Fill

### ARN

arn:aws:ssm:*region*::document/AWSFIS-Run-Disk-Fill

### Document parameters

- **DurationSeconds** – Required. The duration of the disk fill test, in seconds.
- **Percent** – Optional. The percentage of the disk to allocate during the disk fill test. The default is 95%.
- **InstallDependencies** – Optional. If the value is `True`, Systems Manager installs the required dependencies on the target instances if they are not already installed. The default is `True`. The dependencies are **atd**, **kmod** and **fallocate**.

The following is an example of the string you can enter in the console.

```
{"DurationSeconds":"60", "InstallDependencies":"True"}
```

## AWSFIS-Run-IO-Stress

Runs IO stress on an instance using the **stress-ng** tool. Uses the [AWSFIS-Run-IO-Stress](#) SSM document.

### Action type (console only)

aws:ssm:send-command/AWSFIS-Run-IO-Stress

## ARN

arn:aws:ssm:region::document/AWSFIS-Run-IO-Stress

### Document parameters

- **DurationSeconds** – Required. The duration of the IO stress test, in seconds.
- **Workers** – Optional. The number of workers that perform a mix of sequential, random, and memory-mapped read/write operations, forced synchronizing, and cache dropping. Multiple child processes perform different I/O operations on the same file. The default is 1.
- **Percent** – Optional. The percentage of free space on the file system to use during the IO stress test. The default is 80%.
- **InstallDependencies** – Optional. If the value is `True`, Systems Manager installs the required dependencies on the target instances if they are not already installed. The default is `True`. The dependency is `stress-ng`.

The following is an example of the string you can enter in the console.

```
{"Workers": "1", "Percent": "80", "DurationSeconds": "60", "InstallDependencies": "True"}
```

## AWSFIS-Run-Kill-Process

Stops the specified process in the instance, using the `killall` command. Uses the [AWSFIS-Run-Kill-Process](#) SSM document.

### Action type (console only)

aws:ssm:send-command/AWSFIS-Run-Kill-Process

## ARN

arn:aws:ssm:region::document/AWSFIS-Run-Kill-Process

### Document parameters

- **ProcessName** – Required. The name of the process to stop.
- **Signal** – Optional. The signal to send along with the command. The possible values are `SIGTERM` (which the receiver can choose to ignore) and `SIGKILL` (which cannot be ignored). The default is `SIGTERM`.

- **InstallDependencies** – Optional. If the value is `True`, Systems Manager installs the required dependencies on the target instances if they are not already installed. The default is `True`. The dependency is `killall`.

The following is an example of the string you can enter in the console.

```
{"ProcessName":"myapplication", "Signal":"SIGTERM"}
```

## AWSFIS-Run-Memory-Stress

Runs memory stress on an instance using the `stress-ng` tool. Uses the [AWSFIS-Run-Memory-Stress](#) SSM document.

### Action type (console only)

`aws:ssm:send-command/AWSFIS-Run-Memory-Stress`

### ARN

`arn:aws:ssm:region::document/AWSFIS-Run-Memory-Stress`

### Document parameters

- **DurationSeconds** – Required. The duration of the memory stress test, in seconds.
- **Workers** – Optional. The number of virtual memory stressors. The default is 1.
- **Percent** – Required. The percentage of virtual memory to use during the memory stress test.
- **InstallDependencies** – Optional. If the value is `True`, Systems Manager installs the required dependencies on the target instances if they are not already installed. The default is `True`. The dependency is `stress-ng`.

The following is an example of the string you can enter in the console.

```
{"Percent":"80", "DurationSeconds":"60", "InstallDependencies":"True"}
```

## AWSFIS-Run-Network-Blackhole-Port

Drops inbound or outbound traffic for the protocol and port using the `iptables` tool. Uses the [AWSFIS-Run-Network-Blackhole-Port](#) SSM document.

## Action type (console only)

aws:ssm:send-command/AWSFIS-Run-Network-Blackhole-Port

## ARN

arn:aws:ssm:*region*::document/AWSFIS-Run-Network-Blackhole-Port

## Document parameters

- **Protocol** – Required. The protocol. The possible values are `tcp` and `udp`.
- **Port** – Required. The port number.
- **TrafficType** – Optional. The type of traffic. The possible values are `ingress` and `egress`. The default is `ingress`.
- **DurationSeconds** – Required. The duration of the network blackhole test, in seconds.
- **InstallDependencies** – Optional. If the value is `True`, Systems Manager installs the required dependencies on the target instances if they are not already installed. The default is `True`. The dependencies are **atd**, **dig**, **lsof**, and **iptables**.

The following is an example of the string you can enter in the console.

```
{"Protocol":"tcp", "Port":"8080", "TrafficType":"egress", "DurationSeconds":"60",  
  "InstallDependencies":"True"}
```

## AWSFIS-Run-Network-Latency

Adds latency to the network interface using the `tc` tool. Uses the [AWSFIS-Run-Network-Latency](#) SSM document.

## Action type (console only)

aws:ssm:send-command/AWSFIS-Run-Network-Latency

## ARN

arn:aws:ssm:*region*::document/AWSFIS-Run-Network-Latency

## Document parameters

- **Interface** – Optional. The network interface. The default is `eth0`.

- **DelayMilliseconds** – Optional. The delay, in milliseconds. The default is 200.
- **DurationSeconds** – Required. The duration of the network latency test, in seconds.
- **InstallDependencies** – Optional. If the value is `True`, Systems Manager installs the required dependencies on the target instances if they are not already installed. The default is `True`. The dependencies are **atd**, **dig**, and **tc**.

The following is an example of the string you can enter in the console.

```
{"DelayMilliseconds":"200", "Interface":"eth0", "DurationSeconds":"60",  
  "InstallDependencies":"True"}
```

## AWSFIS-Run-Network-Latency-Sources

Adds latency and jitter to the network interface using the **tc** tool for traffic to or from specific sources. Uses the [AWSFIS-Run-Network-Latency-Sources](#) SSM document.

Use the `FlowsPercent` parameter to add latency on a percentage of the connections.

### Action type (console only)

`aws:ssm:send-command/AWSFIS-Run-Network-Latency-Sources`

### ARN

`arn:aws:ssm:region::document/AWSFIS-Run-Network-Latency-Sources`

### Document parameters

- **Interface** – Optional. The network interfaces, separated by commas. `ALL` and `DEFAULT` values are supported. The default is `DEFAULT`, which will target the primary network interface for the Operating System.
- **DelayMilliseconds** – Optional. The delay, in milliseconds. The default is 200.
- **JitterMilliseconds** – Optional. The jitter, in milliseconds. The default is 10.
- **FlowsPercent** – Optional. The percentage of network flows that will be affected by the action. The default is 100%.
- **Sources** – Required. The sources, separated by commas, without spaces. The possible values are: an IPv4 address, an IPv4 CIDR block, a domain name, an AZ name (`us-east-1a`), an AZ ID (`use1-`

az1), ALL, DYNAMODB, and S3. If you specify DYNAMODB or S3, this applies only to the Regional endpoint in the current Region.

- **TrafficType** – Optional. The type of traffic. The possible values are `ingress` and `egress`. The default is `ingress`.
- **DurationSeconds** – Required. The duration of the network latency test, in seconds.
- **InstallDependencies** – Optional. If the value is `True`, Systems Manager installs the required dependencies on the target instances if they are not already installed. The default is `True`. The dependencies are **atd**, **dig**, **jq**, **lsof**, and **tc**.

When using this document, the experiment role requires the following permissions:

- `ec2:DescribeInstances`
- `ec2:DescribeSubnets`

The following is an example of the string you can enter in the console.

```
{"DelayMilliseconds":"200", "JitterMilliseconds":"15",  
  "Sources":"S3,www.example.com,72.21.198.67", "Interface":"eth0",  
  "TrafficType":"egress", "DurationSeconds":"60", "InstallDependencies":"True"}
```

## AWSFIS-Run-Network-Packet-Loss

Adds packet loss to the network interface using the **tc** tool. Uses the [AWSFIS-Run-Network-Packet-Loss](#) SSM document.

### Action type (console only)

`aws:ssm:send-command/AWSFIS-Run-Network-Packet-Loss`

### ARN

`arn:aws:ssm:region::document/AWSFIS-Run-Network-Packet-Loss`

### Document parameters

- **Interface** – Optional. The network interface. The default is `eth0`.
- **LossPercent** – Optional. The percentage of packet loss. The default is 7%.
- **DurationSeconds** – Required. The duration of the network packet loss test, in seconds.

- **InstallDependencies** – Optional. If the value is `True`, Systems Manager installs the required dependencies on the target instances. The default is `True`. The dependencies are `atd`, `lsuf`, `dig`, and `tc`.

The following is an example of the string you can enter in the console.

```
{"LossPercent":"15", "Interface":"eth0", "DurationSeconds":"60",  
  "InstallDependencies":"True"}
```

## AWSFIS-Run-Network-Packet-Loss-Sources

Adds packet loss to the network interface using the `tc` tool for traffic to or from specific sources. Uses the [AWSFIS-Run-Network-Packet-Loss-Sources](#) SSM document.

Use the `FlowsPercent` parameter to inject packet loss on a percentage of the connections.

### Action type (console only)

`aws:ssm:send-command/AWSFIS-Run-Network-Packet-Loss-Sources`

### ARN

`arn:aws:ssm:region::document/AWSFIS-Run-Network-Packet-Loss-Sources`

### Document parameters

- **Interface** – Optional. The network interfaces, separated by commas. `ALL` and `DEFAULT` values are supported. The default is `DEFAULT`, which will target the primary network interface for the Operating System.
- **LossPercent** – Optional. The percentage of packet loss. The default is 7%.
- **FlowsPercent** – Optional. The percentage of network flows that will be affected by the action. The default is 100%.
- **Sources** – Required. The sources, separated by commas, without spaces. The possible values are: an IPv4 address, an IPv4 CIDR block, a domain name, an AZ name (`us-east-1a`), an AZ ID (`use1-az1`), `ALL`, `DYNAMODB`, and `S3`. If you specify `DYNAMODB` or `S3`, this applies only to the Regional endpoint in the current Region.
- **TrafficType** – Optional. The type of traffic. The possible values are `ingress` and `egress`. The default is `ingress`.

- **DurationSeconds** – Required. The duration of the network packet loss test, in seconds.
- **InstallDependencies** – Optional. If the value is `True`, Systems Manager installs the required dependencies on the target instances. The default is `True`. The dependencies are **atd**, **dig**, **jq**, **lsuf**, and **tc**.

When using this document, the experiment role requires the following permissions:

- `ec2:DescribeInstances`
- `ec2:DescribeSubnets`

The following is an example of the string you can enter in the console.

```
{"LossPercent": "15", "Sources": "S3, www.example.com, 72.21.198.67", "Interface": "eth0", "TrafficType": "egress", "DurationSeconds": "60", "InstallDependencies": "True"}
```

## Examples

For an example experiment template, see [the section called “Run a pre-configured AWS FIS SSM document”](#).

For an example tutorial, see [Run CPU stress on an instance](#).

## Limitations

- The following documents cannot run in parallel:
  - `AWSFIS-Run-Network-Blackhole-Port`
  - `AWSFIS-Run-Network-Latency`
  - `AWSFIS-Run-Network-Latency-Sources`
  - `AWSFIS-Run-Network-Packet-Loss`
  - `AWSFIS-Run-Network-Packet-Loss-Sources`

## Rollback scripts

AWS FIS SSM documents automatically create rollback scripts as a safety mechanism to restore system state after fault injection experiments. These scripts ensure that injected faults are removed, even if the action fails or is terminated unexpectedly.

## Rollback script creation

Rollback scripts are created automatically when fault injection experiments begin.

### Creation details

- **Location** – Scripts are created in the `/var/lib/amazon/ssm/` directory.
- **Naming pattern** – `FAULT_NAME-FAULT_IDENTIFIER-Rollback.sh` where `FAULT_IDENTIFIER` is a randomly generated 32-character string
- **Timing** – Created at the beginning of each fault injection experiment, before fault injection starts.
- **Content** – Contains all necessary environment variables and commands to reverse the specific fault.

For example, a network latency experiment might create a rollback script at `/var/lib/amazon/ssm/NetworkLatency-abc123-Rollback.sh`.

## Rollback logging

Rollback scripts implement dual logging to capture all rollback activities for troubleshooting and audit purposes.

### Log file locations

When a rollback script executes, it creates logs in two locations:

- **Temporary files** – `/tmp/aws-fis-rollback-TIMESTAMP-PID.log`
- **System logs** – Sent to syslog with facility `local0.info`

### Log file naming

Temporary log files use the following naming convention:

```
/tmp/aws-fis-rollback-YYYY-MM-DDTHH:MM:SSZ-PID.log
```

Where `YYYY-MM-DDTHH:MM:SSZ` is the UTC timestamp and `PID` is the process ID of the rollback script.

## Syslog configuration

Rollback logs are sent to syslog with the following configuration:

- **Tag** – `aws-fis-rollback`
- **Priority** – `local0.info`
- **Format** – `[YYYY-MM-DDTHH:MM:SSZ] log_message`

## To view rollback logs

Use the following command to view all rollback logs from the systemd journal:

```
sudo journalctl -t aws-fis-rollback
```

## Troubleshooting

Use the following procedure to troubleshoot issues.

### To troubleshoot issues with SSM documents

1. Open the AWS Systems Manager console at <https://console.aws.amazon.com/systems-manager/>.
2. In the navigation pane, choose **Node Management, Run Command**.
3. On the **Command history** tab, use the filters to locate the run of the document.
4. Choose the ID of the command to open its details page.
5. Choose the ID of the instance. Review the output and errors for each step.

## Use the AWS FIS `aws:ecs:task` actions

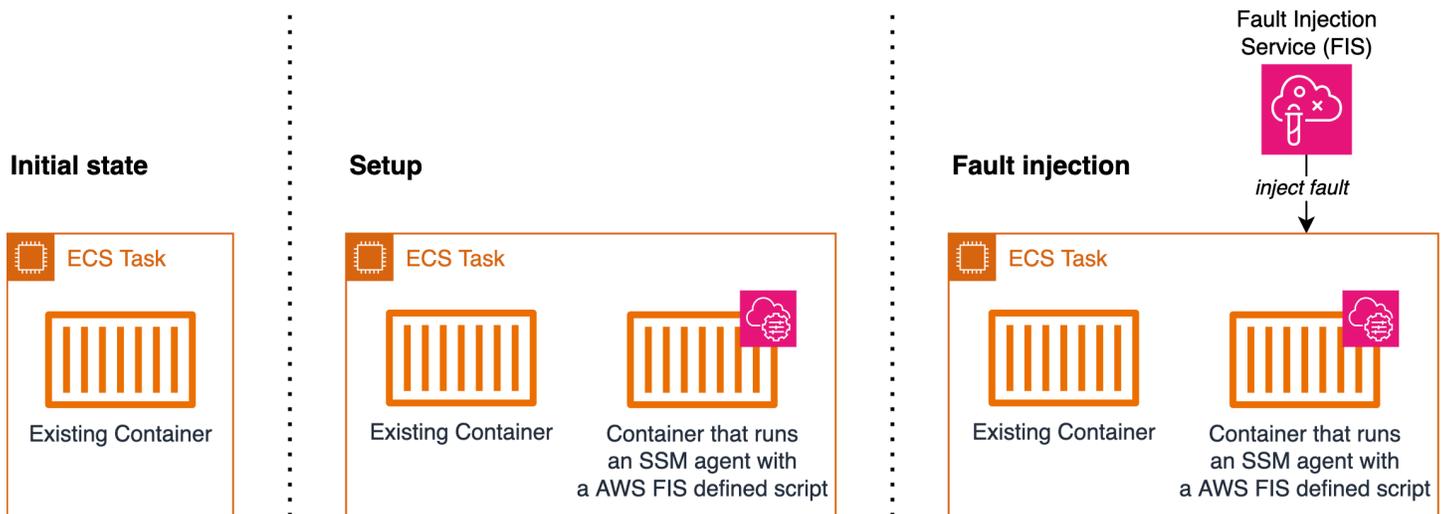
You can use the `aws:ecs:task` actions to inject faults into your Amazon ECS tasks. Amazon EC2 and Fargate capacity types are supported.

These actions use [AWS Systems Manager \(SSM\) documents](#) to inject faults. To use `aws:ecs:task` actions, you will need to add a container with an SSM Agent to your Amazon Elastic Container Service (Amazon ECS) task definition. The container runs an [AWS FIS defined script](#) that registers the Amazon ECS task as Managed Instance in the SSM service. Additionally, the script retrieves task

metadata to add tags to the Managed Instance. The setup will allow AWS FIS to resolve the target task. This paragraph refers to the **Setup** in the diagram below.

When you run an AWS FIS experiment targeting `aws:ecs:task`, AWS FIS maps the target Amazon ECS tasks you specify in an AWS FIS experiment template to a set of SSM managed instances using a resource tag, `ECS_TASK_ARN`. The tag value is the ARN of the associated Amazon ECS task where the SSM documents should be executed. This paragraph refers to the **Fault Injection** in the diagram below.

The following diagram exemplifies the setup and fault injection on a task with one existing container.



## Actions

- [the section called "aws:ecs:task-cpu-stress"](#)
- [the section called "aws:ecs:task-io-stress"](#)
- [the section called "aws:ecs:task-kill-process"](#)
- [the section called "aws:ecs:task-network-blackhole-port"](#)
- [the section called "aws:ecs:task-network-latency"](#)
- [the section called "aws:ecs:task-network-packet-loss"](#)

## Limitations

- The following actions cannot run in parallel:
  - `aws:ecs:task-network-blackhole-port`

- `aws:ecs:task-network-latency`
- `aws:ecs:task-network-packet-loss`
- If you enabled Amazon ECS Exec, you must disable it before you can use these actions.
- The SSM document execution might have Status Cancelled even if the experiment has State Completed. When executing Amazon ECS actions, the customer-provided duration is used both for the action duration in the experiment and the Amazon EC2 Systems Manager (SSM) document duration. After the action is initiated, it takes some time for the SSM document to start running. Consequently, by the time the specified action duration is reached, the SSM document may still have a few seconds remaining to complete its execution. When the experiment action duration is reached, the action is stopped, and the SSM document execution is cancelled. The fault injection was successful.

## Requirements

- Add the following permissions to the AWS FIS [experiment role](#):
  - `ecs:DescribeTasks`
  - `ssm:SendCommand`
  - `ssm:ListCommands`
  - `ssm:CancelCommand`
- Add the following permissions to the Amazon ECS [task IAM role](#):
  - `ssm:CreateActivation`
  - `ssm:AddTagsToResource`
  - `iam:PassRole`

Note that you can specify the ARN of the managed instance role as the resource for `iam:PassRole`.

- Create an Amazon ECS [task execution IAM role](#) and add the [AmazonECSTaskExecutionRolePolicy](#) managed policy.
- In the task definition, set the environment variable `MANAGED_INSTANCE_ROLE_NAME` to the name of the [managed instance role](#). This is the role that will be attached to the tasks registered as managed instances in SSM.
- Add the following permissions to the managed instance role:
  - `ssm>DeleteActivation`

- `ssm:DeregisterManagedInstance`
- Add the [AmazonSSMManagedInstanceCore](#) managed policy to the managed instance role.
- Add an SSM agent container to the Amazon ECS task definition. The command script registers Amazon ECS tasks as managed instances.

```
{
  "name": "amazon-ssm-agent",
  "image": "public.ecr.aws/amazon-ssm-agent/amazon-ssm-agent:latest",
  "cpu": 0,
  "links": [],
  "portMappings": [],
  "essential": false,
  "entryPoint": [],
  "command": [
    "/bin/bash",
    "-c",
    "set -e; dnf upgrade -y; dnf install jq procps awscli -y; term_handler()
    { echo \"Deleting SSM activation $ACTIVATION_ID\"; if ! aws ssm delete-
    activation --activation-id $ACTIVATION_ID --region $ECS_TASK_REGION; then
    echo \"SSM activation $ACTIVATION_ID failed to be deleted\" 1>&2; fi;
    MANAGED_INSTANCE_ID=$(jq -e -r .ManagedInstanceID /var/lib/amazon/ssm/registration);
    echo \"Deregistering SSM Managed Instance $MANAGED_INSTANCE_ID\"; if ! aws
    ssm deregister-managed-instance --instance-id $MANAGED_INSTANCE_ID --region
    $ECS_TASK_REGION; then echo \"SSM Managed Instance $MANAGED_INSTANCE_ID
    failed to be deregistered\" 1>&2; fi; kill -SIGTERM $$SSM_AGENT_PID; }; trap
    term_handler SIGTERM SIGINT; if [[ -z $MANAGED_INSTANCE_ROLE_NAME ]]; then
    echo \"Environment variable MANAGED_INSTANCE_ROLE_NAME not set, exiting\"
    1>&2; exit 1; fi; if ! ps ax | grep amazon-ssm-agent | grep -v grep > /dev/
    null; then if [[ -n $ECS_CONTAINER_METADATA_URI_V4 ]] ; then echo \"Found ECS
    Container Metadata, running activation with metadata\"; TASK_METADATA=$(curl
    \"${ECS_CONTAINER_METADATA_URI_V4}/task\"); ECS_TASK_AVAILABILITY_ZONE=$(echo
    $TASK_METADATA | jq -e -r '.AvailabilityZone'); ECS_TASK_ARN=$(echo $TASK_METADATA
    | jq -e -r '.TaskARN'); ECS_TASK_REGION=$(echo $ECS_TASK_AVAILABILITY_ZONE | sed
    's/.$//'); ECS_TASK_AVAILABILITY_ZONE_REGEX='^(af|ap|ca|cn|eu|me|sa|us|us-gov)-
    (central|north|(north(east|west))|south|south(east|west)|east|west)-[0-9]{1}[a-z]
    {1}$'; if ! [[ $ECS_TASK_AVAILABILITY_ZONE =~ $ECS_TASK_AVAILABILITY_ZONE_REGEX ]];
    then echo \"Error extracting Availability Zone from ECS Container Metadata,
    exiting\" 1>&2; exit 1; fi; ECS_TASK_ARN_REGEX='^arn:(aws|aws-cn|aws-us-gov):ecs:
    [a-z0-9-]+:[0-9]{12}:task/[a-zA-Z0-9-]+/[a-zA-Z0-9]+$'; if ! [[ $ECS_TASK_ARN
    =~ $ECS_TASK_ARN_REGEX ]]; then echo \"Error extracting Task ARN from ECS
    Container Metadata, exiting\" 1>&2; exit 1; fi; CREATE_ACTIVATION_OUTPUT=
    $(aws ssm create-activation --iam-role $MANAGED_INSTANCE_ROLE_NAME --
```

```
tags Key=ECS_TASK_AVAILABILITY_ZONE,Value=$ECS_TASK_AVAILABILITY_ZONE
  Key=ECS_TASK_ARN,Value=$ECS_TASK_ARN Key=FAULT_INJECTION_SIDEDECAR,Value=true --
region $ECS_TASK_REGION); ACTIVATION_CODE=$(echo $CREATE_ACTIVATION_OUTPUT | jq
-e -r .ActivationCode); ACTIVATION_ID=$(echo $CREATE_ACTIVATION_OUTPUT | jq -e
-r .ActivationId); if ! amazon-ssm-agent -register -code $ACTIVATION_CODE -id
$ACTIVATION_ID -region $ECS_TASK_REGION; then echo \"Failed to register with AWS
Systems Manager (SSM), exiting\" 1>&2; exit 1; fi; amazon-ssm-agent & SSM_AGENT_PID=
$!; wait $SSM_AGENT_PID; else echo \"ECS Container Metadata not found, exiting\"
1>&2; exit 1; fi; else echo \"SSM agent is already running, exiting\" 1>&2; exit 1;
fi"
  ],
  "environment": [
    {
      "name": "MANAGED_INSTANCE_ROLE_NAME",
      "value": "SSManagedInstanceRole"
    }
  ],
  "environmentFiles": [],
  "mountPoints": [],
  "volumesFrom": [],
  "secrets": [],
  "dnsServers": [],
  "dnsSearchDomains": [],
  "extraHosts": [],
  "dockerSecurityOptions": [],
  "dockerLabels": {},
  "ulimits": [],
  "logConfiguration": {},
  "systemControls": []
}
```

For a more readable version of the script, see [the section called “Reference version of the script”](#).

- Enable the Amazon ECS Fault Injection APIs, by setting the `enableFaultInjection` field in the Amazon ECS task definition:

```
"enableFaultInjection": true,
```

- When using the `aws:ecs:task-network-blackhole-port`, `aws:ecs:task-network-latency`, or `aws:ecs:task-network-packet-loss` actions on Fargate tasks, the action must have the `useEcsFaultInjectionEndpoints` parameter set to `true`.

- When using the `aws:ecs:task-kill-process`, `aws:ecs:task-network-blackhole-port`, `aws:ecs:task-network-latency`, or `aws:ecs:task-network-packet-loss` actions, the Amazon ECS task definition must have `pidMode` set to `task`.
- When using the `aws:ecs:task-network-blackhole-port`, `aws:ecs:task-network-latency`, or `aws:ecs:task-network-packet-loss` actions on tasks with EC2 launch type, the [networking option in the task definition](#) must be set to `awsvpc` or `host`.

## Reference version of the script

The following is a more readable version of the script in the Requirements section, for your reference.

```
#!/usr/bin/env bash

# This is the activation script used to register ECS tasks as Managed Instances in SSM
# The script retrieves information from the ECS task metadata endpoint to add three
# tags to the Managed Instance
# - ECS_TASK_AVAILABILITY_ZONE: To allow customers to target Managed Instances / Tasks
# in a specific Availability Zone
# - ECS_TASK_ARN: To allow customers to target Managed Instances / Tasks by using the
# Task ARN
# - FAULT_INJECTION_SIDE CAR: To make it clear that the tasks were registered as
# managed instance for fault injection purposes. Value is always 'true'.
# The script will leave the SSM Agent running in the background
# When the container running this script receives a SIGTERM or SIGINT signal, it will
# do the following cleanup:
# - Delete SSM activation
# - Deregister SSM managed instance

set -e # stop execution instantly as a query exits while having a non-zero

dnf upgrade -y
dnf install jq procps awscli -y

term_handler() {
    echo "Deleting SSM activation $ACTIVATION_ID"
    if ! aws ssm delete-activation --activation-id $ACTIVATION_ID --region
$ECS_TASK_REGION; then
        echo "SSM activation $ACTIVATION_ID failed to be deleted" 1>&2
    fi
}
```

```

MANAGED_INSTANCE_ID=$(jq -e -r .ManagedInstanceID /var/lib/amazon/ssm/registration)
echo "Deregistering SSM Managed Instance $MANAGED_INSTANCE_ID"
if ! aws ssm deregister-managed-instance --instance-id $MANAGED_INSTANCE_ID --region
$ECS_TASK_REGION; then
    echo "SSM Managed Instance $MANAGED_INSTANCE_ID failed to be deregistered" 1>&2
fi

kill -SIGTERM $SSM_AGENT_PID
}
trap term_handler SIGTERM SIGINT

# check if the required IAM role is provided
if [[ -z $MANAGED_INSTANCE_ROLE_NAME ]] ; then
    echo "Environment variable MANAGED_INSTANCE_ROLE_NAME not set, exiting" 1>&2
    exit 1
fi

# check if the agent is already running (it will be if ECS Exec is enabled)
if ! ps ax | grep amazon-ssm-agent | grep -v grep > /dev/null; then

# check if ECS Container Metadata is available
if [[ -n $ECS_CONTAINER_METADATA_URI_V4 ]] ; then

    # Retrieve info from ECS task metadata endpoint
    echo "Found ECS Container Metadata, running activation with metadata"
    TASK_METADATA=$(curl "${ECS_CONTAINER_METADATA_URI_V4}/task")
    ECS_TASK_AVAILABILITY_ZONE=$(echo $TASK_METADATA | jq -e -r '.AvailabilityZone')
    ECS_TASK_ARN=$(echo $TASK_METADATA | jq -e -r '.TaskARN')
    ECS_TASK_REGION=$(echo $ECS_TASK_AVAILABILITY_ZONE | sed 's/.$//')

    # validate ECS_TASK_AVAILABILITY_ZONE
    ECS_TASK_AVAILABILITY_ZONE_REGEX='^(af|ap|ca|cn|eu|me|sa|us|us-gov)-(central|north|
(north(east|west))|south|south(east|west)|east|west)-[0-9]{1}[a-z]{1}$'
    if ! [[ $ECS_TASK_AVAILABILITY_ZONE =~ $ECS_TASK_AVAILABILITY_ZONE_REGEX ]] ; then
        echo "Error extracting Availability Zone from ECS Container Metadata, exiting"
1>&2
        exit 1
    fi

    # validate ECS_TASK_ARN
    ECS_TASK_ARN_REGEX='^arn:(aws|aws-cn|aws-us-gov):ecs:[a-z0-9-]+:[0-9]{12}:task/[a-
zA-Z0-9-_-]+/[a-zA-Z0-9-]+$'
    if ! [[ $ECS_TASK_ARN =~ $ECS_TASK_ARN_REGEX ]] ; then
        echo "Error extracting Task ARN from ECS Container Metadata, exiting" 1>&2

```

```

    exit 1
fi

# Create activation tagging with Availability Zone and Task ARN
CREATE_ACTIVATION_OUTPUT=$(aws ssm create-activation \
    --iam-role $MANAGED_INSTANCE_ROLE_NAME \
    --tags Key=ECS_TASK_AVAILABILITY_ZONE,Value=$ECS_TASK_AVAILABILITY_ZONE
Key=ECS_TASK_ARN,Value=$ECS_TASK_ARN Key=FAULT_INJECTION_SIDE CAR,Value=true \
    --region $ECS_TASK_REGION)

ACTIVATION_CODE=$(echo $CREATE_ACTIVATION_OUTPUT | jq -e -r .ActivationCode)
ACTIVATION_ID=$(echo $CREATE_ACTIVATION_OUTPUT | jq -e -r .ActivationId)

# Register with AWS Systems Manager (SSM)
if ! amazon-ssm-agent -register -code $ACTIVATION_CODE -id $ACTIVATION_ID -region
$ECS_TASK_REGION; then
    echo "Failed to register with AWS Systems Manager (SSM), exiting" 1>&2
    exit 1
fi

# the agent needs to run in the background, otherwise the trapped signal
# won't execute the attached function until this process finishes
amazon-ssm-agent &
SSM_AGENT_PID=$!

# need to keep the script alive, otherwise the container will terminate
wait $SSM_AGENT_PID

else
    echo "ECS Container Metadata not found, exiting" 1>&2
    exit 1
fi

else
    echo "SSM agent is already running, exiting" 1>&2
    exit 1
fi

```

## Example experiment template

The following is an example experiment template for the [the section called “aws:ecs:task-cpu-stress”](#) action.

```

{
  "description": "Run CPU stress on the target ECS tasks",
  "targets": {
    "myTasks": {
      "resourceType": "aws:ecs:task",
      "resourceArns": [
        "arn:aws:ecs:us-east-1:111122223333:task/my-
cluster/09821742c0e24250b187dfed8EXAMPLE"
      ],
      "selectionMode": "ALL"
    }
  },
  "actions": {
    "EcsTask-cpu-stress": {
      "actionId": "aws:ecs:task-cpu-stress",
      "parameters": {
        "duration": "PT1M"
      },
      "targets": {
        "Tasks": "myTasks"
      }
    }
  },
  "stopConditions": [
    {
      "source": "none",
    }
  ],
  "roleArn": "arn:aws:iam::111122223333:role/fis-experiment-role",
  "tags": {}
}

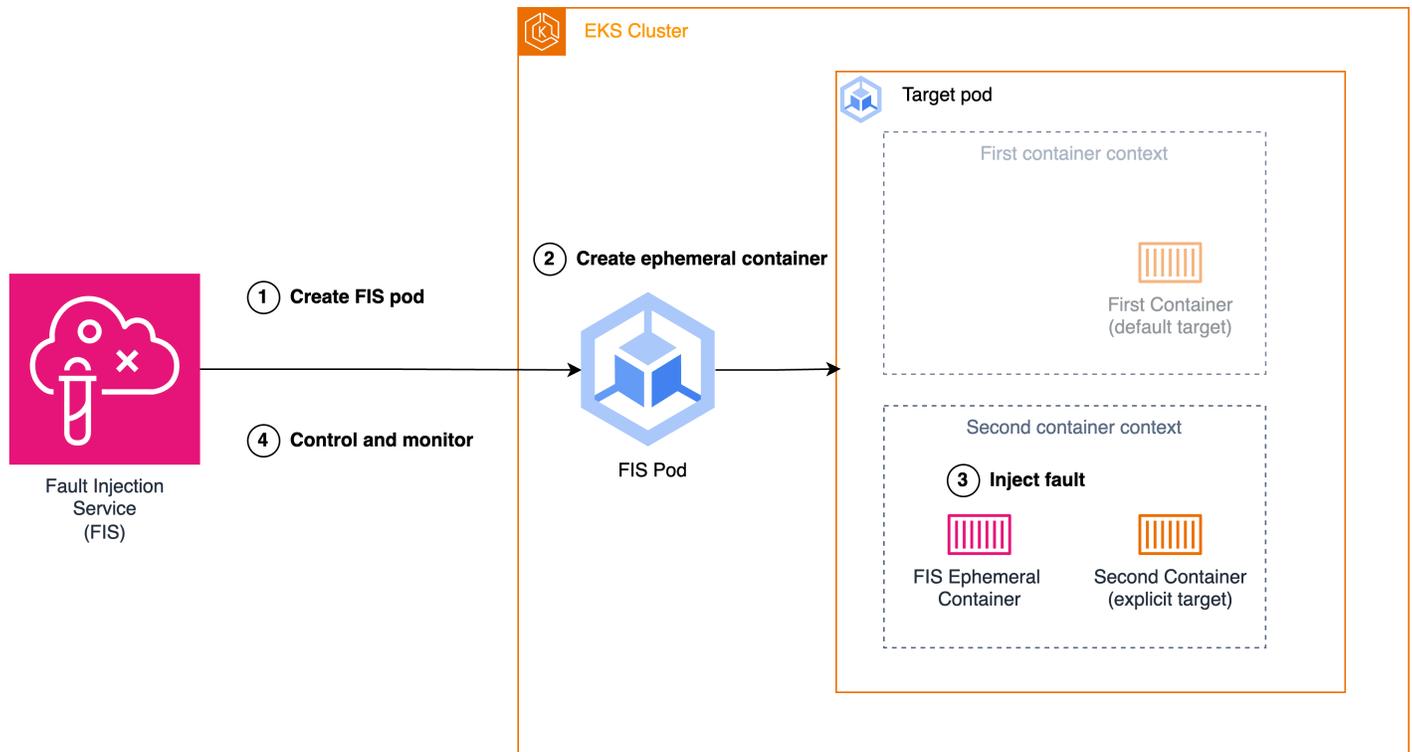
```

## Use the AWS FIS `aws:eks:pod` actions

You can use the `aws:eks:pod` actions to inject faults into the Kubernetes Pods running in your EKS clusters.

When an action is initiated, FIS retrieves the [FIS Pod container image](#). This image is then used to create a Pod in the targeted EKS cluster. The newly-created Pod is responsible for injecting, controlling and monitoring the fault. For all FIS EKS actions, except for [aws:eks:pod-delete](#), the fault injection is achieved through the use of [ephemeral containers](#), a Kubernetes feature that

allows for the creation of temporary containers within an existing Pod. The ephemeral container is started in the same namespace as the target container and executes the desired fault injection tasks. If no target container is specified, the first container in the Pod specification is selected as target.



1. FIS creates the FIS Pod in the target cluster specified in the experiment template.
2. The FIS Pod creates an ephemeral container in the Target Pod in the same namespace as target container.
3. The ephemeral container injects faults in the namespace of the target container.
4. The FIS Pod controls and monitors the fault injection of the ephemeral container and FIS controls and monitors the FIS Pod.

Upon experiment completion or if an error occurs, the ephemeral container and the FIS Pod are removed.

## Actions

- [the section called "aws:eks:pod-cpu-stress"](#)
- [the section called "aws:eks:pod-delete"](#)

- [the section called "aws:eks:pod-io-stress"](#)
- [the section called "aws:eks:pod-memory-stress"](#)
- [the section called "aws:eks:pod-network-blackhole-port"](#)
- [the section called "aws:eks:pod-network-latency"](#)
- [the section called "aws:eks:pod-network-packet-loss"](#)

## Limitations

- The following actions do not work with AWS Fargate:
  - aws:eks:pod-network-blackhole-port
  - aws:eks:pod-network-latency
  - aws:eks:pod-network-packet-loss
- The following actions do not support the bridge [network mode](#):
  - aws:eks:pod-network-blackhole-port
  - aws:eks:pod-network-latency
  - aws:eks:pod-network-packet-loss
- The following actions require root permissions within the ephemeral container.
  - aws:eks:pod-network-blackhole-port
  - aws:eks:pod-network-latency
  - aws:eks:pod-network-packet-loss

The ephemeral container will inherit its permissions from the security context of the target Pod. If you need to run the containers in the Pod as non-root user, you can set separate security contexts for the containers in the target Pod.

- You can't identify targets of type **aws:eks:pod** in your experiment template using resource ARNs or resource tags. You must identify targets using the required resource parameters.
- The actions aws:eks:pod-network-latency and aws:eks:pod-network-packet-loss should not be run in parallel and target the same Pod. Depending on the value of the maxErrors parameter you specify, the action may end in completed or in failed state:
  - If maxErrorsPercent is 0 (default), the action will end in failed state.
  - Otherwise, the failure will add up to the maxErrorsPercent budget. If the number of failed injections do not reach the provided maxErrors, the action will end up in completed state.

- You can identify these failures from the logs of the injected ephemeral container in the target Pod. It will fail with `Exit Code: 16`.
- The action `aws:eks:pod-network-blackhole-port` should not be run in parallel with other actions that target the same Pod and use the same `trafficType`. Parallel actions using different traffic types are supported.
- FIS can only monitor the status of fault injection when the `securityContext` of the target Pods is set to `readOnlyRootFilesystem: false`. Without this configuration, all EKS Pod actions will fail.

## Requirements

- Install the AWS CLI on your computer. This is needed only if you'll use the AWS CLI to create IAM roles. For more information, see [Installing or updating the AWS CLI](#).
- Install **kubectl** on your computer. This is needed only to interact with the EKS cluster to configure or monitor the target application. For more information, see <https://kubernetes.io/docs/tasks/tools/>.
- The minimum supported version of EKS is 1.23.

## Create an experiment role

To run an experiment, you need to configure an IAM role for the experiment. For more information, see [the section called "Experiment role"](#). The required permissions for this role depend on the action you're using. Refer to the [AWS FIS actions that target `aws:eks:pod`](#) to find the necessary permissions for your action.

## Configure the Kubernetes service account

Configure a Kubernetes service account to run experiments with targets in the specified Kubernetes namespace. In the following example, the service account is *myserviceaccount* and the namespace is *default*. Note that *default* is one of the standard Kubernetes namespaces.

### To configure your Kubernetes service account

1. Create a file named `rbac.yaml` and add the following.

```
kind: ServiceAccount
```

```
apiVersion: v1
metadata:
  namespace: default
  name: myserviceaccount
---
kind: Role
apiVersion: rbac.authorization.k8s.io/v1
metadata:
  namespace: default
  name: role-experiments
rules:
- apiGroups: [""]
  resources: ["configmaps"]
  verbs: [ "get", "create", "patch", "delete"]
- apiGroups: [""]
  resources: ["pods"]
  verbs: ["create", "list", "get", "delete", "deletecollection"]
- apiGroups: [""]
  resources: ["pods/ephemeralcontainers"]
  verbs: ["update"]
- apiGroups: [""]
  resources: ["pods/exec"]
  verbs: ["create"]
- apiGroups: ["apps"]
  resources: ["deployments"]
  verbs: ["get"]
---
apiVersion: rbac.authorization.k8s.io/v1
kind: RoleBinding
metadata:
  name: bind-role-experiments
  namespace: default
subjects:
- kind: ServiceAccount
  name: myserviceaccount
  namespace: default
- apiGroup: rbac.authorization.k8s.io
  kind: User
  name: fis-experiment
roleRef:
  kind: Role
  name: role-experiments
```

```
apiGroup: rbac.authorization.k8s.io
```

2. Run the following command.

```
kubectl apply -f rbac.yaml
```

## Grant IAM users and roles access to Kubernetes APIs

Follow the steps explained in [Associate IAM Identities with Kubernetes Permissions](#) in the **EKS** documentation.

### Option 1: Create access entries

We recommend using **Access Entries**. You can use the following command to create an Access Entry which associates the IAM role with the Kubernetes user *fis-experiment*. For more information, see [Grant IAM users access to Kubernetes with EKS access entries](#).

```
aws eks create-access-entry \  
    --principal-arn arn:aws:iam::123456789012:role/fis-experiment-role \  
    --username fis-experiment \  
    --cluster-name my-cluster
```

#### Important

In order to leverage access entries, the authentication mode of the EKS cluster must be configured to either the `API_AND_CONFIG_MAP` or `API` mode.

### Option 2: Add entries to the aws-auth ConfigMap

You can also use the following command to create an identity mapping. For more information, see [Manage IAM users and roles](#) in the **eksctl** documentation.

```
eksctl create iamidentitymapping \  
    --arn arn:aws:iam::123456789012:role/fis-experiment-role \  
    --username fis-experiment \  
    --cluster my-cluster
```

**⚠ Important**

Leveraging the eksctl toolkit to configure identity mappings will result in the creation of entries within the aws-auth ConfigMap. It is important to note that these generated entries do not support the inclusion of a path component. Consequently, the ARN provided as input must not contain a path segment (e.g., `arn:aws:iam::123456789012:role/service-role/fis-experiment-role`).

## Pod container images

The Pod container images provided by AWS FIS are hosted in Amazon ECR. When you reference an image from Amazon ECR, you must use the full image URI.

The Pod container image is also available in the [AWS ECR Public Gallery](#).

AWS Region	Image URI
US East (Ohio)	<code>051821878176.dkr.ecr.us-east-2.amazonaws.com/aws-fis-pod:0.1</code>
US East (N. Virginia)	<code>731367659002.dkr.ecr.us-east-1.amazonaws.com/aws-fis-pod:0.1</code>
US West (N. California)	<code>080694859247.dkr.ecr.us-west-1.amazonaws.com/aws-fis-pod:0.1</code>
US West (Oregon)	<code>864386544765.dkr.ecr.us-west-2.amazonaws.com/aws-fis-pod:0.1</code>
Africa (Cape Town)	<code>056821267933.dkr.ecr.af-south-1.amazonaws.com/aws-fis-pod:0.1</code>
Asia Pacific (Hong Kong)	<code>246405402639.dkr.ecr.ap-east-1.amazonaws.com/aws-fis-pod:0.1</code>
Asia Pacific (Mumbai)	<code>524781661239.dkr.ecr.ap-south-1.amazonaws.com/aws-fis-pod:0.1</code>

<b>AWS Region</b>	<b>Image URI</b>
Asia Pacific (Osaka)	148336246925.dkr.ecr.ap-northeast-3.amazonaws.com/aws-fis-pod:0.1
Asia Pacific (Seoul)	526524659354.dkr.ecr.ap-northeast-2.amazonaws.com/aws-fis-pod:0.1
Asia Pacific (Singapore)	316401638346.dkr.ecr.ap-southeast-1.amazonaws.com/aws-fis-pod:0.1
Asia Pacific (Sydney)	488104106298.dkr.ecr.ap-southeast-2.amazonaws.com/aws-fis-pod:0.1
Asia Pacific (Tokyo)	635234321696.dkr.ecr.ap-northeast-1.amazonaws.com/aws-fis-pod:0.1
Canada (Central)	490658072207.dkr.ecr.ca-central-1.amazonaws.com/aws-fis-pod:0.1
Europe (Frankfurt)	713827034473.dkr.ecr.eu-central-1.amazonaws.com/aws-fis-pod:0.1
Europe (Ireland)	205866052826.dkr.ecr.eu-west-1.amazonaws.com/aws-fis-pod:0.1
Europe (London)	327424803546.dkr.ecr.eu-west-2.amazonaws.com/aws-fis-pod:0.1
Europe (Milan)	478809367036.dkr.ecr.eu-south-1.amazonaws.com/aws-fis-pod:0.1
Europe (Paris)	154605889247.dkr.ecr.eu-west-3.amazonaws.com/aws-fis-pod:0.1
Europe (Spain)	395402409451.dkr.ecr.eu-south-2.amazonaws.com/aws-fis-pod:0.1

AWS Region	Image URI
Europe (Stockholm)	263175118295.dkr.ecr.eu-north-1.amazonaws.com/ aws-fis-pod:0.1
Europe (Zurich)	604225987275.dkr.ecr.eu-central-2.amazonaws.com/ aws-fis-pod:0.1
Middle East (Bahrain)	065825543785.dkr.ecr.me-south-1.amazonaws.com/ aws-fis-pod:0.1
Middle East (UAE)	438374459301.dkr.ecr.me-central-1.amazonaws.com/ aws-fis-pod:0.1
South America (São Paulo)	767113787785.dkr.ecr.sa-east-1.amazonaws.com/ aws-fis-pod:0.1
AWS GovCloud (US-East)	246533647532.dkr.ecr.us-gov-east-1.amazonaws.com/ aws-fis-pod:0.1
AWS GovCloud (US-West)	246529956514.dkr.ecr.us-gov-west-1.amazonaws.com/ aws-fis-pod:0.1

## Example experiment template

The following is an example experiment template for the [the section called “aws:eks:pod-network-latency”](#) action.

```
{
  "description": "Add latency and jitter to the network interface for the target EKS
  Pods",
  "targets": {
    "myPods": {
      "resourceType": "aws:eks:pod",
      "parameters": {
        "clusterIdentifier": "mycluster",
        "namespace": "default",
        "selectorType": "labelSelector",
        "selectorValue": "mylabel=mytarget"
      }
    },
  },
}
```

```

        "selectionMode": "COUNT(3)"
    }
},
"actions": {
    "EksPod-latency": {
        "actionId": "aws:eks:pod-network-latency",
        "description": "Add latency",
        "parameters": {
            "kubernetesServiceAccount": "myserviceaccount",
            "duration": "PT5M",
            "delayMilliseconds": "200",
            "jitterMilliseconds": "10",
            "sources": "0.0.0.0/0"
        },
        "targets": {
            "Pods": "myPods"
        }
    }
},
"stopConditions": [
    {
        "source": "none",
    }
],
"roleArn": "arn:aws:iam::111122223333:role/fis-experiment-role",
"tags": {
    "Name": "EksPodNetworkLatency"
}
}

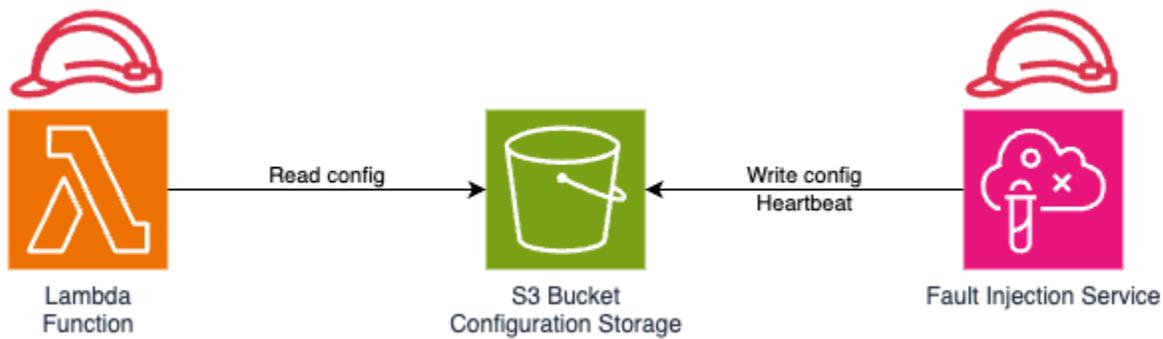
```

## Use the AWS FIS `aws:lambda:function` actions

You can use the `aws:lambda:function` actions to inject faults into invocations of your AWS Lambda functions.

These actions use an AWS FIS managed extension to inject faults. To use `aws:lambda:function` actions, you will need to attach the extension as a layer to your Lambda functions and configure an Amazon S3 bucket to communicate between AWS FIS and the extension.

When you run an AWS FIS experiment targeting `aws:lambda:function`, AWS FIS reads the Amazon S3 configuration from your Lambda function and writes fault injection information to the specified Amazon S3 location, as shown in the diagram below.



## Actions

- [the section called “aws:lambda:invocation-add-delay”](#)
- [the section called “aws:lambda:invocation-error”](#)
- [the section called “aws:lambda:invocation-http-integration-response”](#)

## Limitations

- The AWS FIS Lambda extension cannot be used with functions that use response streaming. Even when no faults are applied, the AWS FIS Lambda extension will suppress streaming configurations. For more information, see [Response streaming for Lambda functions](#) in the *AWS Lambda user guide*.

## Prerequisites

Before using AWS FIS Lambda actions, ensure that you have completed these one-time tasks:

- **Create an Amazon S3 bucket in the region you plan to start an experiment from** - You can use a single Amazon S3 bucket for multiple experiments and share the bucket between multiple AWS accounts. However, you must have a separate bucket for each AWS Region.
- **Create an IAM policy to grant read access for the Lambda extension to the Amazon S3 bucket** - In the following template, replace `my-config-distribution-bucket` with the name of the Amazon S3 bucket you created above and `FisConfigs` with the name of a folder in your Amazon S3 bucket you want to use.

JSON

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "AllowListingConfigLocation",
      "Effect": "Allow",
      "Action": ["s3:ListBucket"],
      "Resource": ["arn:aws:s3::my-config-distribution-bucket"],
      "Condition": {
        "StringLike": {
          "s3:prefix": ["FisConfigs/*"]
        }
      }
    },
    {
      "Sid": "AllowReadingObjectFromConfigLocation",
      "Effect": "Allow",
      "Action": "s3:GetObject",
      "Resource": ["arn:aws:s3::my-config-distribution-bucket/FisConfigs/
*"]
    }
  ]
}

```

- **Create an IAM policy to grant write access for the AWS FIS experiment to the Amazon S3 bucket** - In the following template, replace `my-config-distribution-bucket` with the name of the Amazon S3 bucket you created above and `FisConfigs` with the name of a folder in your Amazon S3 bucket you want to use.

JSON

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "AllowFisToWriteAndDeleteFaultConfigurations",
      "Effect": "Allow",
      "Action": [
        "s3:PutObject",
        "s3>DeleteObject"
      ],
      "Resource": "arn:aws:s3::my-config-distribution-bucket/FisConfigs/*"
    }
  ]
}

```

```
    },
    {
      "Sid": "AllowFisToInspectLambdaFunctions",
      "Effect": "Allow",
      "Action": [
        "lambda:GetFunction"
      ],
      "Resource": "*"
    },
    {
      "Sid": "AllowFisToDoTagLookups",
      "Effect": "Allow",
      "Action": [
        "tag:GetResources"
      ],
      "Resource": "*"
    }
  ]
}
```

## Configure Lambda functions

Follow the steps below for every Lambda function that you want to impact:

1. Attach the Amazon S3 read access policy created above to the Lambda function.
2. Attach the AWS FIS extension as a layer to the function. For more information on the layer ARNs, see [Available versions of the AWS FIS extension for Lambda](#).
3. Set the `AWS_FIS_CONFIGURATION_LOCATION` variable to the ARN of the Amazon S3 configuration folder, for example `arn:aws:s3:::my-config-distribution-bucket/FisConfigs/`.
4. Set the `AWS_LAMBDA_EXEC_WRAPPER` variable to `/opt/aws-fis/bootstrap`.

## Configure an AWS FIS experiment

Before running your experiment, ensure that you have attached the Amazon S3 write access policy that you created in the prerequisites to the experiment roles that will use AWS FIS Lambda actions. For more information on how to set up an AWS FIS experiment, see [Managing AWS FIS experiment templates](#).

## Logging

The AWS FIS Lambda extension writes logs to the console and CloudWatch logs. Logging can be configured using the `AWS_FIS_LOG_LEVEL` variable. Supported values are `INFO`, `WARN`, and `ERROR`. Logs will be written in the log format configured for your Lambda function.

The following is an example of a log in text format:

```
2024-08-09T18:51:38.599984Z INFO AWS FIS EXTENSION - extension enabled 1.0.1
```

The following is an example of a log in JSON format:

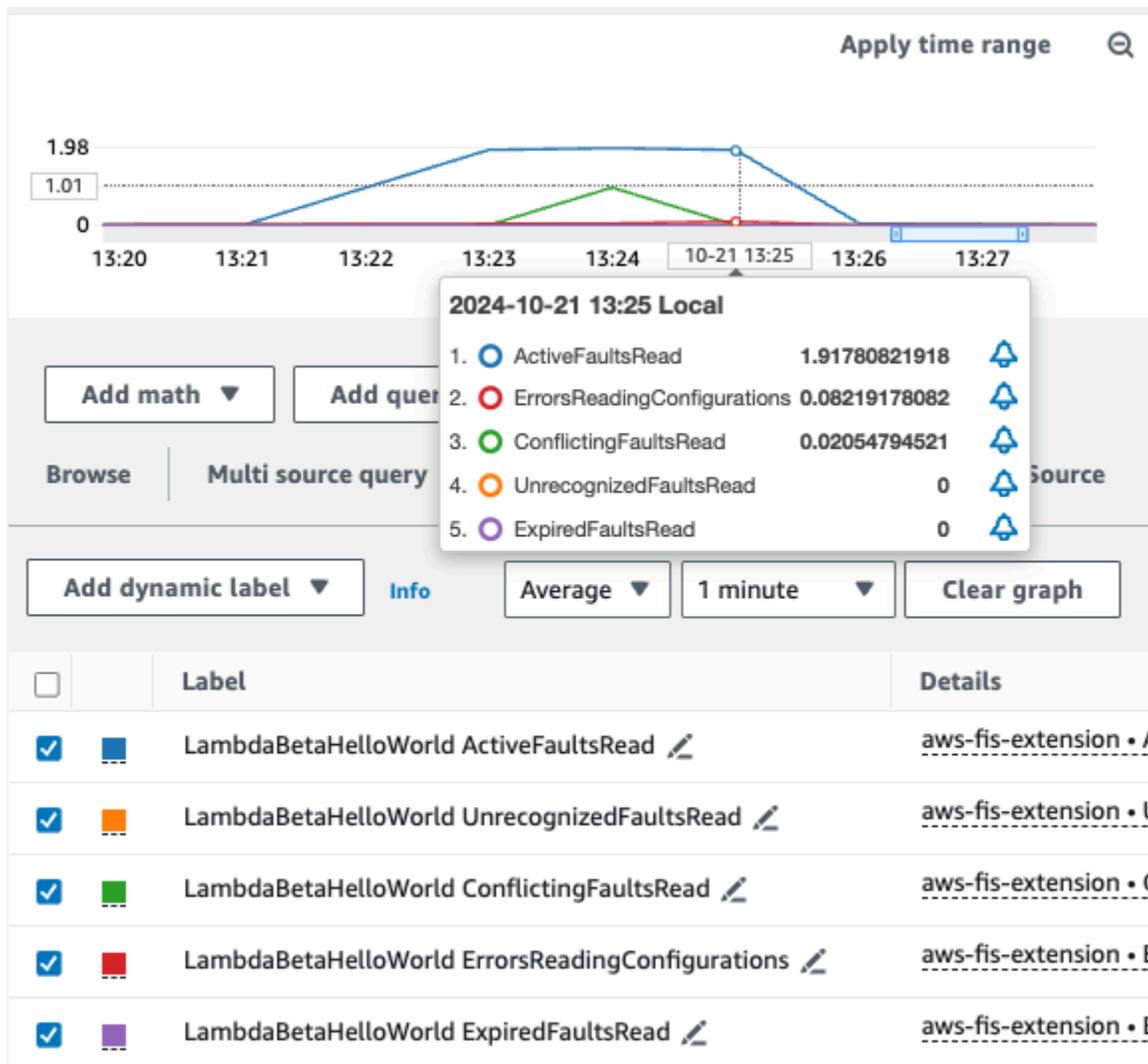
```
{
  "timestamp": "2024-10-08T17:15:36.953905Z",
  "level": "INFO",
  "fields": {
    "message": "AWS FIS EXTENSION - adding 5000 milliseconds of latency to function invocation",
    "requestId": "0608bf70-908f-4a17-bbfe-3782cd783d8b"
  }
}
```

The emitted logs can be used with Amazon CloudWatch metric filters to generate custom metrics. For more information on metric filters, see [Creating metrics from log events using filters](#) in the *Amazon CloudWatch Logs user guide*.

## Using CloudWatch Embedded Metric Format (EMF)

You can configure the AWS FIS Lambda extension to emit EMF logs by setting the `AWS_FIS_EXTENSION_METRICS` variable to `all`. By default, the extension does not emit EMF logs, and `AWS_FIS_EXTENSION_METRICS` defaults to `none`. EMF logs are published in the `aws-fis-extension` namespace on the CloudWatch console.

Within the `aws-fis-extension` namespace, you can select certain metrics to be displayed in a graph. The example below shows some of the available metrics in the `aws-fis-extension` namespace.



## Advanced topics

This section provides additional information on how AWS FIS works with the Lambda extension and special use cases.

### Topics

- [Understanding polling](#)
- [Understanding concurrency](#)
- [Understanding invocation percentage](#)
- [Special considerations for SnapStart](#)
- [Special considerations for fast infrequent functions](#)
- [Configuring multiple extensions using Lambda Runtime API proxy](#)

- [Using AWS FIS with container runtimes](#)
- [AWS FIS Lambda environment variables](#)

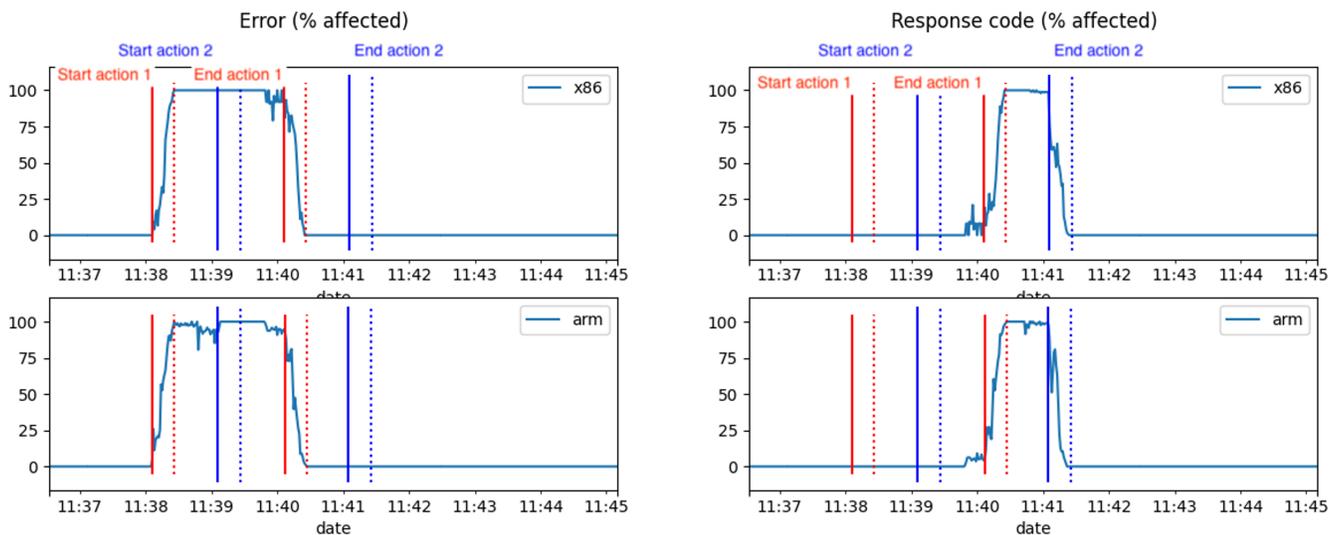
## Understanding polling

You may notice a ramp-up period of up to 60s before faults begin to affect all invocations. This is because the Lambda extension polls for configuration information infrequently while waiting for an experiment to start. You can adjust the polling interval by setting the `AWS_FIS_SLOW_POLL_INTERVAL_SECONDS` environment variable (default 60s). A lower value will poll more often but incur greater performance impact and cost. You may also notice a ramp-down period of up to 20s after the fault has been injected. This is because the extension polls more frequently while experiments are running.

## Understanding concurrency

You may target the same Lambda functions with multiple actions concurrently. If the actions are all different from each other, then all actions will be applied. For example, you can add an initial delay before returning an error. If two identical or conflicting actions are applied to the same function, then only the action with the earliest start date will be applied.

The figure below shows two conflicting actions, **aws:lambda:invocation-error** and **aws:lambda:invocation-http-integration-response**, overlapping. Initially, **aws:lambda:invocation-error** ramps up at 11:38 and runs for 2 minutes. Then, **aws:lambda:invocation-http-integration-response** attempts to start at 11:39, but does not come into effect until 11:40 after the first action has concluded. To maintain experiment timing, **aws:lambda:invocation-http-integration-response** still finishes at the originally intended time of 11:41.



## Understanding invocation percentage

The AWS Fault Injection Service Lambda actions use an **aws:lambda:function** target which allows you to select one or more AWS Lambda function ARNs. Using these ARNs, the AWS Fault Injection Service Lambda actions can inject faults in every invocation of the selected Lambda function. To allow you to inject faults into only a fraction of invocations, each action allows you to specify an `invocationPercentage` parameter with values from 0 to 100. Using the `invocationPercentage` parameter, you can ensure that actions are concurrent even for invocation percentages below 100%.

## Special considerations for SnapStart

AWS Lambda functions with SnapStart enabled will have a higher likelihood of waiting the full duration of `AWS_FIS_SLOW_POLL_INTERVAL_SECONDS` before picking up the first fault configuration, even if an experiment is already running. This is because Lambda SnapStart uses a single snapshot as the initial state for multiple execution environments and persists temporary storage. For the AWS Fault Injection Service Lambda extension it will persist polling frequency and skip the initial configuration check on initialization of the execution environment. For more information on Lambda SnapStart, see [Improving startup performance with Lambda SnapStart](#) in the *AWS Lambda user guide*.

## Special considerations for fast infrequent functions

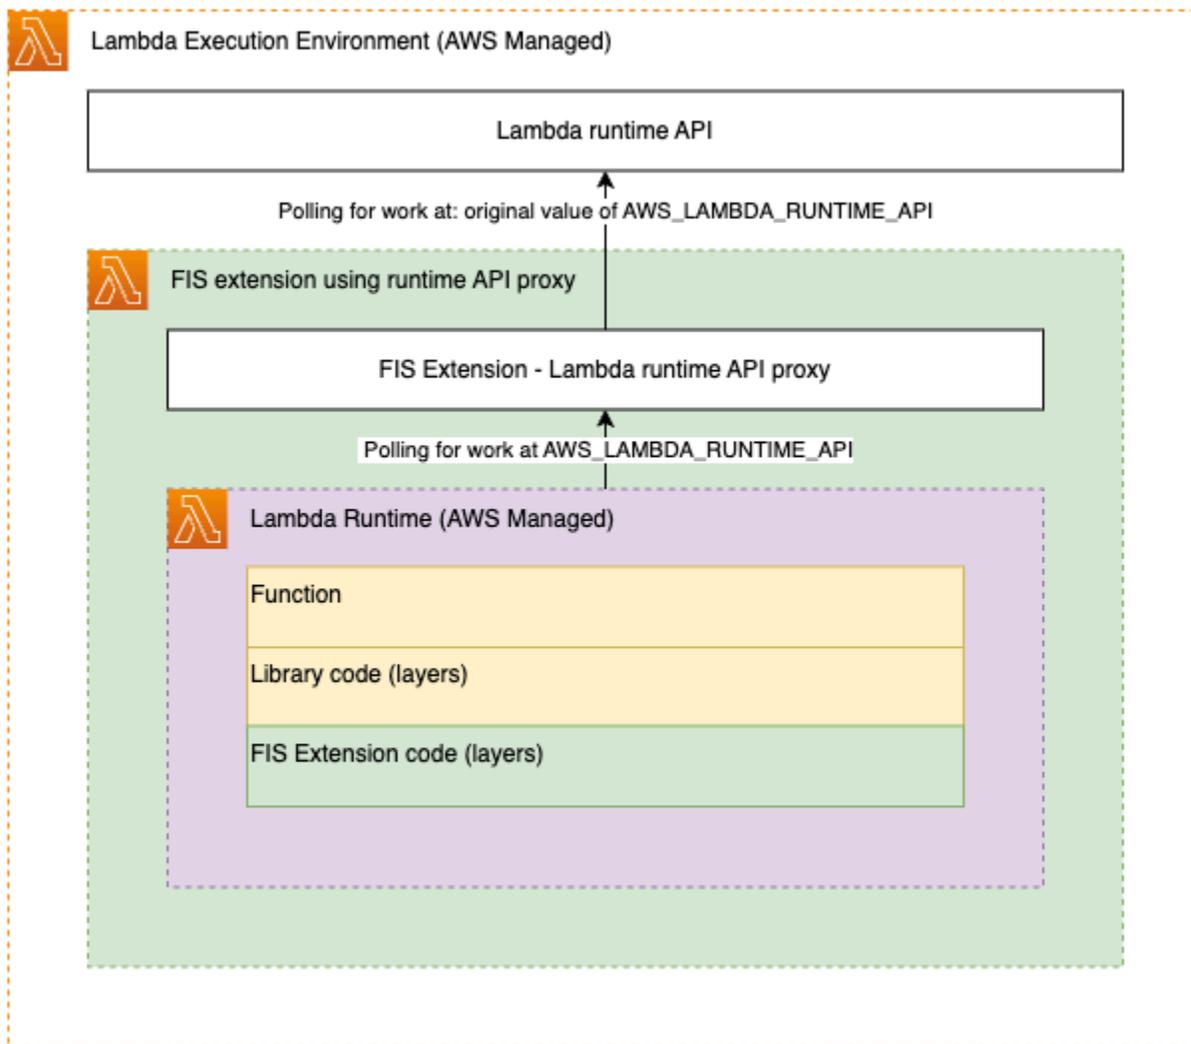
If your Lambda function runs for less than the average poll duration of 70 milliseconds then the polling thread may need multiple invocations to obtain fault configurations. If the function runs

infrequently, for example once every 15 minutes, then the poll will never be completed. To ensure the polling thread can finish, set the `AWS_FIS_POLL_MAX_WAIT_MILLISECONDS` parameter. The extension will wait up to the duration that you set for an in-flight poll to finish before starting the function. Note that this will increase the billed function duration and lead to an additional delay on some invocations.

## Configuring multiple extensions using Lambda Runtime API proxy

The Lambda extension uses the AWS Lambda Runtime API proxy to intercept function invocations before they reach the runtime. It does this by exposing a proxy for the AWS Lambda Runtime API to the runtime and advertising its location in the `AWS_LAMBDA_RUNTIME_API` variable.

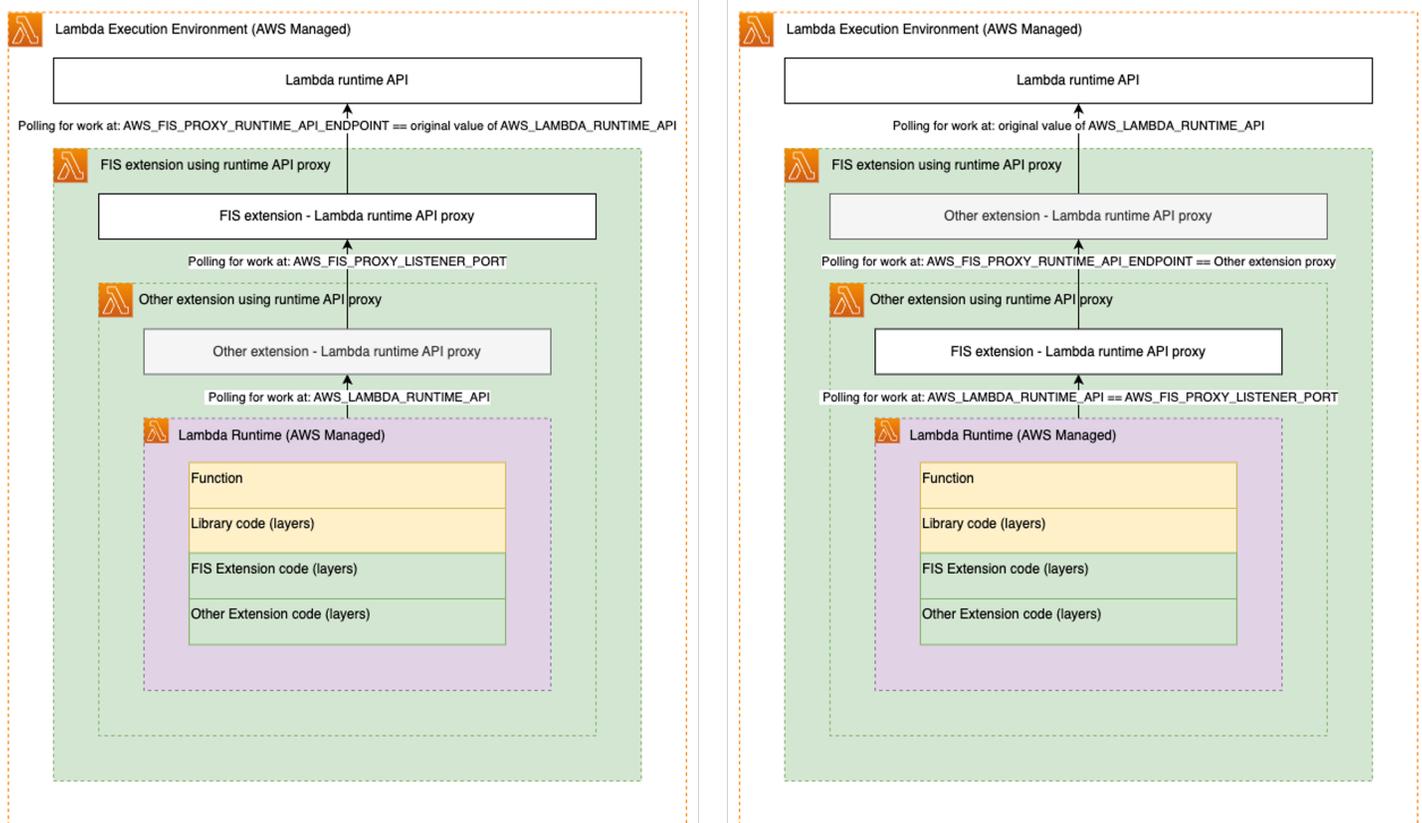
The following diagram shows the configuration for a single extension using the Lambda Runtime API proxy:



To use the AWS FIS Lambda extension with another extension using the AWS Lambda Runtime API proxy pattern, you will need to chain the proxies using a custom bootstrap script. The AWS FIS Lambda extension accepts the following environment variables:

- `AWS_FIS_PROXY_RUNTIME_API_ENDPOINT` - Takes a string in the form `127.0.0.1:9876` representing the local IP and listener port for the AWS Lambda Runtime API. This could be the original value of `AWS_LAMBDA_RUNTIME_API` or the location of another proxy.
- `AWS_FIS_PROXY_LISTENER_PORT` - Takes a port number on which the AWS FIS extension should start its own proxy, by default `9100`.

With these settings you can chain the AWS FIS extension with another extension using the Lambda Runtime API proxy in two different orders.



For more information on the AWS Lambda Runtime API proxy, see [Enhancing runtime security and governance with the AWS Lambda Runtime API proxy extension](#) and [Using the Lambda runtime API for custom runtimes](#) in the *AWS Lambda user guide*.

## Using AWS FIS with container runtimes

For AWS Lambda functions using container images that accept the `AWS_LAMBDA_RUNTIME_API` environment variable, you can package the AWS FIS Lambda extension into your container image by following the steps below:

1. Determine the ARN of the layer from which to extract the extension. For more information on how to find the ARN, see [Configure Lambda functions](#).
2. Use the AWS Command Line Interface (CLI) to request details about the extension `aws lambda get-layer-version-by-arn --arn fis-extension-arn`. The response will contain a `Location` field containing a pre-signed URL from which you can download the FIS extension as a ZIP file.
3. Unzip the content of the extension into `/opt` of your Docker filesystem. The following is an example Dockerfile based on the NodeJS Lambda runtime:

```
# extension installation #
FROM amazon/aws-lambda-nodejs:12 AS builder
COPY extension.zip extension.zip
RUN yum install -y unzip
RUN mkdir -p /opt
RUN unzip extension.zip -d /opt
RUN rm -f extension.zip
FROM amazon/aws-lambda-nodejs:12
WORKDIR /opt
COPY --from=builder /opt .
# extension installation finished #
# JS example. Modify as required by your runtime
WORKDIR ${LAMBDA_TASK_ROOT}
COPY index.js package.json .
RUN npm install
CMD [ "index.handler" ]
```

For more information on container images, see [Create a Lambda function using a container image](#) in the *AWS Lambda user guide*.

## AWS FIS Lambda environment variables

The following is a list of environment variables for the AWS FIS Lambda extension

- `AWS_FIS_CONFIGURATION_LOCATION` - Required. Location where AWS FIS will write active fault configurations and the extension will read fault configurations. The locations should be in Amazon S3 ARN format including a bucket and path. For example, `arn:aws:s3:::my-fis-config-bucket/FisConfigs/`.
- `AWS_LAMBDA_EXEC_WRAPPER` - Required. Location of the AWS Lambda [wrapper script](#) used to configure the AWS FIS Lambda extension. This should be set to the `/opt/aws-fis/bootstrap` script that is included with the extension.
- `AWS_FIS_LOG_LEVEL` - Optional. Log level for messages emitted by AWS FIS Lambda extension. Supported values are `INFO`, `WARN`, and `ERROR`. If not set, AWS FIS extension will default to `INFO`.
- `AWS_FIS_EXTENSION_METRICS` - Optional. Possible values are `all` and `none`. If set to `all` the extension will emit EMF metrics under the `aws-fis-extension` namespace.
- `AWS_FIS_SLOW_POLL_INTERVAL_SECONDS` - Optional. If set will override the polling interval (in seconds) while the extension is not injecting faults and waiting for a fault configuration to be added to configuration location. Defaults to `60`.
- `AWS_FIS_PROXY_RUNTIME_API_ENDPOINT` - Optional. If set will override the value of `AWS_LAMBDA_RUNTIME_API` to define where the AWS FIS extension interacts with the AWS Lambda runtime API to control function invocation. Expects `IP:PORT`, for example, `127.0.0.1:9000`. For more information on `AWS_LAMBDA_RUNTIME_API`, see [Using the Lambda runtime API for custom runtimes](#) in the *AWS Lambda user guide*.
- `AWS_FIS_PROXY_LISTENER_PORT` - Optional. Defines the port on which the AWS FIS Lambda extension exposes an AWS Lambda runtime API proxy that can be used by another extension or the runtime. Defaults to `9100`.
- `AWS_FIS_POLL_MAX_WAIT_MILLISECONDS` - Optional. If set to non-zero value, this variable defines the number of milliseconds the extension will wait for an in-flight async poll to finish before evaluating fault configurations and starting the invocation of the runtime. Defaults to `0`.

## Available versions of the AWS FIS extension for Lambda

This section includes information about the AWS FIS Lambda extension versions. The extension supports Lambda functions developed for the x86-64 and ARM64 (Graviton2) platforms. Your Lambda function must be configured to use the specific Amazon Resource Name (ARN) for the AWS Region where it is currently hosted. You can view AWS Region and ARN details below.

### Topics

- [AWS FIS Lambda extension release notes](#)

- [Access Guide for Lambda Extension ARNs](#)

## AWS FIS Lambda extension release notes

The following table describes changes made to recent versions of the AWS FIS Lambda extension

Version	Launch date	Notes
1.0.0	2024-10-29	Initial release

## Access Guide for Lambda Extension ARNs

You must have at least one parameter in your AWS account and AWS Region before you can search for public parameters using the console. To discover public parameters, see [Discovering public parameters in Parameter Store](#).

### Console Access:

1. Open the AWS Systems Manager console at <https://console.aws.amazon.com/systems-manager/>.
2. In the navigation pane, choose **Parameter Store**.
3. Choose the **Public parameters** tab.
4. Choose the **Select a service** dropdown. From the dropdown options, choose `fis`.
5. (Optional) Filter the parameters you selected by entering more information into the search bar. For arm64 architectures, filter parameters by entering `"arm64"`. For x86\_64 architectures, filter parameters by entering `"x86_64"`.
6. Choose the public parameter you want to use.
7. From the parameter details, locate the ARN value. Copy the ARN to use in configuring layer extensions on your target Lambda functions.

### AWS CLI Access:

### SSM Parameter Names

The following SSM parameter names are available for different architectures:

1. arm64: /aws/service/fis/lambda-extension/AWS-FIS-extension-arm64/1.x.x
2. x86\_64: /aws/service/fis/lambda-extension/AWS-FIS-extension-x86\_64/1.x.x

## AWS CLI Command Format

To retrieve the extension ARNs, use the following AWS CLI command format where *parameterName* is the name for the architecture and *region* is the target AWS Region:

```
aws ssm get-parameter --name parameterName --region region
```

## Example Usage

```
aws ssm get-parameter --name /aws/service/fis/lambda-extension/AWS-FIS-extension-x86_64/1.x.x --region ap-southeast-2
```

## Response Format

The command returns a JSON object containing the parameter details like the following. The ARN of the Lambda layer is included in the *Value* field of the *Parameter* object. Copy the ARN to use in configuring layer extensions on your target Lambda functions.

```
{
  "Parameter": {
    "Name": "/aws/service/fis/lambda-extension/AWS-FIS-extension-x86_64/1.x.x",
    "Type": "String",
    "Value": "arn:aws:lambda:ap-southeast-2:211125361907:layer:aws-fis-extension-x86_64:9",
    "Version": 1,
    "LastModifiedDate": "2025-01-02T15:13:54.465000-05:00",
    "ARN": "arn:aws:ssm:ap-southeast-2::parameter/aws/service/fis/lambda-extension/AWS-FIS-extension-x86_64/1.x.x",
    "DataType": "text"
  }
}
```

## Programmatic Access:

Retrieve these public parameters programmatically when building or configuring your Lambda functions using Infrastructure as Code (IaC). This approach helps maintain your Lambda functions with the latest layer version ARN without requiring manual code updates that would be necessary if the AWS FIS extension layer ARN were hardcoded. The following resources show how to retrieve public parameters using common IaC platforms:

- [Get public parameters using the AWS SDK](#)
- [Get public parameters from AWS Systems Manager Parameter Store using the AWS CDK](#)
- [Get public parameters using Terraform](#)

# Managing AWS FIS experiment templates

You can create and manage experiment templates using the AWS FIS console or the command line. An experiment template contains one or more actions to run on specified targets during an experiment. It also contains the stop conditions that prevent the experiment from going out of bounds. For more information on the components of an experiment template, see [Experiment template components](#). After you create an experiment template, you can use it to run an experiment.

## Tasks

- [Create an experiment template](#)
- [View experiment templates](#)
- [Generate a target preview from an experiment template](#)
- [Start an experiment from a template](#)
- [Update an experiment template](#)
- [Tag experiment templates](#)
- [Delete an experiment template](#)
- [Example AWS FIS experiment templates](#)

## Create an experiment template

Before you begin, complete the following tasks:

- [Plan your experiment](#).
- Create an IAM role that grants the AWS FIS service permission to perform actions on your behalf. For more information, see [IAM roles for AWS FIS experiments](#).
- Ensure that you have access to AWS FIS. For more information, see [AWS FIS policy examples](#).

### To create an experiment template using the console

1. Open the AWS FIS console at <https://console.aws.amazon.com/fis/>.
2. In the navigation pane, choose **Experiment templates**.
3. Choose **Create experiment template**.
4. For **Step 1, Specify template details**, do the following:

- a. For **Description and name**, enter a description for the template, such as Amazon S3 Network Disrupt Connectivity.
  - b. (Optional) For **Account targeting**, choose **Multiple accounts** to configure a multi-account experiment template.
  - c. Choose **Next**, and move to **Step 2, Specify actions and targets**.
5. For **Actions**, specify the set of actions for the template. For each action, choose **Add action** and complete the following:

- For **Name**, enter a name for the action.

Allowed characters are alphanumeric characters, hyphens (-), and underscores(\_). The name must start with a letter. No spaces are allowed. Each action name must be unique in this template.

- (Optional) For **Description**, enter a description for the action. The maximum length is 512 characters.
- (Optional) For **Start after**, select another action defined in this template that must be completed before the current action starts. Otherwise, the action runs at the start of the experiment.
- For **Action type**, choose the AWS FIS action.
- For **Target**, choose a target that you defined in the **Targets** section. If you haven't defined a target for this action yet, AWS FIS creates a new target for you.
- For **Action parameters**, specify the parameters for the action. This section appears only if the AWS FIS action has parameters.
- Choose **Save**.

6. For **Targets**, define the target resources on which to carry out the actions. You must specify at least one resource ID or one resource tag as a target. Choose **Edit** to edit the target that AWS FIS created for you in the previous step, or choose **Add target**. For each target, do the following:

- For **Name**, enter a name for the target.

Allowed characters are alphanumeric characters, hyphens (-), and underscores(\_). The name must start with a letter. No spaces are allowed. Each target name must be unique in this template.

- For **Resource type**, choose a resource type that is supported for the action.

- For **Target method**, do one of the following:
    - Choose **Resource IDs** and then choose or add the resource IDs.
    - Choose **Resource tags, filters, and parameters** and then add the tags and filters that you need. For more information, see [the section called “Identify target resources”](#).
  - For **Selection mode**, choose **Count** to run the action on the specified number of identified targets or choose **Percent** to run the action on the specified percentage of identified targets. By default, the action runs on all identified targets.
  - Choose **Save**.
7. To update an action with the target that you created, find the action under **Actions**, choose **Edit**, and then update **Target**. You can use the same target for multiple actions.
  8. (Optional) For **experiment options**, select the behavior of the empty target resolution mode.
  9. Choose **Next** to move to **Step 3, Configure service access**.
  10. For **Service Access**, choose **Use an existing IAM role**, and then choose the IAM role that you created as described in the prerequisites for this tutorial. If your role is not displayed, verify that it has the required trust relationship. For more information, see [the section called “Experiment role”](#).
  11. (Multi-account experiments only) For **Target account configurations**, add a Role ARN and optional description for each target account. To upload the target account role ARNs with a CSV file, choose **Upload role ARNs for all target accounts** and then choose **Choose .CSV file**.
  12. Choose **Next** to move to **Step 4, Configure optional settings**.
  13. (Optional) For **Stop conditions**, select the Amazon CloudWatch alarms for the stop conditions. For more information, see [Stop conditions for AWS FIS](#).
  14. (Optional) For **Logs**, configure the destination option. To send logs to an S3 bucket, choose **Send to an Amazon S3 bucket** and enter the bucket name and prefix. To send logs to CloudWatch Logs, choose **Send to CloudWatch Logs** and enter the log group.
  15. (Optional) For **Tags**, choose **Add new tag** and specify a tag key and tag value. The tags that you add are applied to your experiment template, not the experiments that are run using the template.
  16. Choose **Next** to move to **Step 5, Review and create**.
  17. Review the template and choose **Create experiment template**. When prompted for confirmation, enter `create`, Then choose **Create experiment template**.

## To create an experiment template using the CLI

Use the [create-experiment-template](#) command.

You can load an experiment template from a JSON file.

Use the `--cli-input-json` parameter.

```
aws fis create-experiment-template --cli-input-json fileb://<path-to-json-file>
```

For more information, see [Generating a CLI skeleton template](#) in the *AWS Command Line Interface User Guide*. For example templates, see [Example AWS FIS experiment templates](#).

## View experiment templates

You can view the experiment templates that you created.

### To view an experiment template using the console

1. Open the AWS FIS console at <https://console.aws.amazon.com/fis/>.
2. In the navigation pane, choose **Experiment templates**.
3. To view information about a specific template, select the **Experiment template ID**.
4. In the **Details** section, you can view the description and stop conditions for the template.
5. To view the actions for the experiment template, choose **Actions**.
6. To view the targets for the experiment template, choose **Targets**.
7. To view the tags for the experiment template, choose **Tags**.

### To view an experiment template using the CLI

Use the [list-experiment-templates](#) command to get a list of experiment templates, and use the [get-experiment-template](#) command to get information about a specific experiment template.

## Generate a target preview from an experiment template

Before you start an experiment, you can generate a target preview to verify that your experiment template is configured to target the expected resources. The resources that are targeted when you begin the actual experiment may be different from those in the preview, as resources may

be removed, updated, or sampled randomly. When you generate a target preview, you start an experiment that skips all actions.

### Note

Generating a target preview does not allow you to verify that you have the necessary permissions to take actions on your resources.

## To start a target preview using the console

1. Open the AWS FIS console at <https://console.aws.amazon.com/fis/>.
2. In the navigation pane, choose **Experiment templates**.
3. To view the targets for the experiment template, choose **Targets**.
4. To verify your target resources for the experiment template, choose **Generate Preview**. When you run an experiment, this target preview will be automatically updated with the targets from the most recent experiment.

## To start a target preview using the CLI

- Run the following [start-experiment](#) command. Replace the values in italics with your own values.

```
aws fis start-experiment \  
  --experiment-options actionsMode=skip-all \  
  --experiment-template-id EXTxxxxxxxxxx
```

## Start an experiment from a template

After you have created an experiment template, you can start experiments using that template.

When you start an experiment, we create a snapshot of the specified template and use that snapshot to run the experiment. Therefore, if the experiment template is updated or deleted while the experiment is running, those changes have no impact on the running experiment.

When you start an experiment, AWS FIS creates a service-linked role on your behalf. For more information, see [Use service-linked roles for AWS Fault Injection Service](#).

After you start the experiment, you can stop it at any time. For more information, see [Stop an experiment](#).

### To start an experiment using the console

1. Open the AWS FIS console at <https://console.aws.amazon.com/fis/>.
2. In the navigation pane, choose **Experiment templates**.
3. Select the experiment template, and choose **Start experiment**.
4. (Optional) To add a tag to your experiment, choose **Add new tag** and enter a tag key and a tag value.
5. Choose **Start experiment**. When prompted for confirmation, enter **start** and choose **Start experiment**.

### To start an experiment using the CLI

Use the [start-experiment](#) command.

## Update an experiment template

You can update an existing experiment template. When you update an experiment template, the changes do not affect any running experiments that use the template.

### To update an experiment template using the console

1. Open the AWS FIS console at <https://console.aws.amazon.com/fis/>.
2. In the navigation pane, choose **Experiment templates**.
3. Select the experiment template, and choose **Actions, Update experiment template**.
4. Modify the template details as needed, and choose **Update experiment template**.

### To update an experiment template using the CLI

Use the [update-experiment-template](#) command.

## Tag experiment templates

You can apply your own tags to experiment templates to help you organize them. You can also implement [tag-based IAM policies](#) to control access to experiment templates.

## To tag an experiment template using the console

1. Open the AWS FIS console at <https://console.aws.amazon.com/fis/>.
2. In the navigation pane, choose **Experiment templates**.
3. Select the experiment template and choose **Actions, Manage tags**.
4. To add a new tag, choose **Add new tag**, and then specify a key and value.  
  
To remove a tag, choose **Remove** for the tag.
5. Choose **Save**.

## To tag an experiment template using the CLI

Use the [tag-resource](#) command.

## Delete an experiment template

If you no longer need an experiment template, you can delete it. When you delete an experiment template, any running experiments that use the template are not affected. The experiment continues to run until completed or stopped. However, experiment templates that are deleted are not available for viewing from the **Experiments** page in the console.

## To delete an experiment template using the console

1. Open the AWS FIS console at <https://console.aws.amazon.com/fis/>.
2. In the navigation pane, choose **Experiment templates**.
3. Select the experiment template, and choose **Actions, Delete experiment template**.
4. When prompted for confirmation, enter **delete** and choose **Delete experiment template**.

## To delete an experiment template using the CLI

Use the [delete-experiment-template](#) command.

## Example AWS FIS experiment templates

If you're using the AWS FIS API or a command line tool to create an experiment template, you can construct the template in JavaScript Object Notation (JSON). For more information about the components of an experiment template, see [AWS FIS experiment template components](#).

To create an experiment using one of the example templates, save it to a JSON file (for example, `my-template.json`), replace the placeholder values in *italics* with your own values, and then run the following [create-experiment-template](#) command.

```
aws fis create-experiment-template --cli-input-json file://my-template.json
```

## Example templates

- [Stop EC2 instances based on filters](#)
- [Stop a specified number of EC2 instances](#)
- [Run a pre-configured AWS FIS SSM document](#)
- [Run a predefined Automation runbook](#)
- [Throttle API actions on EC2 instances with the target IAM role](#)
- [Stress test CPU of pods in a Kubernetes cluster](#)
- [Provisioned throughput exception for specified number of Kinesis Data Streams](#)
- [Experiment role permissions example](#)

## Stop EC2 instances based on filters

The following example stops all running Amazon EC2 instances in the specified Region with the specified tag in the specified VPC. It restarts them after two minutes.

```
{
  "tags": {
    "Name": "StopEC2InstancesWithFilters"
  },
  "description": "Stop and restart all instances in us-east-1b with the tag env=prod
in the specified VPC",
  "targets": {
    "myInstances": {
      "resourceType": "aws:ec2:instance",
      "resourceTags": {
        "env": "prod"
      },
    },
    "filters": [
      {
        "path": "Placement.AvailabilityZone",
        "values": ["us-east-1b"]
      },
    ],
  },
}
```

```

        {
            "path": "State.Name",
            "values": ["running"]
        },
        {
            "path": "VpcId",
            "values": [ "vpc-aabbcc11223344556" ]
        }
    ],
    "selectionMode": "ALL"
}
},
"actions": {
    "StopInstances": {
        "actionId": "aws:ec2:stop-instances",
        "description": "stop the instances",
        "parameters": {
            "startInstancesAfterDuration": "PT2M"
        },
        "targets": {
            "Instances": "myInstances"
        }
    }
},
"stopConditions": [
    {
        "source": "aws:cloudwatch:alarm",
        "value": "arn:aws:cloudwatch:us-east-1:111122223333:alarm:alarm-name"
    }
],
"roleArn": "arn:aws:iam::111122223333:role/role-name"
}

```

## Stop a specified number of EC2 instances

The following example stops three instances with the specified tag. AWS FIS selects the specific instances to stop at random. It restarts these instances after two minutes.

```

{
    "tags": {
        "Name": "StopEC2InstancesByCount"
    },
    "description": "Stop and restart three instances with the specified tag",

```

```

"targets": {
  "myInstances": {
    "resourceType": "aws:ec2:instance",
    "resourceTags": {
      "env": "prod"
    },
    "selectionMode": "COUNT(3)"
  }
},
"actions": {
  "StopInstances": {
    "actionId": "aws:ec2:stop-instances",
    "description": "stop the instances",
    "parameters": {
      "startInstancesAfterDuration": "PT2M"
    },
    "targets": {
      "Instances": "myInstances"
    }
  }
},
"stopConditions": [
  {
    "source": "aws:cloudwatch:alarm",
    "value": "arn:aws:cloudwatch:us-east-1:111122223333:alarm:alarm-name"
  }
],
"roleArn": "arn:aws:iam::111122223333:role/role-name"
}

```

## Run a pre-configured AWS FIS SSM document

The following example runs a CPU fault injection for 60 seconds on the specified EC2 instance using a pre-configured AWS FIS SSM document, [AWSFIS-Run-CPU-Stress](#). AWS FIS monitors the experiment for two minutes.

```

{
  "tags": {
    "Name": "CPUStress"
  },
  "description": "Run a CPU fault injection on the specified instance",
  "targets": {
    "myInstance": {

```

```

        "resourceType": "aws:ec2:instance",
        "resourceArns": ["arn:aws:ec2:us-east-1:111122223333:instance/instance-
id"],
        "selectionMode": "ALL"
    }
},
"actions": {
    "CPUStress": {
        "actionId": "aws:ssm:send-command",
        "description": "run cpu stress using ssm",
        "parameters": {
            "duration": "PT2M",
            "documentArn": "arn:aws:ssm:us-east-1::document/AWSFIS-Run-CPU-Stress",
            "documentParameters": "{\"DurationSeconds\": \"60\",
\\\"InstallDependencies\\\": \\\"True\\\", \\\"CPU\\\": \\\"0\\\"}"
        },
        "targets": {
            "Instances": "myInstance"
        }
    }
},
"stopConditions": [
    {
        "source": "aws:cloudwatch:alarm",
        "value": "arn:aws:cloudwatch:us-east-1:111122223333:alarm:alarm-name"
    }
],
"roleArn": "arn:aws:iam::111122223333:role/role-name"
}

```

## Run a predefined Automation runbook

The following example publishes a notification to Amazon SNS using a runbook provided by Systems Manager, [AWS-PublishSNSNotification](#). The role must have permissions to publish notifications to the specified SNS topic.

```

{
    "description": "Publish event through SNS",
    "stopConditions": [
        {
            "source": "none"
        }
    ],
}

```

```

"targets": {
},
"actions": {
  "sendToSns": {
    "actionId": "aws:ssm:start-automation-execution",
    "description": "Publish message to SNS",
    "parameters": {
      "documentArn": "arn:aws:ssm:us-east-1::document/AWS-
PublishSNSNotification",
      "documentParameters": "{\"Message\": \"Hello, world\", \"TopicArn\":
\\\"arn:aws:sns:us-east-1:111122223333:topic-name\\\"}\",
      "maxDuration": "PT1M"
    },
    "targets": {
    }
  }
},
"roleArn": "arn:aws:iam::111122223333:role/role-name"
}

```

## Throttle API actions on EC2 instances with the target IAM role

The following example throttles 100% of the API calls specified in the action definition for API calls made by the IAM role(s) specified in the target definition.

### Note

If you wish to target EC2 instances that are members of an Auto Scaling group, please use the **aws:ec2:asg-insufficient-instance-capacity-error** action, and target by Auto Scaling group instead. For more information, see [aws:ec2:asg-insufficient-instance-capacity-error](#).

```

{
  "tags": {
    "Name": "ThrottleEC2APIActions"
  },
  "description": "Throttle the specified EC2 API actions on the specified IAM role",
  "targets": {
    "myRole": {
      "resourceType": "aws:iam:role",
      "resourceArns": ["arn:aws:iam::111122223333:role/role-name"],

```

```

        "selectionMode": "ALL"
    }
},
"actions": {
    "ThrottleAPI": {
        "actionId": "aws:fis:inject-api-throttle-error",
        "description": "Throttle APIs for 5 minutes",
        "parameters": {
            "service": "ec2",
            "operations": "DescribeInstances,DescribeVolumes",
            "percentage": "100",
            "duration": "PT2M"
        },
        "targets": {
            "Roles": "myRole"
        }
    }
},
"stopConditions": [
    {
        "source": "aws:cloudwatch:alarm",
        "value": "arn:aws:cloudwatch:us-east-1:111122223333:alarm:alarm-name"
    }
],
"roleArn": "arn:aws:iam::111122223333:role/role-name"
}

```

## Stress test CPU of pods in a Kubernetes cluster

The following example uses Chaos Mesh to stress test the CPU of pods in an Amazon EKS Kubernetes cluster for one minute.

```

{
    "description": "ChaosMesh StressChaos example",
    "targets": {
        "Cluster-Target-1": {
            "resourceType": "aws:eks:cluster",
            "resourceArns": [
                "arn:aws:eks:arn:aws::111122223333:cluster/cluster-id"
            ],
            "selectionMode": "ALL"
        }
    }
},

```

```

"actions": {
  "TestCPUStress": {
    "actionId": "aws:eks:inject-kubernetes-custom-resource",
    "parameters": {
      "maxDuration": "PT2M",
      "kubernetesApiVersion": "chaos-mesh.org/v1alpha1",
      "kubernetesKind": "StressChaos",
      "kubernetesNamespace": "default",
      "kubernetesSpec": "{\"selector\":{\"namespaces\":[\"default\"],\"labelSelectors\":{\"run\":\"nginx\"}},\"mode\":\"all\",\"stressors\":{\"cpu\":{\"workers\":1,\"load\":50}},\"duration\":\"1m\"}"
    },
    "targets": {
      "Cluster": "Cluster-Target-1"
    }
  }
},
"stopConditions": [{
  "source": "none"
}],
"roleArn": "arn:aws:iam::111122223333:role/role-name",
"tags": {}
}

```

The following example uses Litmus to stress test the CPU of pods in an Amazon EKS Kubernetes cluster for one minute.

```

{
  "description": "Litmus CPU Hog",
  "targets": {
    "MyCluster": {
      "resourceType": "aws:eks:cluster",
      "resourceArns": [
        "arn:aws:eks:arn:aws::111122223333:cluster/cluster-id"
      ],
      "selectionMode": "ALL"
    }
  },
  "actions": {
    "MyAction": {
      "actionId": "aws:eks:inject-kubernetes-custom-resource",
      "parameters": {
        "maxDuration": "PT2M",

```

```

        "kubernetesApiVersion": "litmuschaos.io/v1alpha1",
        "kubernetesKind": "ChaosEngine",
        "kubernetesNamespace": "litmus",
        "kubernetesSpec": "{ \"engineState\": \"active\", \"appinfo\":
{ \"appns\": \"default\", \"applabel\": \"run=nginx\", \"appkind\": \"deployment\" },
 \"chaosServiceAccount\": \"litmus-admin\", \"experiments\": [ { \"name\": \"pod-cpu-hog
\", \"spec\": { \"components\": { \"env\": [ { \"name\": \"TOTAL_CHAOS_DURATION\", \"value\":
\"60\" }, { \"name\": \"CPU_CORES\", \"value\": \"1\" }, { \"name\": \"PODS_AFFECTED_PERC\",
\"value\": \"100\" }, { \"name\": \"CONTAINER_RUNTIME\", \"value\": \"docker\" }, { \"name\":
\"SOCKET_PATH\", \"value\": \"/var/run/docker.sock\" } ] }, \"probe\": [ ] } } ], \"annotationCheck
\": \"false\" }"
    },
    "targets": {
        "Cluster": "MyCluster"
    }
}
},
"stopConditions": [ {
    "source": "none"
} ],
"roleArn": "arn:aws:iam::111122223333:role/role-name",
"tags": {}
}

```

## Provisioned throughput exception for specified number of Kinesis Data Streams

The following example throws a provisioned throughput exception for 100% of requests up to five Kinesis Data Streams with the specified tag. AWS FIS selects the streams to affect at random. After 5 minutes the fault is removed.

```

{
  "description": "Kinesis stream experiment",
  "targets": {
    "KinesisStreams-Target-1": {
      "resourceType": "aws:kinesis:stream",
      "resourceTags": {
        "tag-key": "tag-value"
      },
      "selectionMode": "COUNT(5)"
    }
  },
}

```

```

"actions": {
  "kinesis": {
    "actionId": "aws:kinesis:stream-provisioned-throughput-exception",
    "description": "my-stream",
    "parameters": {
      "duration": "PT5M",
      "percentage": "100",
      "service": "kinesis"
    },
    "targets": {
      "KinesisStreams": "KinesisStreams-Target-1"
    }
  }
},
"stopConditions": [
  {
    "source": "none"
  }
],
"roleArn": "arn:aws:iam::111122223333:role/role-name",
"tags": {},
"experimentOptions": {
  "accountTargeting": "single-account",
  "emptyTargetResolutionMode": "fail"
}
}

```

## Experiment role permissions example

The following permission allows you to run the `aws:kinesis:stream-provisioned-throughput-exception` and `aws:kinesis:stream-expired-iterator-exception` actions on a specific stream impacting 50% of requests.

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": "kinesis:InjectApiError",
      "Resource": "*"
      "Condition": {
        "ForAllValues:StringEquals": {
          "kinesis:FisActionId": [

```

```
        "aws:kinesis:stream-provisioned-throughput-exception",
        "aws:kinesis:stream-expired-iterator-exception"
    ],
    "kinesis:FisTargetArns": [
        "arn:aws:kinesis:us-east-1:111122223333:stream/stream-name"
    ],
    },
    "NumericEquals": {
        "kinesis:FisInjectPercentage": "50"
    }
}
},
{
    "Action": [
        "kinesis:DescribeStreamSummary",
    ],
    "Resource": "*",
    "Effect": "Allow"
}
]
}
```

# Managing your AWS FIS experiments

AWS FIS enables you to perform fault injection experiments on your AWS workloads. To get started, create an [experiment template](#). After you create an experiment template, you can use it to start an experiment.

An experiment is finished when one of the following occurs:

- All [actions](#) in the template completed successfully.
- A [stop condition](#) is triggered.
- An action cannot be completed because of an error. For example, if the [target](#) cannot be found.
- The experiment is [stopped manually](#).

You cannot resume a stopped or failed experiment. You also cannot rerun a completed experiment. However, you can start a new experiment from the same experiment template. You can optionally update the experiment template before you specify it again in a new experiment.

## Tasks

- [Start an experiment](#)
- [View your experiments](#)
- [Tag an experiment](#)
- [Stop an experiment](#)
- [List resolved targets](#)

## Start an experiment

You start an experiment from an experiment template. For more information, see [Start an experiment from a template](#).

You can schedule your experiments as a one-time task or recurring tasks using Amazon EventBridge. For more information, see [Tutorial: Schedule a recurring experiment](#).

You can monitor your experiment using any of the following features:

- View your experiments in the AWS FIS console. For more information, see [View your experiments](#).

- View Amazon CloudWatch metrics for the target resources in your experiments or view AWS FIS usage metrics. For more information, see [Monitor using CloudWatch](#).
- Enable experiment logging to capture detailed information about your experiment as it runs. For more information see [Experiment logging](#).

## View your experiments

You can view the progress of a running experiment, and you can view experiments that have completed, stopped, or failed.

Stopped, completed, and failed experiments are automatically removed from your account after 120 days.

### To view experiments using the console

1. Open the AWS FIS console at <https://console.aws.amazon.com/fis/>.
2. In the navigation pane, choose **Experiments**.
3. Choose the **Experiment ID** of the experiment to open its details page.
4. Do one or more of the following:
  - Check **Details, State** for the [state of the experiment](#).
  - Choose the **Actions** tab for information about the experiment actions.
  - Choose the **Targets** tab for information about the experiment targets.
  - Choose the **Timeline** tab for a visual representation of the actions based on their start and end times.

### To view experiments using the CLI

Use the [list-experiments](#) command to get a list of experiments, and use the [get-experiment](#) command to get information about a specific experiment.

## Experiment states

An experiment can be in one of the following states:

- **pending** – The experiment is pending.
- **initiating** – The experiment is preparing to start.

- **running** – The experiment is running.
- **completed** – All actions in the experiment completed successfully.
- **stopping** – The stop condition was triggered or the experiment was stopped manually.
- **stopped** – All running or pending actions in the experiment are stopped.
- **failed** – The experiment failed due to an error, such as insufficient permissions or incorrect syntax.
- **cancelled** – The experimented was stopped or prevented from starting due to an engaged safety lever.

## Action states

An action can be in one of the following states:

- **pending** – The action is pending, either because the experiment hasn't started or the action is to start later in the experiment.
- **initiating** – The action is preparing to start.
- **running** – The action is running.
- **completed** – The action completed successfully.
- **cancelled** – The experiment stopped before the action started.
- **skipped** – The action has been skipped.
- **stopping** – The action is stopping.
- **stopped** – All running or pending actions in the experiment are stopped.
- **failed** – The action failed due to a client error, such as insufficient permissions or incorrect syntax.

## Tag an experiment

You can apply tags to experiments to help you organize them. You can also implement [tag-based IAM policies](#) to control access to experiments.

### To tag an experiment using the console

1. Open the AWS FIS console at <https://console.aws.amazon.com/fis/>.
2. In the navigation pane, choose **Experiments**.

3. Select the experiment and choose **Actions, Manage tags**.
4. To add a new tag, choose **Add new tag**, and specify a key and value.

To remove a tag, choose **Remove** for the tag.

5. Choose **Save**.

### To tag an experiment using the CLI

Use the [tag-resource](#) command.

## Stop an experiment

You can stop a running experiment at any time. When you stop an experiment, any post actions that have not been completed for an action are completed before the experiment stops. You cannot resume a stopped experiment.

### To stop an experiment using the console

1. Open the AWS FIS console at <https://console.aws.amazon.com/fis/>.
2. In the navigation pane, choose **Experiments**.
3. Select the experiment, and choose **Stop experiment**.
4. In the confirmation dialog box, choose **Stop experiment**.

### To stop an experiment using the CLI

Use the [stop-experiment](#) command.

## List resolved targets

You can view information for resolved targets for an experiment after target resolution has ended.

### To view resolved targets using the console

1. Open the AWS FIS console at <https://console.aws.amazon.com/fis/>.
2. In the navigation pane, choose **Experiments**.
3. Select the experiment, and choose **Report**.
4. View resolved target information under **Resources**.

## To view resolved targets using the CLI

Use the [list-experiment-resolved-targets](#) command.

# Tutorials for AWS Fault Injection Service

The following tutorials show you how to create and run experiments using AWS Fault Injection Service (AWS FIS).

## Tutorials

- [Tutorial: Test instance stop and start using AWS FIS](#)
- [Tutorial: Run CPU stress on an instance using AWS FIS](#)
- [Tutorial: Test Spot Instance interruptions using AWS FIS](#)
- [Tutorial: Simulate a connectivity event](#)
- [Tutorial: Schedule a recurring experiment](#)

## Tutorial: Test instance stop and start using AWS FIS

You can use AWS Fault Injection Service (AWS FIS) to test how your applications handle instance stop and start. Use this tutorial to create an experiment template that uses the AWS FIS `aws:ec2:stop-instances` action to stop one instance and then a second instance.

## Prerequisites

To complete this tutorial, ensure that you do the following:

- Launch two test EC2 instances in your account. After you launch your instances, note the IDs of both instances.
- Create an IAM role that enables the AWS FIS service to perform the `aws:ec2:stop-instances` action on your behalf. For more information, see [IAM roles for AWS FIS experiments](#).
- Ensure that you have access to AWS FIS. For more information, see [AWS FIS policy examples](#).

## Step 1: Create an experiment template

Create the experiment template using the AWS FIS console. In the template, you specify two actions that will run sequentially for three minutes each. The first action stops one of the test instances, which AWS FIS chooses randomly. The second action stops both test instances.

## To create an experiment template

1. Open the AWS FIS console at <https://console.aws.amazon.com/fis/>.
2. In the navigation pane, choose **Experiment templates**.
3. Choose **Create experiment template**.
4. For **Step 1, Specify template details**, do the following:
  - a. For **Description and name**, enter a description for the template, such as Amazon S3 Network Disrupt Connectivity.
  - b. Choose **Next**, and move to **Step 2, Specify actions and targets**.
5. For **Actions**, do the following:
  - a. Choose **Add action**.
  - b. Enter a name for the action. For example, enter **stopOneInstance**.
  - c. For **Action type**, choose **aws:ec2:stop-instances**.
  - d. For **Target** keep the target that AWS FIS creates for you.
  - e. For **Action parameters, Start instances after duration**, specify 3 minutes (PT3M).
  - f. Choose **Save**.
6. For **Targets**, do the following:
  - a. Choose **Edit** for the target that AWS FIS automatically created for you in the previous step.
  - b. Replace the default name with a more descriptive name. For example, enter **oneRandomInstance**.
  - c. Verify that **Resource type** is **aws:ec2:instance**.
  - d. For **Target method**, choose **Resource IDs**, and then choose the IDs of the two test instances.
  - e. For **Selection mode**, choose **Count**. For **Number of resources**, enter **1**.
  - f. Choose **Save**.
7. Choose **Add target** and do the following:
  - a. Enter a name for the target. For example, enter **bothInstances**.
  - b. For **Resource type**, choose **aws:ec2:instance**.
  - c. For **Target method**, choose **Resource IDs**, and then choose the IDs of the two test instances.

- d. For **Selection mode**, choose **All**.
  - e. Choose **Save**.
8. From the **Actions** section, choose **Add action**. Do the following:
- a. For **Name**, enter a name for the action. For example, enter **stopBothInstances**.
  - b. For **Action type**, choose **aws:ec2:stop-instances**.
  - c. For **Start after**, choose the first action that you added (**stopOneInstance**).
  - d. For **Target**, choose the second target that you added (**bothInstances**).
  - e. For **Action parameters**, **Start instances after duration**, specify 3 minutes (PT3M).
  - f. Choose **Save**.
9. Choose **Next** to move to **Step 3, Configure service access**.
10. For **Service Access**, choose **Use an existing IAM role**, and then choose the IAM role that you created as described in the prerequisites for this tutorial. If your role is not displayed, verify that it has the required trust relationship. For more information, see [the section called "Experiment role"](#).
11. Choose **Next** to move to **Step 4, Configure optional settings**.
12. (Optional) For **Tags**, choose **Add new tag** and specify a tag key and tag value. The tags that you add are applied to your experiment template, not the experiments that are run using the template.
13. Choose **Next** to move to **Step 5, Review and create**.
14. Review the template and choose **Create experiment template**. When prompted for confirmation, enter `create`, Then choose **Create experiment template**.

### (Optional) To view the experiment template JSON

Choose the **Export** tab. The following is an example of the JSON created by the preceding console procedure.

```
{
  "description": "Test instance stop and start",
  "targets": {
    "bothInstances": {
      "resourceType": "aws:ec2:instance",
      "resourceArns": [
        "arn:aws:ec2:region:123456789012:instance/instance_id_1",
        "arn:aws:ec2:region:123456789012:instance/instance_id_2"
      ]
    }
  }
}
```

```

    ],
    "selectionMode": "ALL"
  },
  "oneRandomInstance": {
    "resourceType": "aws:ec2:instance",
    "resourceArns": [
      "arn:aws:ec2:region:123456789012:instance/instance_id_1",
      "arn:aws:ec2:region:123456789012:instance/instance_id_2"
    ],
    "selectionMode": "COUNT(1)"
  }
},
"actions": {
  "stopBothInstances": {
    "actionId": "aws:ec2:stop-instances",
    "parameters": {
      "startInstancesAfterDuration": "PT3M"
    },
    "targets": {
      "Instances": "bothInstances"
    },
    "startAfter": [
      "stopOneInstance"
    ]
  },
  "stopOneInstance": {
    "actionId": "aws:ec2:stop-instances",
    "parameters": {
      "startInstancesAfterDuration": "PT3M"
    },
    "targets": {
      "Instances": "oneRandomInstance"
    }
  }
},
"stopConditions": [
  {
    "source": "none"
  }
],
"roleArn": "arn:aws:iam::123456789012:role/AllowFISEC2Actions",
"tags": {}
}

```

## Step 2: Start the experiment

When you have finished creating your experiment template, you can use it to start an experiment.

### To start an experiment

1. You should be on the details page for the experiment template that you just created. Otherwise, choose **Experiment templates** and then select the ID of the experiment template to open the details page.
2. Choose **Start experiment**.
3. (Optional) To add a tag to your experiment, choose **Add new tag** and enter a tag key and a tag value.
4. Choose **Start experiment**. When prompted for confirmation, enter **start** and choose **Start experiment**.

## Step 3: Track the experiment progress

You can track the progress of a running experiment until the experiment is completed, stopped, or failed.

### To track the progress of an experiment

1. You should be on the details page for the experiment that you just started. Otherwise, choose **Experiments** and then select the ID of the experiment to open the details page.
2. To view the state of the experiment, check **State** in the **Details** pane. For more information, see [experiment states](#).
3. When the state of the experiment is **Running**, go to the next step.

## Step 4: Verify the experiment result

You can verify that the instances were stopped and started by the experiment as expected.

### To verify the result of the experiment

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/> in a new browser tab or window. This allows you to continue to track the progress of the experiment in the AWS FIS console while viewing the result of the experiment in the Amazon EC2 console.

2. In the navigation pane, choose **Instances**.
3. When the state of the first action changes from **Pending** to **Running** (AWS FIS console), the state of one of the target instances changes from **Running** to **Stopped** (Amazon EC2 console).
4. After three minutes, the state of the first action changes to **Completed**, the state of the second action changes to **Running**, and the state of the other target instance changes to **Stopped**.
5. After three minutes, the state of the second action changes to **Completed**, the state of the target instances changes to **Running**, and the state of the experiment changes to **Completed**.

## Step 5: Clean up

If you no longer need the test EC2 instances that you created for this experiment, you can terminate them.

### To terminate the instances

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. In the navigation pane, choose **Instances**.
3. Select both test instances and choose **Instance state, Terminate instance**.
4. When prompted for confirmation, choose **Terminate**.

If you no longer need the experiment template, you can delete it.

### To delete an experiment template using the AWS FIS console

1. Open the AWS FIS console at <https://console.aws.amazon.com/fis/>.
2. In the navigation pane, choose **Experiment templates**.
3. Select the experiment template, and choose **Actions, Delete experiment template**.
4. When prompted for confirmation, enter **delete** and then choose **Delete experiment template**.

## Tutorial: Run CPU stress on an instance using AWS FIS

You can use AWS Fault Injection Service (AWS FIS) to test how your applications handle CPU stress. Use this tutorial to create an experiment template that uses AWS FIS to run a pre-configured

SSM document that runs CPU stress on an instance. The tutorial uses a stop condition to halt the experiment when the CPU utilization of the instance exceeds a configured threshold.

For more information, see [the section called “Pre-configured AWS FIS SSM documents”](#).

## Prerequisites

Before you can use AWS FIS to run CPU stress, complete the following prerequisites.

### Create an IAM role

Create a role and attach a policy that enables AWS FIS to use the `aws:ssm:send-command` action on your behalf. For more information, see [IAM roles for AWS FIS experiments](#).

### Verify access to AWS FIS

Ensure that you have access to AWS FIS. For more information, see [AWS FIS policy examples](#).

### Prepare a test EC2 instance

- Launch an EC2 instance using Amazon Linux 2 or Ubuntu, as required by the pre-configured SSM documents.
- The instance must be managed by SSM. To verify that the instance is managed by SSM, open the [Fleet Manager console](#). If the instance is not managed by SSM, verify that the SSM Agent is installed and that the instance has an attached IAM role with the **AmazonSSMManagedInstanceCore** policy. To verify the installed SSM Agent, connect to your instance and run the following command.

#### Amazon Linux 2

```
yum info amazon-ssm-agent
```

#### Ubuntu

```
apt list amazon-ssm-agent
```

- Enable detailed monitoring for the instance. This provides data in 1-minute periods, for an additional charge. Select the instance and choose **Actions, Monitor and troubleshoot, Manage detailed monitoring**.

## Step 1: Create a CloudWatch alarm for a stop condition

Configure a CloudWatch alarm so that you can stop the experiment if CPU utilization exceeds the threshold that you specify. The following procedure sets the threshold to 50% CPU utilization for the target instance. For more information, see [Stop conditions](#).

### To create an alarm that indicates when CPU utilization exceeds a threshold

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. In the navigation pane, choose **Instances**.
3. Select the target instance and choose **Actions, Monitor and troubleshoot, Manage CloudWatch alarms**.
4. For **Alarm notification**, use the toggle to turn off Amazon SNS notifications.
5. For **Alarm thresholds**, use the following settings:
  - **Group samples by: Maximum**
  - **Type of data to sample: CPU utilization**
  - **Percent: 50**
  - **Period: 1 Minute**
6. When you're done configuring the alarm, choose **Create**.

## Step 2: Create an experiment template

Create the experiment template using the AWS FIS console. In the template, you specify the following action to run: [aws:ssm:send-command/AWSFIS-Run-CPU-Stress](#).

### To create an experiment template

1. Open the AWS FIS console at <https://console.aws.amazon.com/fis/>.
2. In the navigation pane, choose **Experiment templates**.
3. Choose **Create experiment template**.
4. For **Step 1, Specify template details**, do the following:
  - a. For **Description and name**, enter a description for the template.
  - b. Choose **Next**, and move to **Step 2, Specify actions and targets**.
5. For **Actions**, do the following:

- a. Choose **Add action**.
  - b. Enter a name for the action. For example, enter **runCpuStress**.
  - c. For **Action type**, choose **aws:ssm:send-command/AWSFIS-Run-CPU-Stress**. This automatically adds the ARN of the SSM document to **Document ARN**.
  - d. For **Target** keep the target that AWS FIS creates for you.
  - e. For **Action parameters**, **Document parameters**, enter the following:

```
{"DurationSeconds": "120"}
```
  - f. For **Action parameters**, **Duration**, specify 5 minutes (PT5M).
  - g. Choose **Save**.
6. For **Targets**, do the following:
    - a. Choose **Edit** for the target that AWS FIS automatically created for you in the previous step.
    - b. Replace the default name with a more descriptive name. For example, enter **testInstance**.
    - c. Verify that **Resource type** is **aws:ec2:instance**.
    - d. For **Target method**, choose **Resource IDs**, and then choose the ID of the test instance.
    - e. For **Selection mode**, choose **All**.
    - f. Choose **Save**.
  7. Choose **Next** to move to **Step 3, Configure service access**.
  8. For **Service Access**, choose **Use an existing IAM role**, and then choose the IAM role that you created as described in the prerequisites for this tutorial. If your role is not displayed, verify that it has the required trust relationship. For more information, see [the section called "Experiment role"](#).
  9. Choose **Next** to move to **Step 4, Configure optional settings**.
  10. For **Stop conditions**, select the CloudWatch alarm that you created in Step 1.
  11. (Optional) For **Tags**, choose **Add new tag** and specify a tag key and tag value. The tags that you add are applied to your experiment template, not the experiments that are run using the template.
  12. Choose **Next** to move to **Step 5, Review and create**.
  13. Review the template and choose **Create experiment template**. When prompted for confirmation, enter `create`, Then choose **Create experiment template**.

## (Optional) To view the experiment template JSON

Choose the **Export** tab. The following is an example of the JSON created by the preceding console procedure.

```
{
  "description": "Test CPU stress predefined SSM document",
  "targets": {
    "testInstance": {
      "resourceType": "aws:ec2:instance",
      "resourceArns": [
        "arn:aws:ec2:region:123456789012:instance/instance_id"
      ],
      "selectionMode": "ALL"
    }
  },
  "actions": {
    "runCpuStress": {
      "actionId": "aws:ssm:send-command",
      "parameters": {
        "documentArn": "arn:aws:ssm:region::document/AWSFIS-Run-CPU-Stress",
        "documentParameters": "{\"DurationSeconds\": \"120\"}",
        "duration": "PT5M"
      },
      "targets": {
        "Instances": "testInstance"
      }
    }
  },
  "stopConditions": [
    {
      "source": "aws:cloudwatch:alarm",
      "value": "arn:aws:cloudwatch:region:123456789012:alarm:awsec2-instance_id-
GreaterThanOrEqualToThreshold-CPUUtilization"
    }
  ],
  "roleArn": "arn:aws:iam::123456789012:role/AllowFISSSMActions",
  "tags": {}
}
```

## Step 3: Start the experiment

When you have finished creating your experiment template, you can use it to start an experiment.

## To start an experiment

1. You should be on the details page for the experiment template that you just created. Otherwise, choose **Experiment templates** and then select the ID of the experiment template to open the details page.
2. Choose **Start experiment**.
3. (Optional) To add a tag to your experiment, choose **Add new tag** and enter a tag key and a tag value.
4. Choose **Start experiment**. When prompted for confirmation, enter **start**. Choose **Start experiment**.

## Step 4: Track the experiment progress

You can track the progress of a running experiment until the experiment completes, stops, or fails.

### To track the progress of an experiment

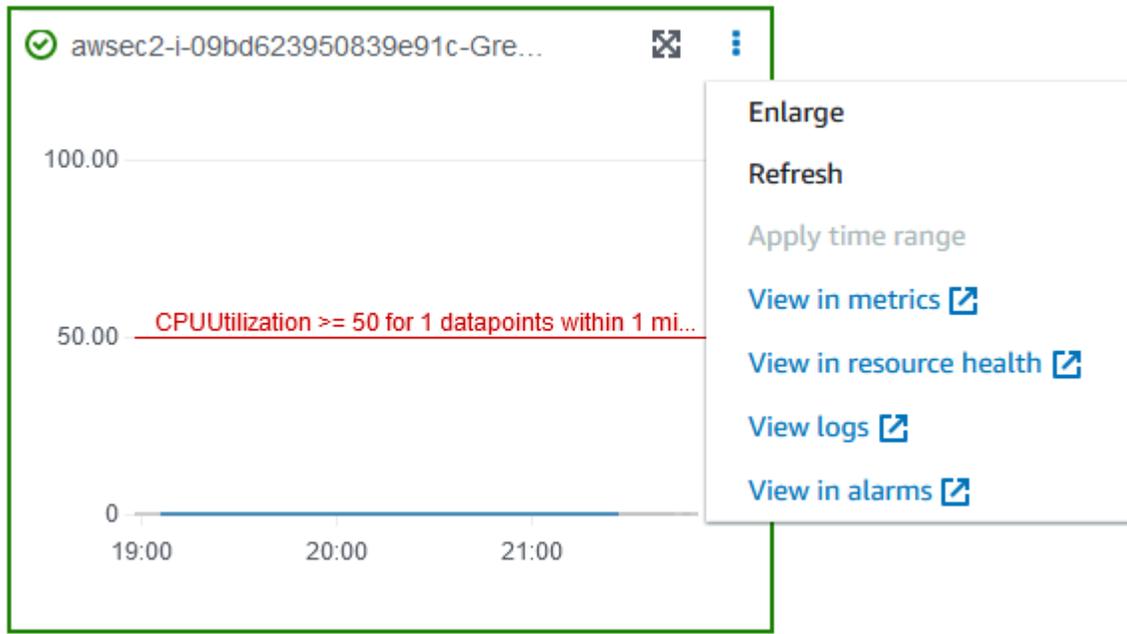
1. You should be on the details page for the experiment that you just started. Otherwise, choose **Experiments** and then select the ID of the experiment to open the details page for the experiment.
2. To view the state of the experiment, check **State** in the **Details** pane. For more information, see [experiment states](#).
3. When the experiment state is **Running**, go to the next step.

## Step 5: Verify the experiment results

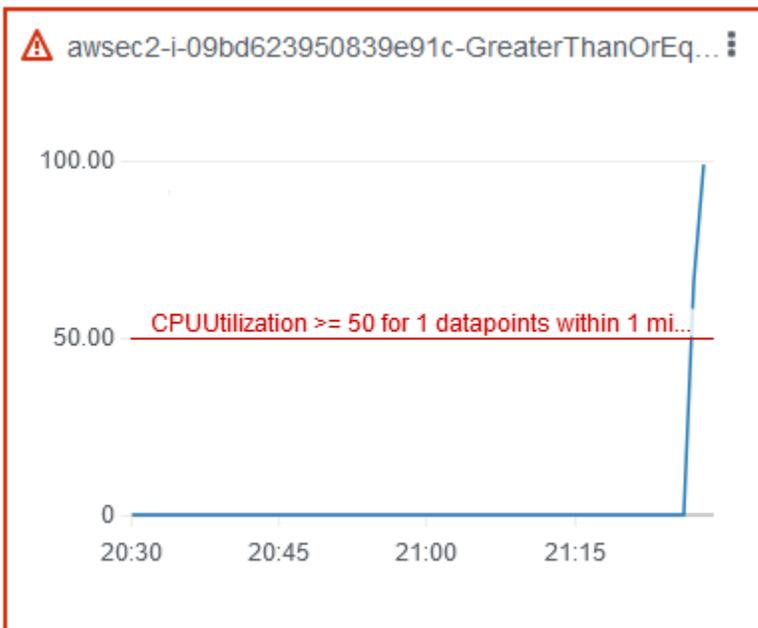
You can monitor the CPU utilization of your instance while the experiment is running. When the CPU utilization reaches the threshold, the alarm is triggered and the experiment is halted by the stop condition.

### To verify the results of the experiment

1. Choose the **Stop conditions** tab. The green border and green checkmark icon indicate that the initial state of the alarm is OK. The red line indicates the alarm threshold. If you prefer a more detailed graph, choose **Enlarge** from the widget menu.



- When CPU utilization exceeds the threshold, the red border and red exclamation point icon in the **Stop conditions** tab indicate that the alarm state changed to ALARM. In the **Details** pane, the state of the experiment is **Stopped**. If you select the state, the message displayed is "Experiment halted by stop condition".



- When CPU utilization decreases below the threshold, the green border and green checkmark icon indicate that the alarm state changed to OK.

4. (Optional) Choose **View in alarms** from the widget menu. This opens the alarm details page in the CloudWatch console, where you can get more detail about the alarm or edit the alarm settings.

## Step 6: Clean up

If you no longer need the test EC2 instance that you created for this experiment, you can terminate it.

### To terminate the instance

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. In the navigation pane, choose **Instances**.
3. Select the test instances and choose **Instance state, Terminate instance**.
4. When prompted for confirmation, choose **Terminate**.

If you no longer need the experiment template, you can delete it.

### To delete an experiment template using the AWS FIS console

1. Open the AWS FIS console at <https://console.aws.amazon.com/fis/>.
2. In the navigation pane, choose **Experiment templates**.
3. Select the experiment template, and choose **Actions, Delete experiment template**.
4. When prompted for confirmation, enter **delete** and then choose **Delete experiment template**.

## Tutorial: Test Spot Instance interruptions using AWS FIS

Spot Instances use spare EC2 capacity that is available, for up to a 90% discount compared to On-Demand pricing. However, Amazon EC2 can interrupt your Spot Instances when it needs the capacity back. When using Spot Instances, you must be prepared for potential interruptions. For more information, see [Spot Instance interruptions](#) in the *Amazon EC2 User Guide*.

You can use AWS Fault Injection Service (AWS FIS) to test how your applications handle a Spot Instance interruption. Use this tutorial to create an experiment template that uses the AWS

FIS `aws:ec2:send-spot-instance-interruptions` action to interrupt one of your Spot Instances.

Alternatively, to initiate the experiment using the Amazon EC2 console, see [Initiate a Spot Instance interruption](#) in the *Amazon EC2 User Guide*.

## Prerequisites

Before you can use AWS FIS to interrupt a Spot Instance, complete the following prerequisites.

### 1. Create an IAM role

Create a role and attach a policy that enables AWS FIS to perform the `aws:ec2:send-spot-instance-interruptions` action on your behalf. For more information, see [IAM roles for AWS FIS experiments](#).

### 2. Verify access to AWS FIS

Ensure that you have access to AWS FIS. For more information, see [AWS FIS policy examples](#).

### 3. (Optional) Create a Spot Instance request

If you'd like a new Spot Instance to use for this experiment, use the [run-instances](#) command to request a Spot Instance. The default is to terminate Spot Instances that are interrupted. If you set the interruption behavior to `stop`, you must also set the type to `persistent`. For this tutorial, do not set the interruption behavior to `hibernate`, as the hibernation process begins immediately.

```
aws ec2 run-instances \  
  --image-id ami-0ab193018fEXAMPLE \  
  --instance-type "t2.micro" \  
  --count 1 \  
  --subnet-id subnet-1234567890abcdef0 \  
  --security-group-ids sg-111222333444aaab \  
  --instance-market-options file://spot-options.json \  
  --query Instances[*].InstanceId
```

The following is an example of the `spot-options.json` file.

```
{  
  "MarketType": "spot",  
  "SpotOptions": {
```

```
    "SpotInstanceType": "persistent",
    "InstanceInterruptionBehavior": "stop"
  }
}
```

The `--query` option in the example command makes it so that the command returns only the instance ID of the Spot Instance. The following is example output.

```
[
  "i-0abcdef1234567890"
]
```

#### 4. Add a tag so that AWS FIS can identify the target Spot Instance

Use the [create-tags](#) command to add the tag `Name=interruptMe` to your target Spot Instance.

```
aws ec2 create-tags \
  --resources i-0abcdef1234567890 \
  --tags Key=Name,Value=interruptMe
```

## Step 1: Create an experiment template

Create the experiment template using the AWS FIS console. In the template, you specify the action that will run. The action interrupts the Spot Instance with the specified tag. If there is more than one Spot Instance with the tag, AWS FIS chooses one of them at random.

### To create an experiment template

1. Open the AWS FIS console at <https://console.aws.amazon.com/fis/>.
2. In the navigation pane, choose **Experiment templates**.
3. Choose **Create experiment template**.
4. For **Step 1, Specify template details**, do the following:
  - a. For **Description and name**, enter a description and a name for the template.
  - b. Choose **Next**, and move to **Step 2, Specify actions and targets**.
5. For **Actions**, do the following:
  - a. Choose **Add action**.

- b. Enter a name for the action. For example, enter **interruptSpotInstance**.
  - c. For **Action type**, choose **aws:ec2:send-spot-instance-interruptions**.
  - d. For **Target** keep the target that AWS FIS creates for you.
  - e. For **Action parameters**, **Duration before interruption**, specify 2 Minutes (PT2M).
  - f. Choose **Save**.
6. For **Targets**, do the following:
  - a. Choose **Edit** for the target that AWS FIS automatically created for you in the previous step.
  - b. Replace the default name with a more descriptive name. For example, enter **oneSpotInstance**.
  - c. Verify that **Resource type** is **aws:ec2:spot-instance**.
  - d. For **Target method**, choose **Resource tags, filters, and parameters**.
  - e. For **Resource tags**, choose **Add new tag**, and enter the tag key and tag value. Use the tag that you added to the Spot Instance to interrupt, as described in the Prerequisites for this tutorial.
  - f. For **Resource filters** choose **Add new filter** and enter **State.Name** as the path and **running** as the value.
  - g. For **Selection mode**, choose **Count**. For **Number of resources**, enter **1**.
  - h. Choose **Save**.
7. Choose **Next** to move to **Step 3, Configure service access**.
8. For **Service Access**, choose **Use an existing IAM role**, and then choose the IAM role that you created as described in the prerequisites for this tutorial. If your role is not displayed, verify that it has the required trust relationship. For more information, see [the section called "Experiment role"](#).
9. Choose **Next** to move to **Step 4, Configure optional settings**.
10. (Optional) For **Tags**, choose **Add new tag** and specify a tag key and tag value. The tags that you add are applied to your experiment template, not the experiments that are run using the template.
11. Choose **Next** to move to **Step 5, Review and create**.
12. Review the template and choose **Create experiment template**. When prompted for confirmation, enter `create`, Then choose **Create experiment template**.

### (Optional) To view the experiment template JSON

Choose the **Export** tab. The following is an example of the JSON created by the preceding console procedure.

```
{
  "description": "Test Spot Instance interruptions",
  "targets": {
    "oneSpotInstance": {
      "resourceType": "aws:ec2:spot-instance",
      "resourceTags": {
        "Name": "interruptMe"
      },
      "filters": [
        {
          "path": "State.Name",
          "values": [
            "running"
          ]
        }
      ],
      "selectionMode": "COUNT(1)"
    }
  },
  "actions": {
    "interruptSpotInstance": {
      "actionId": "aws:ec2:send-spot-instance-interruptions",
      "parameters": {
        "durationBeforeInterruption": "PT2M"
      },
      "targets": {
        "SpotInstances": "oneSpotInstance"
      }
    }
  },
  "stopConditions": [
    {
      "source": "none"
    }
  ],
  "roleArn": "arn:aws:iam::123456789012:role/AllowFISSpotInterruptionActions",
  "tags": {
    "Name": "my-template"
  }
}
```

## Step 2: Start the experiment

When you have finished creating your experiment template, you can use it to start an experiment.

### To start an experiment

1. You should be on the details page for the experiment template that you just created. Otherwise, choose **Experiment templates** and then select the ID of the experiment template to open the details page.
2. Choose **Start experiment**.
3. (Optional) To add a tag to your experiment, choose **Add new tag** and enter a tag key and a tag value.
4. Choose **Start experiment**. When prompted for confirmation, enter **start** and choose **Start experiment**.

## Step 3: Track the experiment progress

You can track the progress of a running experiment until the experiment is completed, stopped, or failed.

### To track the progress of an experiment

1. You should be on the details page for the experiment that you just started. Otherwise, choose **Experiments** and then select the ID of the experiment to open the details page.
2. To view the state of the experiment, check **State** in the **Details** pane. For more information, see [experiment states](#).
3. When the state of the experiment is **Running**, go to the next step.

## Step 4: Verify the experiment result

When the action for this experiment is completed, the following occurs:

- The target Spot Instance receives an [instance rebalance recommendation](#).
- A [Spot Instance interruption notice](#) is issued two minutes before Amazon EC2 terminates or stops your instance.
- After two minutes, the Spot Instance is terminated or stopped.

- A Spot Instance that was stopped by AWS FIS remains stopped until you restart it.

### To verify that the instance was interrupted by the experiment

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. From the navigation pane, open **Spot Requests** and **Instances** in separate browser tabs or windows.
3. For **Spot Requests**, select the Spot Instance request. The initial status is fulfilled. After the experiment completes, the status changes as follows:
  - `terminate` - The status changes to `instance-terminated-by-experiment`.
  - `stop` - The status changes to `marked-for-stop-by-experiment` and then `instance-stopped-by-experiment`.
4. For **Instances**, select the Spot Instance. The initial status is `Running`. Two minutes after you receive the Spot Instance interruption notice, the status changes as follows:
  - `stop` - The status changes to `Stopping` and then `Stopped`.
  - `terminate` - The status changes to `Shutting-down` and then `Terminated`.

## Step 5: Clean up

If you created the test Spot Instance for this experiment with an interruption behavior of `stop` and you no longer need it, you can cancel the Spot Instance request and terminate the Spot Instance.

### To cancel the request and terminate the instance using the AWS CLI

1. Use the [cancel-spot-instance-requests](#) command to cancel the Spot Instance request.

```
aws ec2 cancel-spot-instance-requests --spot-instance-request-ids sir-ksie869j
```

2. Use the [terminate-instances](#) command to terminate the instance.

```
aws ec2 terminate-instances --instance-ids i-0abcdef1234567890
```

If you no longer need the experiment template, you can delete it.

## To delete an experiment template using the AWS FIS console

1. Open the AWS FIS console at <https://console.aws.amazon.com/fis/>.
2. In the navigation pane, choose **Experiment templates**.
3. Select the experiment template, and choose **Actions, Delete experiment template**.
4. When prompted for confirmation, enter **delete** and then choose **Delete experiment template**.

## Tutorial: Simulate a connectivity event

You can use AWS Fault Injection Service (AWS FIS) to simulate a variety of connectivity events. AWS FIS simulates connectivity events by blocking network connections in one of the following ways:

- `all` – Denies all traffic entering and leaving the subnet. Note that this option allows intra-subnet traffic, including traffic to and from network interfaces in the subnet.
- `availability-zone` – Denies intra-VPC traffic to and from subnets in other Availability Zones.
- `dynamodb` – Denies traffic to and from the Regional endpoint for DynamoDB in the current Region.
- `prefix-list` – Denies traffic to and from the specified prefix list.
- `s3` – Denies traffic to and from the Regional endpoint for Amazon S3 in the current Region.
- `s3express` – Denies traffic to and from the zonal endpoint for Amazon S3 Express One Zone in the target subnets' AZ. Target subnets must reside in AZs where S3 Express One Zone is currently available. For more information, see [S3 Express One Zone Availability Zones and Regions](#).
- `vpc` – Denies traffic entering and leaving the VPC.

Use this tutorial to create an experiment template that uses the AWS FIS

`aws:network:disrupt-connectivity` action to introduce connectivity loss with Amazon S3 in a target subnet.

### Topics

- [Prerequisites](#)
- [Step 1: Create an AWS FIS experiment template](#)
- [Step 2: Ping an Amazon S3 endpoint](#)
- [Step 3: Start your AWS FIS experiment](#)

- [Step 4: Track your AWS FIS experiment progress](#)
- [Step 5: Verify Amazon S3 network disruption](#)
- [Step 5: Clean up](#)

## Prerequisites

Before beginning this tutorial, you need a role with the appropriate permissions in your AWS account, and a test Amazon EC2 instance:

### A role with permissions in your AWS account

Create a role and attach a policy that enables AWS FIS to perform the `aws:network:disrupt-connectivity` action on your behalf.

Your IAM role requires the following policy:

- [AWSFaultInjectionSimulatorNetworkAccess](#) – Grants AWS FIS service permission in Amazon EC2 networking and other required services to perform AWS FIS actions related to network infrastructure.

#### Note

For simplicity, this tutorial uses an AWS managed policy. For production use, we recommend that you instead grant only the minimum permissions necessary for your use case.

For more information about how to create an IAM role, see [IAM roles for AWS FIS experiments \(AWS CLI\)](#) or [Creating an IAM role \(console\)](#) in the *IAM User Guide*.

### A test Amazon EC2 instance

Launch and connect to a test Amazon EC2 instance. You can use the following tutorial to launch and connect to an Amazon EC2 instance: [Tutorial: Get started with Amazon EC2 Linux instances](#) in the *Amazon EC2 User Guide*.

## Step 1: Create an AWS FIS experiment template

Create the experiment template by using the AWS FIS AWS Management Console. An AWS FIS template is made up of actions, targets, stop conditions, and an experiment role. For more information about how the templates work, see [Experiment templates for AWS FIS](#).

Before you begin, make sure you have the following ready:

- An IAM role with the correct permissions.
- An Amazon EC2 instance.
- The subnet ID of your Amazon EC2 instance.

### To create an experiment template

1. Open the AWS FIS console at <https://console.aws.amazon.com/fis/>.
2. In the left navigation pane, choose **Experiment templates**.
3. Choose **Create experiment template**.
4. For **Step 1, Specify template details**, do the following:
  - a. For **Description and name**, enter a description for the template, such as Amazon S3 Network Disrupt Connectivity.
  - b. Choose **Next**, and move to **Step 2, Specify actions and targets**.
5. Under **Actions**, choose **Add action**.
  - a. For the **Name**, enter `disruptConnectivity`.
  - b. For **Action type**, select `aws:network:disrupt-connectivity`.
  - c. Under **Action parameters**, set the **Duration** to 2 minutes.
  - d. Under **Scope**, select `s3`.
  - e. At the top, choose **Save**.
6. Under **Targets**, you should see the target that has been created automatically. Choose **Edit**.
  - a. Verify that **Resource type** is `aws:ec2:subnet`.
  - b. Under **Target method**, select **Resource IDs**, and then choose the subnet that you used when creating your Amazon EC2 instance in the [Prerequisites](#) steps.
  - c. Verify that **Selection mode** is **All**.

- d. Choose **Save**.
7. Choose **Next** to move to **Step 3, Configure service access**.
8. Under **Service Access**, select the IAM role that you created as described in the [Prerequisites](#) for this tutorial. If your role is not displayed, verify that it has the required trust relationship. For more information, see [the section called "Experiment role"](#).
9. Choose **Next** to move to **Step 4, Configure optional settings**.
10. (Optional) Under **Stop conditions**, you can select a CloudWatch alarm to stop the experiment if the condition occurs. For more information, see [Stop conditions for AWS FIS](#).
11. (Optional) Under **Logs**, you can select an Amazon S3 bucket, or send logs to CloudWatch for your experiment.
12. Choose **Next** to move to **Step 5, Review and create**.
13. Review the template and choose **Create experiment template**. When prompted for confirmation, enter `create`, Then choose **Create experiment template**.

## Step 2: Ping an Amazon S3 endpoint

Verify that your Amazon EC2 instance is able to reach an Amazon S3 endpoint.

1. Connect to the Amazon EC2 instance that you created in the [Prerequisites](#) steps.

For troubleshooting, see [Troubleshoot connecting to your instance](#) in the *Amazon EC2 User Guide*.

2. Check to see the AWS Region where your instance is located. You can do this in the Amazon EC2 console or by running the following command.

```
hostname
```

For example, if you launched an Amazon EC2 instance in `us-west-2`, you'll see the following output.

```
[ec2-user@ip-172.16.0.0 ~]$ hostname  
ip-172.16.0.0.us-west-2.compute.internal
```

3. Ping an Amazon S3 endpoint in your AWS Region. Replace *AWS Region* with your Region.

```
ping -c 1 s3.AWS Region.amazonaws.com
```

For the output, you should see a successful ping with 0% packet loss, as shown in the following example.

```
PING s3.us-west-2.amazonaws.com (x.x.x.x) 56(84) bytes of data.  
64 bytes from s3-us-west-2.amazonaws.com (x.x.x.x: icmp_seq=1 ttl=249 time=1.30 ms  
  
--- s3.us-west-2.amazonaws.com ping statistics ---  
1 packets transmitted, 1 received, 0% packet loss, time 0ms  
rtt min/avg/max/mdev = 1.306/1.306/1.306/0.000 ms
```

### Step 3: Start your AWS FIS experiment

Start an experiment with the experiment template that you just created.

1. Open the AWS FIS console at <https://console.aws.amazon.com/fis/>.
2. In the left navigation pane, choose **Experiment templates**.
3. Select the ID of the experiment template that you created to open its details page.
4. Choose **Start experiment**.
5. (Optional) In the confirmation page, add tags for your experiment.
6. In the confirmation page, choose **Start experiment**.

### Step 4: Track your AWS FIS experiment progress

You can track the progress of a running experiment until the experiment is completed, stopped, or has failed.

1. You should be on the details page for the experiment that you just started. If you're not, choose **Experiments**, and then select the ID of the experiment to open its details page.
2. To view the state of the experiment, check the **State** in the details pane. For more information, see [Experiment states](#).
3. When the state of the experiment is **Running**, move to the next step.

## Step 5: Verify Amazon S3 network disruption

You can validate the experiment progress by by pinging the Amazon S3 endpoint.

- From your Amazon EC2 instance, ping the Amazon S3 endpoint in your AWS Region. Replace *AWS Region* with your Region.

```
ping -c 1 s3.AWS Region.amazonaws.com
```

For the output, you should see an unsuccessful ping with 100% packet loss, as shown in the following example.

```
ping -c 1 s3.us-west-2.amazonaws.com
PING s3.us-west-2.amazonaws.com (x.x.x.x) 56(84) bytes of data.

--- s3.us-west-2.amazonaws.com ping statistics ---
1 packets transmitted, 0 received, 100% packet loss, time 0ms
```

## Step 5: Clean up

If you no longer need the Amazon EC2 instance that you created for this experiment or the AWS FIS template, you can remove them.

### To remove the Amazon EC2 instance

- Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
- In the navigation pane, choose **Instances**.
- Select the test instance, choose **Instance state**, and then choose **Terminate instance**.
- When prompted for confirmation, choose **Terminate**.

### To delete the experiment template using the AWS FIS console

- Open the AWS FIS console at <https://console.aws.amazon.com/fis/>.
- In the navigation pane, choose **Experiment templates**.
- Select the experiment template, and then choose **Actions, Delete experiment template**.

4. When prompted for confirmation, enter `delete`, and then choose **Delete experiment template**.

## Tutorial: Schedule a recurring experiment

With AWS Fault Injection Service (AWS FIS), you can perform fault injection experiments on your AWS workloads. These experiments run on templates that contain one or more actions to run on specified targets. When you also use Amazon EventBridge, you can schedule your experiments as a one-time task or recurring tasks.

Use this tutorial to create an EventBridge schedule that runs an AWS FIS experiment template every 5 minutes.

### Tasks

- [Prerequisites](#)
- [Step 1: Create an IAM role and policy](#)
- [Step 2: Create an Amazon EventBridge Scheduler](#)
- [Step 3: Verify your experiment](#)
- [Step 4: Clean up](#)

### Prerequisites

Before beginning this tutorial, you must have an AWS FIS experiment template that you want to run on a schedule. If you already have a working experiment template, make note of the template ID and AWS Region. Otherwise, you can create a template by following the instructions in [the section called "Test instance stop and start"](#), and then return to this tutorial.

## Step 1: Create an IAM role and policy

### To create an IAM role and policy

1. Open the IAM console at <https://console.aws.amazon.com/iam/>.
2. In the left navigation pane, choose **Roles**, and then **Create Role**.
3. Choose **Custom trust policy**, and then insert the following snippet to allow Amazon EventBridge Scheduler to assume the role on your behalf.

## JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Principal": {
        "Service": "scheduler.amazonaws.com"
      },
      "Action": "sts:AssumeRole"
    }
  ]
}
```

Choose **Next**.

4. Under **Add permissions**, choose **Create policy**.
5. Choose **JSON**, and then insert the following policy. Replace the *your-experiment-template-id* value with the template ID of your experiment from the Prerequisites steps.

## JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": "fis:StartExperiment",
      "Resource": [
        "arn:aws:fis::*:experiment-template/your-experiment-template-id",
        "arn:aws:fis::*:experiment/*"
      ]
    }
  ]
}
```

You can restrict the scheduler to only run AWS FIS experiment templates that have a specific tag value. For example, the following policy grants the `StartExperiment` permission for all AWS FIS experiments, but restricts the scheduler to only run experiment templates that are tagged `Purpose=Schedule`.

JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": "fis:StartExperiment",
      "Resource": "arn:aws:fis:*:*:experiment/*"
    },
    {
      "Effect": "Allow",
      "Action": "fis:StartExperiment",
      "Resource": "arn:aws:fis:*:*:experiment-template/*",
      "Condition": {
        "StringEquals": {
          "aws:ResourceTag/Purpose": "Schedule"
        }
      }
    }
  ]
}
```

Choose **Next: Tags**.

6. Choose **Next: Review**.
7. Under **Review policy**, name your policy `FIS_RecurringExperiment`, and then choose **Create policy**.
8. Under **Add permissions**, add the new `FIS_RecurringExperiment` policy to your role, and then choose **Next**.
9. Under **Name, review, and create**, name the role `FIS_RecurringExperiment_role`, and then choose **Create role**.

## Step 2: Create an Amazon EventBridge Scheduler

### To create an Amazon EventBridge Scheduler

1. Open the Amazon EventBridge console at <https://console.aws.amazon.com/events/>.
2. In the left navigation pane, choose **Schedules**.
3. Verify that you are in the same AWS Region as your AWS FIS experiment template.
4. Choose **Create schedule**, and fill in the following:
  - Under **Schedule name**, insert `FIS_recurring_experiment_tutorial`.
  - Under **Schedule pattern**, select **Recurring schedule**.
  - Under **Schedule type**, select **Rate-based schedule**.
  - Under **Rate expression**, choose **5 minutes**.
  - Under **Flexible time window**, select **Off**.
  - (Optional) Under **Timeframe**, select your time zone.
  - Choose **Next**.
5. Under **Select target**, choose **All APIs**, and then search for **AWS FIS**.
6. Choose **AWS FIS**, and then select **StartExperiment**.
7. Under **Input**, insert the following JSON payload. Replace the *your-experiment-template-id* value with the template ID of your experiment. The `ClientToken` is a unique identifier for the scheduler. In this tutorial, we are using a context keyword allowed by Amazon EventBridge Scheduler. For more information, see [Adding context attributes](#) in the *Amazon EventBridge User Guide*.

```
{
  "ClientToken": "<aws.scheduler.execution-id>",
  "ExperimentTemplateId": "your-experiment-template-id"
}
```

Choose **Next**.

8. (Optional) Under **Settings**, you can set the **Retry policy**, **Dead-letter queue (DLQ)**, and **Encryption** settings. Alternatively, you can keep the default values.
9. Under **Permissions**, select **Use existing role**, and then search for `FIS_RecurringExperiment_role`.

10. Choose **Next**.
11. Under **Review and create schedule**, review your scheduler details, and then choose **Create schedule**.

## Step 3: Verify your experiment

### To verify that your AWS FIS experiment ran on schedule

1. Open the AWS FIS console at <https://console.aws.amazon.com/fis/>.
2. In the left navigation pane, choose **Experiments**.
3. Five minutes after you create your schedule, you should see your experiment running.

## Step 4: Clean up

### To disable your Amazon EventBridge Scheduler

1. Open the Amazon EventBridge console at <https://console.aws.amazon.com/events/>.
2. In the left navigation pane, choose **Schedules**.
3. Select your newly created scheduler, and then choose **Disable**.

# Working with the AWS FIS scenario library

Scenarios define events or conditions that customers can apply to test the resiliency of their applications, such as the interruption of compute resources on which the application is running. Scenarios are created and owned by AWS, and minimize undifferentiated heavy lifting by providing you with a group of pre-defined targets and fault actions (e.g., stopping 30% of instances in an autoscaling group) for common application impairments.

Scenarios are provided through a console-only scenario library and run using an AWS FIS experiment template. In order to run an experiment using a scenario, you will select the scenario from the library, specify parameters matching your workload details, and save it as an experiment template in your account.

## Topics

- [Viewing a scenario](#)
- [Using a scenario](#)
- [Exporting a scenario](#)
- [Scenarios reference](#)

## Viewing a scenario

To view a scenario using the console:

1. Open the AWS FIS console at <https://console.aws.amazon.com/fis/>.
2. In the navigation pane, choose **Scenario library**.
3. To view information about a specific scenario, select the scenario card to bring up a split panel.
  - In the **Description** tab in the split panel at the bottom of the page, you can view a short description of the scenario. You can also find a short summary of pre-requisites containing a summary of the target resources required and any actions you need to take to prepare the resources for use with the scenario. Finally you can also see additional information about the targets and actions in the scenario as well as the anticipated duration when the experiment runs successfully with default settings.
  - In the **Content** tab in the split panel at the bottom of the page, you can preview a partially populated version of the experiment template that will be created from the scenario.

- In the **Details** tab in the split panel at the bottom of the page, you can find a detailed explanation how the scenario is implemented. This may contain detailed information about how individual aspects of the scenario are approximated. Where applicable you can also read about what metrics to use as stop conditions and to provide observability to learn from the experiment. Finally you will find recommendations how to expand the resulting experiment template.

## Using a scenario

To use a scenario using the console:

1. Open the AWS FIS console at <https://console.aws.amazon.com/fis/>.
2. In the navigation pane, choose **Scenario library**.
3. To view information about a specific scenario, select the scenario card to bring up a split panel
4. To use the scenario, select the scenario card and choose **Create template with scenario**.
5. In the **Create experiment template** view fill in any missing items.
  - a. Some scenarios allow you to edit parameters that are shared across multiple actions or targets. This functionality will be disabled once you make any changes to the scenario, including changes by the shared parameter editing. To use this feature select the **Edit shared parameters** button. Edit parameters in the modal and select the **Save** button.
  - b. Some experiment templates may have missing action or target parameters, highlighted on each action and target card. Select the **Edit** button for each card, add the missing information, and select the **Save** button on the card.
  - c. All templates require a **Service access** execution role. You can choose an existing role or create a new role for this experiment template.
  - d. We recommend defining one or more optional **Stop conditions** by selecting an existing AWS CloudWatch alarm. Learn more about [Stop conditions for AWS FIS](#). If you don't have an alarm configured yet, you can follow the instructions at [Using Amazon CloudWatch Alarms](#) and update the experiment template later.
  - e. We recommend enabling optional experiment **Logs** to Amazon CloudWatch logs or to an Amazon S3 bucket. Learn more about [Experiment logging for AWS FIS](#). If you don't have appropriate resources configured yet, you can update the experiment template later.
6. In the **Create experiment template** select **Create experiment template**.

7. From the **Experiment templates** view of the AWS FIS console select **Start experiment**. Learn more about [Managing AWS FIS experiment templates](#).

## Exporting a scenario

Scenarios are a console-only experience. While similar to experiment templates, scenarios are not complete experiment templates and can not be directly imported into AWS FIS. If you wish to use scenarios as part of your own automation, you can use one of two paths:

1. Follow the steps in [Using a scenario](#) to create a valid AWS FIS experiment template and export that template.
2. Follow the steps in [Viewing a scenario](#) and in step 3, from the **Content** tab, copy and save the scenario content, then add missing parameters manually to create a valid experiment template.

## Scenarios reference

Scenarios included in the scenario library are designed to use [tags](#) where possible and each scenario describes the required tags in the **Prerequisites** and **How it works** sections of the scenario description. You can tag your resources with those pre-defined tags or you can set your own tags using the shared parameter editing experience (see [Using a scenario](#)).

This reference describes the common scenarios in the AWS FIS scenario library. You can also list the supported scenarios using the AWS FIS console.

For more information, see [Working with the AWS FIS scenario library](#).

AWS FIS supports the following Amazon EC2 scenarios. These scenarios target instances using [tags](#). You can use your own tags or use the default tags included in the scenario. Some of these scenarios [use SSM documents](#).

- **EC2 stress: instance failure** - Explore the effect of instance failure by stopping one or more EC2 instances.

Target instances in the current region that have a specific tag attached. In this scenario we will stop those instances and restart them at the end of the action duration, by default 5 min.

- **EC2 stress: Disk** - Explore impact of increased disk utilization on your EC2 based application.

In this scenario we will target EC2 instances in the current region that have a specific tag attached. In this scenario you can customize an increasing amount disk utilization injected on targeted EC2 instances for the action duration, by default 5 min for each disk stress action.

- **EC2 stress: CPU** - Explore impact of increased CPU on your EC2 based application.

In this scenario we will target EC2 instances in the current region that have a specific tag attached. In this scenario you can customize an increasing amount of CPU stress injected on targeted EC2 instances for the action duration, by default 5 min for each CPU stress action.

- **EC2 stress: Memory** - Explore impact of increased memory utilization on your EC2 based application.

In this scenario we will target EC2 instances in the current region that have a specific tag attached. In this scenario you can customize an increasing amount of memory stress injected on targeted EC2 instances for the action duration, by default 5 min for each memory stress action.

- **EC2 stress: Network Latency** - Explore impact of increased network latency on your EC2 based application.

In this scenario we will target EC2 instances in the current region that have a specific tag attached. In this scenario you can customize an increasing amount of network latency injected on targeted EC2 instances for the action duration, by default 5 min for each latency action.

AWS FIS supports the following Amazon EKS scenarios. These scenarios target EKS pods using a Kubernetes application labels. You can use your own labels or use the default labels included in the scenario. For more information about EKS with FIS, see [EKS Pod actions](#).

- **EKS stress: Pod Delete** - Explore the effect of EKS pod failure by deleting one or more pods.

In this scenario we will target pods in the current region that are associated with an application label. In this scenario we will terminate all matched pods. Re-creation of pods will be controlled by kubernetes configuration.

- **EKS stress: CPU** - Explore impact of increased CPU on your EKS based application.

In this scenario we will target pods in the current region that are associated with an application label. In this scenario you can customize an increasing amount of CPU stress injected on targeted EKS pods for the action duration, by default 5 min for each CPU stress action.

- **EKS stress: Disk** - Explore impact of increased disk utilization on your EKS based application.

In this scenario we will target pods in the current region that are associated with an application label. In this scenario you can customize an increasing amount of disk stress injected on targeted EKS pods for the action duration, by default 5 min for each CPU stress action.

- **EKS stress: Memory** - Explore impact of increased memory utilization on your EKS based application.

In this scenario we will target pods in the current region that are associated with an application label. In this scenario you can customize an increasing amount of memory stress injected on targeted EKS pods for the action duration, by default 5 min for each memory stress action.

- **EKS stress: Network latency** - Explore impact of increased network latency on your EKS based application.

In this scenario we will target pods in the current region that are associated with an application label. In this scenario you can customize an increasing amount of network latency injected on targeted EKS pods for the action duration, by default 5 min for each latency action.

AWS FIS supports the following scenarios for single-AZ, multi-AZ and multi-Region applications. These scenarios target multiple resource types.

- **AZ Availability: Power Interruption** - Inject the expected symptoms of a complete interruption of power in an Availability Zone (AZ). Learn more about [AZ Availability: Power Interruption](#).
- **AZ: Application Slowdown** - Add latency between resources within a single Availability Zone (AZ) to slow down an application. Learn more about [AZ: Application Slowdown](#).
- **Cross-AZ: Traffic Slowdown** - Inject packet loss to disrupt and slow down traffic between Availability Zones (AZs). Learn more about [Cross-AZ: Traffic Slowdown](#).
- **Cross-Region: Connectivity** - Block application network traffic from the experiment Region to the destination Region and pause cross-Region data replication. Learn more about using [Cross-Region: Connectivity](#).

AWS FIS supports the following scenarios for Amazon EBS volumes. These scenarios target volumes using tags. You can use your own tags or use the default tags included in the scenario. The target volumes must be in the same Availability Zone. For more information, [Fault testing on Amazon EBS](#).

- **EBS: Sustained Latency** — Explore impact of persistent I/O latency on your application.

In this scenario, we will target volumes in the current Availability Zone that have a specific tag attached. This scenario injects constant latency of 500 ms on 50 percent of read and 100 percent of write operations for a volume, using a single latency action over a 15-minute period. In this scenario, you can customize the amount of latency injected, the percentage of I/O injected, and the duration for the action.

- **EBS: Increasing Latency** — Explore impact of increasing I/O latency on your application.

In this scenario, we will target volumes in the current Availability Zone that have a specific tag attached. This scenario injects increasing latency of 50 ms, 200 ms, 700 ms, 1 second, and 15 seconds on 10 percent of read and 25 percent of write operations for a volume using five latency actions over a 15-minute period. In this scenario, you can customize the amount of latency injected, the percentage of I/O injected, and the action duration, for each latency action.

- **EBS: Intermittent Latency** — Explore impact of intermittent I/O latency spikes on your application.

In this scenario, we will target volumes in the current Availability Zone that have a specific tag attached. This scenario injects three sharp intermittent latency spikes of 30 seconds, 10 seconds, and 20 seconds on 0.1 percent of the read and write I/O operations for a volume, using three latency actions, with intervals of recovery in between each spike over a 15-minute period. In this scenario, you can customize the amount of latency injected, the percentage of I/O injected, and the action duration, for each latency action.

- **EBS: Decreasing Latency** — Explore impact of decreasing I/O latency on your application.

In this scenario, we will target volumes in the current Availability Zone that have a specific tag attached. This scenario injects decreasing latency of 20 seconds, 5 seconds, 900 ms, 300 ms, and 40 ms on 10 percent of read and write operations for a volume, using five latency actions over a 15-minute period. In this scenario, you can customize the amount of latency injected, the percentage of I/O injected, and the action duration, for each latency action.

## AZ Availability: Power Interruption

You can use the AZ Availability: Power Interruption scenario to induce the expected symptoms of a complete interruption of power in an Availability Zone (AZ).

This scenario can be used to demonstrate that multi-AZ applications operate as expected during a single, complete AZ power interruption. It includes loss of zonal compute (Amazon EC2, EKS,

and ECS), no re-scaling of compute in the AZ, subnet connectivity loss, RDS failover, ElastiCache failover, impaired access to S3 Express One Zone directory buckets, and unresponsive EBS volumes. By default, actions for which no targets are found will be skipped.

## Actions

Together, the following actions create many of the expected symptoms of a complete power interruption in a single AZ. AZ Availability: Power Interruption only affects services that are expected to see impact during a single AZ power interruption. By default, the scenario injects power interruption symptoms for 30 minutes and then, for an additional 30 minutes, injects symptoms that may occur during recovery.

### Stop-Instances

During an AZ power interruption, EC2 instances in the affected AZ will shut down. After power is restored instances will reboot. AZ Availability: Power Interruption includes [aws:ec2:stop-instances](#) to stop all instances in the affected AZ for the interruption duration. After the duration, the instances are restarted. Stopping EC2 instances managed by Amazon EKS causes dependent EKS pods to be deleted. Stopping EC2 instances managed by Amazon ECS causes dependent ECS tasks to be stopped.

This action targets EC2 instances running in the affected AZ. By default, it targets instances with a tag named `AzImpairmentPower` with a value of `StopInstances`. You can add this tag to your instances or replace the default tag with your own tag in the experiment template. By default, if no valid instances are found this action will be skipped.

### Stop-ASG-Instances

During an AZ power interruption, EC2 instances managed by an Auto Scaling group in the affected AZ will shut down. After power is restored instances will reboot. AZ Availability: Power Interruption includes [aws:ec2:stop-instances](#) to stop all instances, including those managed by Auto Scaling, in the affected AZ for the interruption duration. After the duration, the instances are restarted.

This action targets EC2 instances running in the affected AZ. By default, it targets instances with a tag named `AzImpairmentPower` with a value of `IceAsg`. You can add this tag to your instances or replace the default tag with your own tag in the experiment template. By default, if no valid instances are found this action will be skipped.

## Pause Instance Launches

During an AZ power interruption, EC2 API calls to provision capacity in the AZ will fail. In particular, the following APIs will be impacted: `ec2:StartInstances`, `ec2:CreateFleet`, and `ec2:RunInstances`. AZ Availability: Power Interruption includes [aws:ec2:api-insufficient-instance-capacity-error](#) to prevent new instances from being provisioned in the affected AZ.

This action targets IAM roles used to provision instances. These must be targeted using an ARN. By default, if no valid IAM roles are found this action will be skipped.

## Pause ASG Scaling

During an AZ power interruption, EC2 API calls made by the Auto Scaling control plane to recover lost capacity in the AZ will fail. In particular, the following APIs will be impacted: `ec2:StartInstances`, `ec2:CreateFleet`, and `ec2:RunInstances`. AZ Availability: Power Interruption includes [aws:ec2:asg-insufficient-instance-capacity-error](#) to prevent new instances from being provisioned in the affected AZ. This also prevents Amazon EKS and Amazon ECS from scaling in the affected AZ.

This action targets Auto Scaling groups. By default, it targets Auto Scaling groups with a tag named `AzImpairmentPower` with a value of `IceAsg`. You can add this tag to your Auto Scaling groups or replace the default tag with your own tag in the experiment template. By default, if no valid Auto Scaling groups are found this action will be skipped.

## Pause Network Connectivity

During an AZ power interruption, networking in the AZ will be unavailable. When this happens some AWS services may take up to a few minutes to update DNS to reflect that private endpoints in the affected AZ are not available. During this time, DNS lookups may return inaccessible IP addresses. AZ Availability: Power Interruption includes [aws:network:disrupt-connectivity](#) to block all network connectivity for all subnets in the affected AZ for 2 minutes. This will force timeouts and DNS refreshes for most applications. Ending the action after 2 minutes allows for subsequent recovery of regional service DNS while the AZ continues to be unavailable.

This action targets subnets. By default, it targets clusters with a tag named `AzImpairmentPower` with a value of `DisruptSubnet`. You can add this tag to your subnets or replace the default tag with your own tag in the experiment template. By default, if no valid subnets are found this action will be skipped.

## Failover RDS

During an AZ power interruption, RDS nodes in the affected AZ will shut down. Single AZ RDS nodes in the affected AZ will be fully unavailable. For multi-AZ clusters, the writer node will failover into an unaffected AZ and reader nodes in the affected AZ will be unavailable. For multi-AZ clusters, AZ Availability: Power Interruption includes [aws:rds:failover-db-cluster](#) to failover if the writer is in the affected AZ.

This action targets RDS clusters. By default, it targets clusters with a tag named `AzImpairmentPower` with a value of `DisruptRds`. You can add this tag to your clusters or replace the default tag with your own tag in the experiment template. By default, if no valid clusters are found this action will be skipped.

## Pause ElastiCache Replication Group

During an AZ power interruption, ElastiCache nodes in the AZ are unavailable. AZ Availability: Power Interruption includes [aws:elasticache:replicationgroup-interrupt-az-power](#) to terminate ElastiCache nodes in the affected AZ. For the duration of the interruption, new instances will not be provisioned in the affected AZ, so the replication group will remain at reduced capacity.

This action targets ElastiCache replication groups. By default, it targets replication groups with a tag named `AzImpairmentPower` with a value of `ElasticacheImpact`. You can add this tag to your replication groups or replace the default tag with your own tag in the experiment template. By default, if no valid replication groups are found this action will be skipped. Note that only replication groups with nodes in the affected AZ will be considered valid targets.

## Start ARC Zonal Autoshift

Five minutes after the AZ power interruption begins, the recovery action `aws:arc:start-zonal-autoshift` automatically shifts resource traffic away from the specified AZ for the remaining 25 minutes of the power interruption. After that duration, traffic shifts back to the original AZ. Note that during a real-world AZ power interruption AWS will detect the impairment and shift resource traffic if autoshift is enabled. While the timing of this shift varies it is estimated to occur five minutes from the impairment commencing.

This action targets Amazon Application Recovery Controller (ARC) autoshift-enabled resources. By default, it targets resources with the tag key `AzImpairmentPower` and value `RecoverAutoshiftResources`. You can add this tag to your resources or replace the default tag with your own tag in the experiment template. For example, you may want to use an application-specific tag. By default, if no valid resources are found this action will be skipped.

## Pause EBS I/O

After an AZ power interruption, once power is restored a very small percentage of instances may experience unresponsive EBS volumes. AZ Availability: Power Interruption includes [aws:ebs:pause-io](#) to leave 1 EBS volume in an unresponsive state.

By default, only volumes set to persist after the instance is terminated are targeted. This action targets volumes with a tag named `AzImpairmentPower` with a value of `APIPauseVolume`. You can add this tag to your volumes or replace the default tag with your own tag in the experiment template. By default, if no valid volumes are found this action will be skipped.

## Disrupt connectivity to S3 Express One Zone directory buckets

During an AZ power interruption, data stored in S3 Express One Zone directory buckets in the AZ is not accessible. AZ Availability: Power Interruption includes [aws:network:disrupt-connectivity](#) to disrupt connectivity between subnets and One Zone directory buckets in the affected AZ for the duration of the experiment, resulting in timeouts to Zonal endpoint data plane API operations. Use this action to test disruption when compute is co-located with storage in an AZ.

This action targets subnets. By default, it targets subnets with a tag named `AzImpairmentPower` with a value of `DisruptSubnet`. You can add this tag to your subnets or replace the default tag with your own tag in the experiment template. By default, if no valid subnets are found this action will be skipped.

## Limitations

- This scenario does not include [stop conditions](#). The correct stop conditions for your application should be added to the experiment template.
- In the targeted AZ, Amazon EKS Pods running on EC2 will be terminated with EC2 worker nodes and starting of new EC2 nodes will be blocked. However, Amazon EKS Pods running on AWS Fargate are not supported.
- In the targeted AZ, Amazon ECS tasks running on EC2 will be terminated with EC2 worker nodes and starting of new EC2 nodes will be blocked. However, Amazon ECS tasks running on AWS Fargate are not supported.
- [Amazon RDS Multi-AZ](#) with two readable standby DB instances is not supported. In this case, the instances will be terminated, RDS will failover, and capacity will immediately be provisioned back in the affected AZ. The readable standby in the affected AZ will remain available.

## Requirements

- Add the required permission to the AWS FIS [experiment role](#).
- Resource tags must be applied to resources that are to be targeted by the experiment. These can use your own tagging convention or the default tags defined in the scenario.

## Permissions

ARC zonal autoshift uses an IAM service-linked role

`AWSServiceRoleForZonalAutoshiftPracticeRun` to perform zonal shift on your behalf. This role uses the IAM managed policy [AWSZonalAutoshiftPracticeRunSLRPolicy](#). You don't need to create the role manually. When you create an experiment template from the AZ Power Interruption scenario in the AWS Management Console, the AWS CLI, or an AWS SDK, ARC creates the service-linked role for you. For more information, see [Using the service-linked role for zonal autoshift in ARC](#).

The following policy grants AWS FIS the necessary permissions to execute an experiment with the AZ Availability: Power Interruption scenario. This policy must be attached to the [experiment role](#).

JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "AllowFISExperimentLoggingActionsCloudwatch",
      "Effect": "Allow",
      "Action": [
        "logs:CreateLogDelivery",
        "logs:PutResourcePolicy",
        "logs:DescribeResourcePolicies",
        "logs:DescribeLogGroups"
      ],
      "Resource": "*"
    },
    {
      "Effect": "Allow",
      "Action": "ec2:CreateTags",
      "Resource": "arn:aws:ec2:*:*:network-acl/*",
      "Condition": {
```

```

        "StringEquals": {
            "ec2:CreateAction": "CreateNetworkAcl",
            "aws:RequestTag/managedByFIS": "true"
        }
    },
    {
        "Effect": "Allow",
        "Action": "ec2:CreateNetworkAcl",
        "Resource": "arn:aws:ec2:*:*:network-acl/*",
        "Condition": {
            "StringEquals": {
                "aws:RequestTag/managedByFIS": "true"
            }
        }
    },
    {
        "Effect": "Allow",
        "Action": [
            "ec2:CreateNetworkAclEntry",
            "ec2>DeleteNetworkAcl"
        ],
        "Resource": [
            "arn:aws:ec2:*:*:network-acl/*",
            "arn:aws:ec2:*:*:vpc/*"
        ],
        "Condition": {
            "StringEquals": {
                "ec2:ResourceTag/managedByFIS": "true"
            }
        }
    },
    {
        "Effect": "Allow",
        "Action": "ec2:CreateNetworkAcl",
        "Resource": "arn:aws:ec2:*:*:vpc/*"
    },
    {
        "Effect": "Allow",
        "Action": [
            "ec2:DescribeVpcs",
            "ec2:DescribeManagedPrefixLists",
            "ec2:DescribeSubnets",
            "ec2:DescribeNetworkAcls"
        ]
    }

```

```

    ],
    "Resource": "*"
  },
  {
    "Effect": "Allow",
    "Action": "ec2:ReplaceNetworkAclAssociation",
    "Resource": [
      "arn:aws:ec2:*:*:subnet/*",
      "arn:aws:ec2:*:*:network-acl/*"
    ]
  },
  {
    "Effect": "Allow",
    "Action": [
      "rds:FailoverDBCluster"
    ],
    "Resource": [
      "arn:aws:rds:*:*:cluster:*"
    ]
  },
  {
    "Effect": "Allow",
    "Action": [
      "rds:RebootDBInstance"
    ],
    "Resource": [
      "arn:aws:rds:*:*:db:*"
    ]
  },
  {
    "Effect": "Allow",
    "Action": [
      "elasticache:DescribeReplicationGroups",
      "elasticache:InterruptClusterAzPower"
    ],
    "Resource": [
      "arn:aws:elasticache:*:*:replicationgroup:*"
    ]
  },
  {
    "Sid": "TargetResolutionByTags",
    "Effect": "Allow",
    "Action": [
      "tag:GetResources"
    ]
  }

```

```

    ],
    "Resource": "*"
  },
  {
    "Effect": "Allow",
    "Action": [
      "ec2:StartInstances",
      "ec2:StopInstances"
    ],
    "Resource": "arn:aws:ec2:*:*:instance/*"
  },
  {
    "Effect": "Allow",
    "Action": [
      "ec2:DescribeInstances"
    ],
    "Resource": "*"
  },
  {
    "Effect": "Allow",
    "Action": [
      "kms:CreateGrant"
    ],
    "Resource": [
      "arn:aws:kms:*:*:key/*"
    ],
    "Condition": {
      "StringLike": {
        "kms:ViaService": "ec2.*.amazonaws.com"
      },
      "Bool": {
        "kms:GrantIsForAWSResource": "true"
      }
    }
  },
  {
    "Effect": "Allow",
    "Action": [
      "ec2:DescribeVolumes"
    ],
    "Resource": "*"
  },
  {
    "Effect": "Allow",

```

```

    "Action": [
      "ec2:PauseVolumeIO"
    ],
    "Resource": "arn:aws:ec2:*:*:volume/*"
  },
  {
    "Sid": "AllowInjectAPI",
    "Effect": "Allow",
    "Action": [
      "ec2:InjectApiError"
    ],
    "Resource": [
      "*"
    ],
    "Condition": {
      "ForAnyValue:StringEquals": {
        "ec2:FisActionId": [
          "aws:ec2:api-insufficient-instance-capacity-error",
          "aws:ec2:asg-insufficient-instance-capacity-error"
        ]
      }
    }
  },
  {
    "Sid": "DescribeAsg",
    "Effect": "Allow",
    "Action": [
      "autoscaling:DescribeAutoScalingGroups"
    ],
    "Resource": [
      "*"
    ]
  }
]
}

```

## Scenario Content

The following content defines the scenario. This JSON can be saved and used to create an [experiment template](#) using the [create-experiment-template](#) command from the AWS Command

Line Interface (AWS CLI). For the most recent version of the scenario, visit the scenario library in the FIS console.

```
{
  "targets": {
    "IAM-role": {
      "resourceType": "aws:iam:role",
      "resourceArns": [],
      "selectionMode": "ALL"
    },
    "EBS-Volumes": {
      "resourceType": "aws:ec2:ebs-volume",
      "resourceTags": {
        "AzImpairmentPower": "ApiPauseVolume"
      },
      "selectionMode": "COUNT(1)",
      "parameters": {
        "availabilityZoneIdentifier": "us-east-1a"
      },
      "filters": [
        {
          "path": "Attachments.DeleteOnTermination",
          "values": [
            "false"
          ]
        }
      ]
    },
    "EC2-Instances": {
      "resourceType": "aws:ec2:instance",
      "resourceTags": {
        "AzImpairmentPower": "StopInstances"
      },
      "filters": [
        {
          "path": "State.Name",
          "values": [
            "running"
          ]
        },
        {
          "path": "Placement.AvailabilityZone",
          "values": [
```

```

        "us-east-1a"
      ]
    }
  ],
  "selectionMode": "ALL"
},
"ASG": {
  "resourceType": "aws:ec2:autoscaling-group",
  "resourceTags": {
    "AzImpairmentPower": "IceAsg"
  },
  "selectionMode": "ALL"
},
"ASG-EC2-Instances": {
  "resourceType": "aws:ec2:instance",
  "resourceTags": {
    "AzImpairmentPower": "IceAsg"
  },
  "filters": [
    {
      "path": "State.Name",
      "values": [
        "running"
      ]
    },
    {
      "path": "Placement.AvailabilityZone",
      "values": [
        "us-east-1a"
      ]
    }
  ],
  "selectionMode": "ALL"
},
"Subnet": {
  "resourceType": "aws:ec2:subnet",
  "resourceTags": {
    "AzImpairmentPower": "DisruptSubnet"
  },
  "filters": [
    {
      "path": "AvailabilityZone",
      "values": [
        "us-east-1a"
      ]
    }
  ]
}

```

```

        ]
      }
    ],
    "selectionMode": "ALL",
    "parameters": {}
  },
  "RDS-Cluster": {
    "resourceType": "aws:rds:cluster",
    "resourceTags": {
      "AzImpairmentPower": "DisruptRds"
    },
    "selectionMode": "ALL",
    "parameters": {
      "writerAvailabilityZoneIdentifiers": "us-east-1a"
    }
  },
  "ElastiCache-Cluster": {
    "resourceType": "aws:elasticache:replicationgroup",
    "resourceTags": {
      "AzImpairmentPower": "DisruptElasticache"
    },
    "selectionMode": "ALL",
    "parameters": {
      "availabilityZoneIdentifier": "us-east-1a"
    }
  }
},
"actions": {
  "Pause-Instance-Launches": {
    "actionId": "aws:ec2:api-insufficient-instance-capacity-error",
    "parameters": {
      "availabilityZoneIdentifiers": "us-east-1a",
      "duration": "PT30M",
      "percentage": "100"
    },
    "targets": {
      "Roles": "IAM-role"
    }
  },
  "Pause-EBS-IO": {
    "actionId": "aws:ebs:pause-volume-io",
    "parameters": {
      "duration": "PT30M"
    }
  },

```

```
    "targets": {
      "Volumes": "EBS-Volumes"
    },
    "startAfter": [
      "Stop-Instances",
      "Stop-ASG-Instances"
    ]
  },
  "Stop-Instances": {
    "actionId": "aws:ec2:stop-instances",
    "parameters": {
      "completeIfInstancesTerminated": "true",
      "startInstancesAfterDuration": "PT30M"
    },
    "targets": {
      "Instances": "EC2-Instances"
    }
  },
  "Pause-ASG-Scaling": {
    "actionId": "aws:ec2:asg-insufficient-instance-capacity-error",
    "parameters": {
      "availabilityZoneIdentifiers": "us-east-1a",
      "duration": "PT30M",
      "percentage": "100"
    },
    "targets": {
      "AutoScalingGroups": "ASG"
    }
  },
  "Stop-ASG-Instances": {
    "actionId": "aws:ec2:stop-instances",
    "parameters": {
      "completeIfInstancesTerminated": "true",
      "startInstancesAfterDuration": "PT30M"
    },
    "targets": {
      "Instances": "ASG-EC2-Instances"
    }
  },
  "Pause-network-connectivity": {
    "actionId": "aws:network:disrupt-connectivity",
    "parameters": {
      "duration": "PT2M",
      "scope": "all"
    }
  }
}
```

```
    },
    "targets": {
      "Subnets": "Subnet"
    }
  },
  "Failover-RDS": {
    "actionId": "aws:rds:failover-db-cluster",
    "parameters": {},
    "targets": {
      "Clusters": "RDS-Cluster"
    }
  },
  "Pause-ElastiCache": {
    "actionId": "aws:elasticache:replicationgroup-interrupt-az-power",
    "parameters": {
      "duration": "PT30M"
    },
    "targets": {
      "ReplicationGroups": "ElastiCache-Cluster"
    }
  }
},
"stopConditions": [
  {
    "source": "aws:cloudwatch:alarm",
    "value": ""
  }
],
"roleArn": "",
"tags": {
  "Name": "AZ Impairment: Power Interruption"
},
"logConfiguration": {
  "logSchemaVersion": 2
},
"experimentOptions": {
  "accountTargeting": "single-account",
  "emptyTargetResolutionMode": "skip"
},
"description": "Affect multiple resource types in a single AZ, targeting by tags and explicit ARNs, to approximate power interruption in one AZ."
}
```

## AZ: Application Slowdown

You can use the AZ: Application Slowdown scenario to introduce additional latency between resources within a single Availability Zone (AZ). This latency creates many of the symptoms of an application slowdown, a partial disruption, sometimes known as a gray failure. It adds latency to network flows between target resources. Network flows represent the traffic between computing resources — the data packets carrying requests, responses, and other communications between your servers, containers, and services. The scenario can help to validate observability setups, tune alarm thresholds, discover application sensitivity to slowdowns, and practice critical operational decisions like AZ evacuation.

By default, the scenario adds 200ms of latency to 100% of network flows between target resources within the selected AZ for a duration of 30 minutes. You can use the **Edit shared parameters** dialog in the AWS FIS console to adjust the following parameters at the scenario level, which then apply to the underlying actions:

- Availability Zone - you can select the AZ to impair in the scenario.
- Milliseconds (ms) of latency - adjust this based on your application's sensitivity and needs. You can set latency lower for more sensitive applications or higher to test timeout handling, for example. Consider using multiples of your current application latency as a baseline.
- Flows percentage - reduce to impair a subset of traffic. For example, you can add 200ms latency affecting 25% of the network flows for even more subtle testing.
- Duration - set how long the experiment runs. You can shorten for quicker tests, or run longer sustained tests. For example, set the duration to 2 hours to test your recovery mechanisms under impaired conditions.
- Resource targeting - you can define target resources for the overall scenario using tags (for EC2 instances or ECS tasks on EC2 or Fargate) or labels (for EKS pods on EC2). You can specify your own tags and labels, or use the defaults provided in the scenario. If you don't wish to use tags or labels, you can edit the action to target resources by specifying other parameters.
- Customization - If you don't want to target EC2 or ECS resources, you can leave the actions with default tags. The experiment won't find any resources to target and the action will be skipped. However, if you don't want to target EKS resources, you should remove the EKS action and target from the scenario completely, as it requires an EKS cluster identifier to be provided. For even more granular customization, you can modify individual actions in the experiment template directly.

## Actions

Together, the following actions help create many of the symptoms of an application slowdown in a single AZ by introducing additional latency on the network flows, which then propagates through the application. These actions run in parallel, each adding 200ms latency for 30 minutes by default. After this period, latency returns to normal levels. The scenario needs at least one of the following resource types to run: EC2 instance, ECS task, or EKS pod.

### ECS Network Latency

AZ: Application Slowdown includes [aws:ecs:task-network-latency](#) to introduce latency for ECS tasks. The action targets tasks in the selected AZ. By default, it targets tasks with a [tag](#) named `AZApplicationSlowdown` with a value of `LatencyForECS`. You can replace the default tag with your own, or add the scenario tag to your tasks. If no valid tasks are found this action will be skipped. Before running an experiment on ECS, you should follow the [setup steps for ECS task actions](#).

### EKS Network Latency

AZ: Application Slowdown includes [aws:eks:pod-network-latency](#) to introduce latency for EKS pods. The action targets pods in the selected AZ. By default, it targets pods within a cluster that have labels with the format `key=value`. The default label provided is `AZApplicationSlowdown=LatencyForEKS`. You can replace the default label with your own, or add this label to your pods. If no valid pods are found this action will be skipped. Before running an experiment on EKS, you should follow the [setup steps for EKS pod actions](#).

### EC2 Network Latency

AZ: Application Slowdown uses the [aws:ssm:send-command](#) action to run the [AWSFIS-Run-Network-Latency-Sources](#) document to introduce latency for EC2 instances. The action targets instances in the selected AZ. By default, it targets instances with a [tag](#) named `AZApplicationSlowdown` with a value of `LatencyForEC2`. You can replace the default tag with your own, or add this tag to your instances. If no valid instances are found this action will be skipped. Before running an experiment on EC2 using SSM, you should [configure the AWS Systems Manager agent](#).

## Limitations

- This scenario does not include [stop conditions](#). The correct stop conditions for your application should be added to the experiment template.

## Requirements

- Add the required permissions to the AWS FIS [experiment role](#).
- You need to target one or more resources from any of the following 3 types within the selected AZ: EC2 instances, ECS tasks, or EKS pods.
- All targets of the scenario must be in the same VPC.

## Permissions

To run this scenario you need an IAM role with a trust policy that allows FIS to assume the role and the managed policies for the resource types you target in the experiment: EC2, ECS, and EKS. When you create an experiment template from the AZ: Application Slowdown scenario, FIS creates the role for you with the trust policy and the following AWS managed policies:

- [AWSFaultInjectionSimulatorEC2Access](#)
- [AWSFaultInjectionSimulatorECSAccess](#)
- [AWSFaultInjectionSimulatorEKSAccess](#)

If you're using an existing [IAM role](#) to run the AZ: Application Slowdown scenario, you can attach the following policy to grant AWS FIS the necessary permissions:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "DescribeTasks",
      "Effect": "Allow",
      "Action": "ecs:DescribeTasks",
      "Resource": "*"
    },
    {
      "Sid": "DescribeContainerInstances",
      "Effect": "Allow",
      "Action": "ecs:DescribeContainerInstances",
      "Resource": "arn:aws:ecs:*:*:container-instance/*/*"
    },
    {
      "Sid": "DescribeInstances",
      "Effect": "Allow",
```

```

    "Action": "ec2:DescribeInstances",
    "Resource": "*"
  },
  {
    "Sid": "DescribeSubnets",
    "Effect": "Allow",
    "Action": "ec2:DescribeSubnets",
    "Resource": "*"
  },
  {
    "Sid": "DescribeCluster",
    "Effect": "Allow",
    "Action": "eks:DescribeCluster",
    "Resource": "arn:aws:eks:*:*:cluster/*"
  },
  {
    "Sid": "TargetResolutionByTags",
    "Effect": "Allow",
    "Action": "tag:GetResources",
    "Resource": "*"
  },
  {
    "Sid": "SendCommand",
    "Effect": "Allow",
    "Action": [
      "ssm:SendCommand"
    ],
    "Resource": [
      "arn:aws:ec2:*:*:instance/*",
      "arn:aws:ssm:*:*:managed-instance/*",
      "arn:aws:ssm:*:*:document/*"
    ]
  },
  {
    "Sid": "ListCommands",
    "Effect": "Allow",
    "Action": [
      "ssm:ListCommands"
    ],
    "Resource": "*"
  },
  {
    "Sid": "CancelCommand",
    "Effect": "Allow",

```

```

        "Action": [
            "ssm:CancelCommand"
        ],
        "Resource": "*"
    }
]
}

```

## Scenario Content

The following content defines the scenario. This JSON can be saved and used to create an [experiment template](#) using the [create-experiment-template](#) command from the AWS Command Line Interface (AWS CLI). For the most recent version of the scenario, visit the scenario library in the FIS console and go to the **Content** tab.

```

{
  "tags": {
    "Name": "AZ: Application Slowdown"
  },
  "description": "Add latency between resources within a single AZ.",
  "actions": {
    "LatencyForEKS": {
      "actionId": "aws:eks:pod-network-latency",
      "parameters": {
        "delayMilliseconds": "200",
        "duration": "PT30M",
        "flowsPercent": "100",
        "interface": "DEFAULT",
        "kubernetesServiceAccount": "fis-service-account",
        "sources": "us-east-1a"
      },
      "targets": {
        "Pods": "TargetsForEKS"
      }
    },
    "LatencyForEC2": {
      "actionId": "aws:ssm:send-command",
      "parameters": {
        "duration": "PT30M",
        "documentArn": "arn:aws:ssm:us-east-1::document/AWSFIS-Run-Network-Latency-Sources",

```

```

        "documentParameters": "{\"DelayMilliseconds\": \"200\", \"Sources\": \"us-
east-1a\", \"Interface\": \"DEFAULT\", \"TrafficType\": \"egress\", \"DurationSeconds\":
1800\", \"FlowsPercent\": \"100\", \"InstallDependencies\": \"True\"}"
    },
    "targets": {
        "Instances": "TargetsForEC2"
    }
},
"LatencyForECS": {
    "actionId": "aws:ecs:task-network-latency",
    "parameters": {
        "delayMilliseconds": "200",
        "duration": "PT30M",
        "flowsPercent": "100",
        "installDependencies": "true",
        "sources": "us-east-1a",
        "useEcsFaultInjectionEndpoints": "true"
    },
    "targets": {
        "Tasks": "TargetsForECS"
    },
    "startAfter": []
}
},
"targets": {
    "TargetsForEKS": {
        "parameters": {
            "availabilityZoneIdentifier": "us-east-1a",
            "clusterIdentifier": "",
            "namespace": "default",
            "selectorType": "labelSelector",
            "selectorValue": "AZApplicationSlowdown=LatencyForEKS"
        },
        "resourceType": "aws:eks:pod",
        "selectionMode": "ALL"
    },
    "TargetsForEC2": {
        "filters": [
            {
                "path": "Placement.AvailabilityZone",
                "values": [
                    "us-east-1a"
                ]
            }
        ]
    }
}

```

```
    ],
    "resourceTags": {
      "AZApplicationSlowdown": "LatencyForEC2"
    },
    "resourceType": "aws:ec2:instance",
    "selectionMode": "ALL"
  },
  "TargetsForECS": {
    "filters": [
      {
        "path": "AvailabilityZone",
        "values": [
          "us-east-1a"
        ]
      }
    ],
    "resourceTags": {
      "AZApplicationSlowdown": "LatencyForECS"
    },
    "resourceType": "aws:ecs:task",
    "selectionMode": "ALL"
  }
},
"experimentOptions": {
  "accountTargeting": "single-account",
  "emptyTargetResolutionMode": "skip"
},
"stopConditions": [
  {
    "source": "none"
  }
]
}
```

## Cross-AZ: Traffic Slowdown

You can use the Cross-AZ: Traffic Slowdown scenario to inject packet loss to disrupt and slow down traffic between Availability Zones (AZs). The packet loss impairs cross-AZ communication, a partial disruption, sometimes known as a gray failure. It injects packet loss on network flows between target resources. Network flows represent the traffic between computing resources — the data packets carrying requests, responses, and other communications between your servers, containers, and services. The scenario can help to validate observability setups, tune alarm thresholds,

discover application sensitivity and dependencies in cross-AZ communication, and practice critical operational decisions like AZ evacuation.

By default, the scenario injects 15% packet loss to 100% of outbound network flows for target resources from the selected AZ for a duration of 30 minutes. You can use the **Edit shared parameters** dialog in the AWS FIS console to adjust the following parameters at the scenario level, which then apply to the underlying actions:

- **Availability Zone** - you can select the AZ to impair, and packet loss will be injected from that AZ to the other AZs within the Region.
- **Packet Loss** - adjust the packet loss lower for subtle disruption testing, such as 5%, or higher to test severe communication degradation and recovery mechanisms, such as 50%, or even 100% for total connectivity impact.
- **Flows percentage** - reduce to impair a subset of traffic. For example, you can inject 15% packet loss affecting 25% of the network flows for even more subtle testing.
- **Duration** - set how long the experiment runs. You can shorten for quicker tests, or run longer sustained tests. For example, set the duration to 2 hours to help test recovery mechanisms under impaired conditions.
- **Resource targeting** - you can define target resources for the overall scenario using tags (for EC2 instances or ECS tasks on EC2 or Fargate) or labels (for EKS pods on EC2). You can specify your own tags and labels, or use the defaults provided in the scenario. If you don't wish to use tags or labels, you can edit the action to target resources by specifying other parameters.
- **Customization** - If you don't want to target EC2 or ECS resources, you can leave the actions with default tags. The experiment won't find any resources to target and the action will be skipped. However, if you don't want to target EKS resources, you should remove the EKS action and target from the scenario completely, as it requires an EKS cluster identifier to be provided. For even more granular customization, you can modify individual actions in the experiment template directly.

## Actions

Together, the following actions help create the symptoms of a traffic slowdown between AZs by introducing packet loss on outbound communication from the target AZ to other AZs in the region at the network layer. These actions run in parallel, each injecting 15% packet loss for 30 minutes by default. After this period, communication returns to normal. The scenario needs at least one of the following resource types in the selected AZ to run: EC2 instance, ECS task, or EKS pod.

## ECS Network Packet Loss

Cross-AZ: Traffic Slowdown includes [aws:ecs:task-network-packet-loss](#) to inject packet loss for ECS tasks. The action targets tasks in the selected AZ, and impairs their outbound communication to all other AZs in the Region. You can further customize the scope of the impact by editing the action and adding or removing AZs from the Sources field. By default, it targets tasks with a [tag](#) named `CrossAZTrafficSlowdown` with a value of `PacketLossForECS`. You can replace the default tag with your own, or add the scenario tag to your tasks. If no valid tasks are found this action will be skipped. Before running an experiment on ECS, you should follow the [setup steps for ECS task actions](#).

## EKS Network Packet Loss

Cross-AZ: Traffic Slowdown includes [aws:eks:pod-network-packet-loss](#) to inject packet loss for EKS pods. The action targets pods in the selected AZ, and impairs their outbound communication to all other AZs in the region. You can further customize the scope of the impact by editing the action and adding or removing AZs from the Sources field. By default, it targets pods within a cluster that have labels with the format `key=value`. The default label provided is `CrossAZTraffic=PacketLossForEKS`. You can replace the default label with your own, or add this label to your pods. If no valid pods are found this action will be skipped. Before running an experiment on EKS, you should follow the [setup steps for EKS pod actions](#).

## EC2 Network Packet Loss

Cross-AZ: Traffic Slowdown uses the [aws:ssm:send-command](#) action to run the [AWSFIS-Run-Network-Packet-Loss-Sources](#) document to inject packet loss for EC2 instances, and impairs their outbound communication to all other AZs in the Region. You can further customize the scope of the impact by editing the action and adding or removing AZs from the Sources field. The action targets instances in the selected AZ. By default, it targets instances with a [tag](#) named `CrossAZTrafficSlowdown` with a value of `PacketLossForEC2`. You can replace the default tag with your own, or add this tag to your instances. If no valid instances are found this action will be skipped. Before running an experiment on EC2 using SSM, you should [configure the AWS Systems Manager agent](#).

## Limitations

- This scenario does not include [stop conditions](#). The correct stop conditions for your application should be added to the experiment template.

## Requirements

- Add the required permissions to the AWS FIS [experiment role](#).
- You need to target one or more resources from any of the following 3 types within the selected AZ: EC2 instances, ECS tasks, or EKS pods.
- All targets of the scenario must be in the same VPC.

## Permissions

To run this scenario you need an IAM role with a trust policy that allows FIS to assume the role and the managed policies for the resource types you target in the experiment: EC2, ECS, and EKS. When you create an experiment template from the Cross-AZ: Traffic Slowdown scenario, FIS creates the role for you with the trust policy and the following AWS managed policies:

- [AWSFaultInjectionSimulatorEC2Access](#)
- [AWSFaultInjectionSimulatorECSAccess](#)
- [AWSFaultInjectionSimulatorEKSAccess](#)

If you're using an existing [IAM role](#) to run the Cross-AZ: Traffic Slowdown scenario, you can attach the following policy to grant AWS FIS the necessary permissions:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "DescribeTasks",
      "Effect": "Allow",
      "Action": "ecs:DescribeTasks",
      "Resource": "*"
    },
    {
      "Sid": "DescribeContainerInstances",
      "Effect": "Allow",
      "Action": "ecs:DescribeContainerInstances",
      "Resource": "arn:aws:ecs:*:*:container-instance/*/*"
    },
    {
      "Sid": "DescribeInstances",
      "Effect": "Allow",
```

```

    "Action": "ec2:DescribeInstances",
    "Resource": "*"
  },
  {
    "Sid": "DescribeSubnets",
    "Effect": "Allow",
    "Action": "ec2:DescribeSubnets",
    "Resource": "*"
  },
  {
    "Sid": "DescribeCluster",
    "Effect": "Allow",
    "Action": "eks:DescribeCluster",
    "Resource": "arn:aws:eks:*:*:cluster/*"
  },
  {
    "Sid": "TargetResolutionByTags",
    "Effect": "Allow",
    "Action": "tag:GetResources",
    "Resource": "*"
  },
  {
    "Sid": "SendCommand",
    "Effect": "Allow",
    "Action": [
      "ssm:SendCommand"
    ],
    "Resource": [
      "arn:aws:ec2:*:*:instance/*",
      "arn:aws:ssm:*:*:managed-instance/*",
      "arn:aws:ssm:*:*:document/*"
    ]
  },
  {
    "Sid": "ListCommands",
    "Effect": "Allow",
    "Action": [
      "ssm:ListCommands"
    ],
    "Resource": "*"
  },
  {
    "Sid": "CancelCommand",
    "Effect": "Allow",

```

```

        "Action": [
            "ssm:CancelCommand"
        ],
        "Resource": "*"
    }
}

```

## Scenario Content

The following content defines the scenario. This JSON can be saved and used to create an [experiment template](#) using the [create-experiment-template](#) command from the AWS Command Line Interface (AWS CLI). For the most recent version of the scenario, visit the scenario library in the FIS console and go to the **Content** tab.

```

{
  "tags": {
    "Name": "Cross-AZ: Traffic Slowdown"
  },
  "description": "Inject packet loss to disrupt and slow down traffic between AZs.",
  "actions": {
    "PacketLossForEC2": {
      "actionId": "aws:ssm:send-command",
      "parameters": {
        "duration": "PT30M",
        "documentArn": "arn:aws:ssm:us-east-1::document/AWSFIS-Run-Network-
Packet-Loss-Sources",
        "documentParameters": "{\"Sources\": \"us-east-1b,us-east-1c,us-
east-1d,us-east-1e,us-east-1f\", \"LossPercent\": \"15\", \"Interface\": \"DEFAULT\",
\"TrafficType\": \"egress\", \"DurationSeconds\": \"1800\", \"FlowsPercent\": \"100\",
\"InstallDependencies\": \"True\"}"
      },
      "targets": {
        "Instances": "TargetsForEC2"
      }
    },
    "PacketLossForECS": {
      "actionId": "aws:ecs:task-network-packet-loss",
      "parameters": {
        "sources": "us-east-1b,us-east-1c,us-east-1d,us-east-1e,us-east-1f",
        "lossPercent": "15",
        "duration": "PT30M",
        "flowsPercent": "100",

```



```

        ]
      }
    ],
    "resourceTags": {
      "CrossAZTrafficSlowdown": "PacketLossForECS"
    },
    "resourceType": "aws:ecs:task",
    "selectionMode": "ALL"
  },
  "TargetsForEKS": {
    "parameters": {
      "availabilityZoneIdentifier": "us-east-1a",
      "clusterIdentifier": "",
      "namespace": "default",
      "selectorType": "labelSelector",
      "selectorValue": "CrossAZTrafficSlowdown=PacketLossForEKS"
    },
    "resourceType": "aws:eks:pod",
    "selectionMode": "ALL"
  }
},
"experimentOptions": {
  "accountTargeting": "single-account",
  "emptyTargetResolutionMode": "skip"
},
"stopConditions": [
  {
    "source": "none"
  }
]
}

```

## Cross-Region: Connectivity

You can use the Cross-Region: Connectivity scenario to block application network traffic from the experiment Region to the destination Region and pause cross-Region replication for Amazon S3 and Amazon DynamoDB multi-Region global tables. Cross Region: Connectivity affects outbound application traffic from the Region in which you run the experiment (*experiment Region*). Stateless inbound traffic from the Region you wish to isolate from the *experiment region* (*destination Region*) may not be blocked. Traffic from AWS managed services may not be blocked.

This scenario can be used to demonstrate that multi-Region applications operate as expected when resources in the destination Region are not accessible from the experiment Region. It includes blocking network traffic from the experiment Region to the destination Region by targeting transit gateways and route tables. It also pauses cross-Region replication for S3 and DynamoDB global tables. By default, actions for which no targets are found will be skipped.

## Actions

Together, the following actions block cross-Region connectivity for the included AWS services. The actions are run in parallel. By default, the scenario blocks traffic for 3 hours, which you can increase up to a maximum 12 Hour duration.

### Disrupt Transit Gateway Connectivity

Cross Region: Connectivity includes [aws:network:transit-gateway-disrupt-cross-region-connectivity](#) to block cross-Region network traffic from VPCs in the *experiment Region* to VPCs in the *destination Region* connected by a transit gateway. This does not affect access to VPC endpoints within the *experiment Region* but will block traffic from the *experiment Region* destined for a VPC endpoint in the *destination Region*.

This action targets transit gateways connecting the *experiment Region* and the *destination Region*. By default, it targets transit gateways with a [tag](#) named `DisruptTransitGateway` with a value of `Allowed`. You can add this tag to your transit gateways or replace the default tag with your own tag in the experiment template. By default, if no valid transit gateways are found this action will be skipped.

### Disrupt Subnet Connectivity

Cross Region: Connectivity includes [aws:network:route-table-disrupt-cross-region-connectivity](#) to block cross-Region network traffic from VPCs in the *experiment Region* to public AWS IP blocks in the *destination Region*. These public IP blocks include AWS service endpoints in the *destination Region*, e.g. the S3 regional endpoint, and AWS IP blocks for managed services, e.g. the IP addresses used for load balancers and Amazon API Gateway. This action also blocks network connectivity over cross-Region VPC Peering connections from the *experiment Region* to the *destination Region*. It does not affect access to VPC endpoints in the *experiment Region* but will block traffic from the *experiment Region* destined for a VPC endpoint in the *destination Region*.

This action targets subnets in the experiment Region. By default, it targets subnets with a [tag](#) named `DisruptSubnet` with a value of `Allowed`. You can add this tag to your subnets or replace

the default tag with your own tag in the experiment template. By default, if no valid subnets are found this action will be skipped.

### Disrupt VPC Endpoint Connectivity

Cross Region: Connectivity includes [aws:network:disrupt-vpc-endpoint](#) disrupt connectivity to a service associated with the target VPC endpoints. For example, if a VPC endpoint creates a Private Link to com.amazonaws.us-east-1.ec2, then the connectivity to that service will be disrupted.

This action targets VPC endpoints in the experiment region. By default, it targets interface VPC endpoints with a [tag](#) named `DisruptVpcEndpoint` with a value `Allowed`. You can add this tag to your VPC endpoints or replace the default tag with your own tag in the experiment template. By default, if no valid VPC endpoints are found this action will be skipped.

### Pause S3 Replication

Cross Region: Connectivity includes [aws:s3:bucket-pause-replication](#) to pause S3 replication from the *experiment Region* to the *destination Region* for the targeted buckets. Replication from the *destination Region* to the *experiment Region* will be unaffected. After the scenario ends, bucket replication will resume from the point it was paused. Note that the time it takes for replication to keep all objects in sync will vary based on the duration of the experiment, and the rate of object upload to the bucket.

This action targets S3 buckets in the experiment Region with [Cross-Region Replication](#) (CRR) enabled to an S3 bucket in the destination Region. By default, it targets buckets with a [tag](#) named `DisruptS3` with a value of `Allowed`. You can add this tag to your buckets or replace the default tag with your own tag in the experiment template. By default, if no valid buckets are found this action will be skipped.

### Pause DynamoDB Replication

Cross-Region: Connectivity includes [aws:dynamodb:global-table-pause-replication](#) to pause replication between the experiment Region and all other Regions, including the destination Region. This prevents replication into and out of the *experiment Region* but does not affect replication between other Regions. After the scenario ends, table replication will resume from the point it was paused. Note that the time it takes for replication to keep all data in sync will vary based on the duration of the experiment and the rate of changes to the table.

This action targets both DynamoDB multi-Region strongly and eventually consistent global tables in the experiment Region. By default, it targets tables with a [tag](#) named `DisruptDynamoDb` with a value of `Allowed`. You can add this tag to your tables or replace the default tag with your own

tag in the experiment template. By default, if no valid global tables are found this action will be skipped.

## Pause MemoryDB multi-Region Replication

Cross-Region: Connectivity includes [aws:memorydb:multi-region-cluster-pause-replication](#) to pause replication from the regional member cluster in the experiment Region to the rest of the clusters in the targeted multi-Region cluster. Replication between other regional member clusters will be unaffected. After the scenario ends, replication will resume from the point it was paused. Note that the time for replication to sync data between the member clusters will vary based on the duration of the experiment, and the rate of data written to the clusters.

This action targets MemoryDB Multi-Region clusters with a regional member in the experiment Region. By default, it targets multi-Region clusters with a [tag](#) named `DisruptMemoryDB` with a value of `Allowed`. You can add this tag to your multi-Region clusters or replace the default tag with your own tag in the experiment template. By default, if no valid clusters are found this action will be skipped.

## Limitations

- This scenario does not include [stop conditions](#). The correct stop conditions for your application should be added to the experiment template.

## Requirements

- Add the required permission to the AWS FIS [experiment role](#).
- Resource tags must be applied to resources that are to be targeted by the experiment. These can use your own tagging convention or the default tags defined in the scenario.

## Permissions

The following policy grants AWS FIS the necessary permissions to execute an experiment with the Cross-Region: Connectivity scenario. This policy must be attached to the [experiment role](#).

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "RouteTableDisruptConnectivity1",
```

```

    "Effect": "Allow",
    "Action": "ec2:CreateRouteTable",
    "Resource": "arn:aws:ec2:*:*:route-table/*",
    "Condition": {
      "StringEquals": {
        "aws:RequestTag/managedByFIS": "true"
      }
    }
  },
  {
    "Sid": "RouteTableDisruptConnectivity2",
    "Effect": "Allow",
    "Action": "ec2:CreateRouteTable",
    "Resource": "arn:aws:ec2:*:*:vpc/*"
  },
  {
    "Sid": "RouteTableDisruptConnectivity21",
    "Effect": "Allow",
    "Action": "ec2:CreateTags",
    "Resource": "arn:aws:ec2:*:*:route-table/*",
    "Condition": {
      "StringEquals": {
        "ec2:CreateAction": "CreateRouteTable",
        "aws:RequestTag/managedByFIS": "true"
      }
    }
  },
  {
    "Sid": "RouteTableDisruptConnectivity3",
    "Effect": "Allow",
    "Action": "ec2:CreateTags",
    "Resource": "arn:aws:ec2:*:*:network-interface/*",
    "Condition": {
      "StringEquals": {
        "ec2:CreateAction": "CreateNetworkInterface",
        "aws:RequestTag/managedByFIS": "true"
      }
    }
  },
  {
    "Sid": "RouteTableDisruptConnectivity4",
    "Effect": "Allow",
    "Action": "ec2:CreateTags",
    "Resource": "arn:aws:ec2:*:*:prefix-list*",

```

```

    "Condition": {
      "StringEquals": {
        "ec2:CreateAction": "CreateManagedPrefixList",
        "aws:RequestTag/managedByFIS": "true"
      }
    }
  },
  {
    "Sid": "RouteTableDisruptConnectivity5",
    "Effect": "Allow",
    "Action": "ec2:DeleteRouteTable",
    "Resource": [
      "arn:aws:ec2:*:*:route-table/*",
      "arn:aws:ec2:*:*:vpc/*"
    ],
    "Condition": {
      "StringEquals": {
        "ec2:ResourceTag/managedByFIS": "true"
      }
    }
  },
  {
    "Sid": "RouteTableDisruptConnectivity6",
    "Effect": "Allow",
    "Action": "ec2:CreateRoute",
    "Resource": "arn:aws:ec2:*:*:route-table/*",
    "Condition": {
      "StringEquals": {
        "ec2:ResourceTag/managedByFIS": "true"
      }
    }
  },
  {
    "Sid": "RouteTableDisruptConnectivity7",
    "Effect": "Allow",
    "Action": "ec2:CreateNetworkInterface",
    "Resource": "arn:aws:ec2:*:*:network-interface/*",
    "Condition": {
      "StringEquals": {
        "aws:RequestTag/managedByFIS": "true"
      }
    }
  },
  {

```

```

    "Sid": "RouteTableDisruptConnectivity8",
    "Effect": "Allow",
    "Action": "ec2:CreateNetworkInterface",
    "Resource": [
      "arn:aws:ec2:*:*:subnet/*",
      "arn:aws:ec2:*:*:security-group/*"
    ]
  },
  {
    "Sid": "RouteTableDisruptConnectivity9",
    "Effect": "Allow",
    "Action": "ec2:DeleteNetworkInterface",
    "Resource": "arn:aws:ec2:*:*:network-interface/*",
    "Condition": {
      "StringEquals": {
        "ec2:ResourceTag/managedByFIS": "true"
      }
    }
  },
  {
    "Sid": "RouteTableDisruptConnectivity10",
    "Effect": "Allow",
    "Action": "ec2:CreateManagedPrefixList",
    "Resource": "arn:aws:ec2:*:*:prefix-list/*",
    "Condition": {
      "StringEquals": {
        "aws:RequestTag/managedByFIS": "true"
      }
    }
  },
  {
    "Sid": "RouteTableDisruptConnectivity11",
    "Effect": "Allow",
    "Action": [
      "ec2:DeleteManagedPrefixList",
      "ec2:ModifyManagedPrefixList"
    ],
    "Resource": "arn:aws:ec2:*:*:prefix-list/*",
    "Condition": {
      "StringEquals": {
        "aws:ResourceTag/managedByFIS": "true"
      }
    }
  }
},

```

```

{
  "Sid": "EC2DescribeResources",
  "Effect": "Allow",
  "Action": [
    "ec2:DescribeNetworkInterfaces",
    "ec2:DescribeVpcs",
    "ec2:DescribeVpcPeeringConnections",
    "ec2:DescribeManagedPrefixLists",
    "ec2:DescribeSubnets",
    "ec2:DescribeRouteTables",
    "ec2:DescribeVpcEndpoints",
    "ec2:DescribeTransitGatewayPeeringAttachments",
    "ec2:DescribeTransitGatewayAttachments",
    "ec2:DescribeTransitGateways",
    "ec2:DescribeSecurityGroups"
  ],
  "Resource": "*"
},
{
  "Sid": "RouteTableDisruptConnectivity14",
  "Effect": "Allow",
  "Action": "ec2:ReplaceRouteTableAssociation",
  "Resource": [
    "arn:aws:ec2:*:*:subnet/*",
    "arn:aws:ec2:*:*:route-table/*"
  ]
},
{
  "Sid": "RouteTableDisruptConnectivity15",
  "Effect": "Allow",
  "Action": "ec2:GetManagedPrefixListEntries",
  "Resource": "arn:aws:ec2:*:*:prefix-list/*"
},
{
  "Sid": "RouteTableDisruptConnectivity16",
  "Effect": "Allow",
  "Action": "ec2:AssociateRouteTable",
  "Resource": [
    "arn:aws:ec2:*:*:subnet/*",
    "arn:aws:ec2:*:*:route-table/*"
  ]
},
{
  "Sid": "RouteTableDisruptConnectivity17",

```

```

    "Effect": "Allow",
    "Action": "ec2:DisassociateRouteTable",
    "Resource": "arn:aws:ec2:*:*:route-table/*",
    "Condition": {
      "StringEquals": {
        "ec2:ResourceTag/managedByFIS": "true"
      }
    }
  },
  {
    "Sid": "RouteTableDisruptConnectivity18",
    "Effect": "Allow",
    "Action": "ec2:DisassociateRouteTable",
    "Resource": "arn:aws:ec2:*:*:subnet/*"
  },
  {
    "Sid": "RouteTableDisruptConnectivity19",
    "Effect": "Allow",
    "Action": "ec2:ModifyVpcEndpoint",
    "Resource": "arn:aws:ec2:*:*:route-table/*",
    "Condition": {
      "StringEquals": {
        "ec2:ResourceTag/managedByFIS": "true"
      }
    }
  },
  {
    "Sid": "TransitGatewayDisruptConnectivity1",
    "Effect": "Allow",
    "Action": [
      "ec2:DisassociateTransitGatewayRouteTable",
      "ec2:AssociateTransitGatewayRouteTable"
    ],
    "Resource": [
      "arn:aws:ec2:*:*:transit-gateway-route-table/*",
      "arn:aws:ec2:*:*:transit-gateway-attachment/*"
    ]
  },
  {
    "Sid": "S3CrossRegion1",
    "Effect": "Allow",
    "Action": "s3:ListAllMyBuckets",
    "Resource": "*"
  },

```

```

    {
      "Sid": "S3CrossRegion3",
      "Effect": "Allow",
      "Action": "s3:PauseReplication",
      "Resource": "arn:aws:s3:::*",
      "Condition": {
        "StringLike": {
          "s3:DestinationRegion": "*"
        }
      }
    },
    {
      "Sid": "S3CrossRegion4",
      "Effect": "Allow",
      "Action": [
        "s3:GetReplicationConfiguration",
        "s3:PutReplicationConfiguration"
      ],
      "Resource": "arn:aws:s3:::*",
      "Condition": {
        "BoolIfExists": {
          "s3:isReplicationPauseRequest": "true"
        }
      }
    },
    {
      "Sid": "DynamoDbPauseReplication",
      "Effect": "Allow",
      "Action": [
        "dynamodb:DescribeTable",
        "dynamodb:PutResourcePolicy",
        "dynamodb:GetResourcePolicy",
        "dynamodb>DeleteResourcePolicy"
      ],
      "Resource": [
        "arn:aws:dynamodb:*:*:table/*"
      ]
    },
    {
      "Sid": "DynamoDbMrscPauseReplication",
      "Effect": "Allow",
      "Action": [
        "dynamodb:InjectError"
      ],
    }
  ]
}

```

```

    "Resource": ["*"]
  },
  {
    "Sid": "ResolveResourcesViaTags",
    "Effect": "Allow",
    "Action": "tag:GetResources",
    "Resource": "*"
  },
  {
    "Sid": "MemDbCrossRegion",
    "Effect": "Allow",
    "Action": [
      "memorydb:DescribeMultiRegionClusters",
      "memorydb:PauseMultiRegionClusterReplication"
    ],
    "Resource": [
      "arn:aws:memorydb:*:*:multiregioncluster/*"
    ]
  },
  {
    "Sid": "DisruptVPCE1",
    "Effect": "Allow",
    "Action": "ec2:CreateSecurityGroup",
    "Resource": [
      "arn:aws:ec2:*:*:vpc/*",
      "arn:aws:ec2:*:*:security-group/*"
    ]
  },
  {
    "Sid": "DisruptVPCE2",
    "Effect": "Allow",
    "Action": "ec2:CreateTags",
    "Resource": "arn:aws:ec2:*:*:security-group/*",
    "Condition": {
      "StringEquals": {
        "ec2:CreateAction": "CreateSecurityGroup",
        "aws:RequestTag/managedByFIS": "true"
      }
    }
  },
  {
    "Sid": "DisruptVPCE3",
    "Effect": "Allow",
    "Action": [

```

```

        "ec2:DeleteSecurityGroup",
        "ec2:RevokeSecurityGroupEgress"
    ],
    "Resource": "arn:aws:ec2:*:*:security-group/*",
    "Condition": {
        "StringEquals": {
            "aws:ResourceTag/managedByFIS": "true"
        }
    }
},
{
    "Sid": "DisruptVPCE4",
    "Effect": "Allow",
    "Action": "vpce:AllowMultiRegion",
    "Resource": "arn:aws:ec2:*:*:vpc-endpoint/*"
},
{
    "Sid": "ModifyVPCE",
    "Effect": "Allow",
    "Action": "ec2:ModifyVpcEndpoint",
    "Resource": [
        "arn:aws:ec2:*:*:vpc-endpoint/*",
        "arn:aws:ec2:*:*:security-group/*"
    ]
}
]
}

```

## Scenario Content

The following content defines the scenario. This JSON can be saved and used to create an [experiment template](#) using the [create-experiment-template](#) command from the AWS Command Line Interface (AWS CLI). For the most recent version of the scenario, visit the scenario library in the FIS console.

```

{
    "targets": {
        "Transit-Gateway": {
            "resourceType": "aws:ec2:transit-gateway",
            "resourceTags": {
                "TgwTag": "TgwValue"
            },
            "selectionMode": "ALL"
        }
    }
}

```

```

    },
    "Subnet": {
      "resourceType": "aws:ec2:subnet",
      "resourceTags": {
        "SubnetKey": "SubnetValue"
      },
      "selectionMode": "ALL",
      "parameters": {}
    },
    "VPC-Endpoint": {
      "resourceType": "aws:ec2:vpc-endpoint",
      "resourceTags": {
        "DisruptPrivateLink": "Allowed"
      },
      "selectionMode": "ALL"
    },
    "S3-Bucket": {
      "resourceType": "aws:s3:bucket",
      "resourceTags": {
        "S3Impact": "Allowed"
      },
      "selectionMode": "ALL"
    },
    "DynamoDB-Global-Table": {
      "resourceType": "aws:dynamodb:global-table",
      "resourceTags": {
        "DisruptDynamoDb": "Allowed"
      },
      "selectionMode": "ALL"
    },
    "MemoryDB-Multi-Region-Cluster": {
      "resourceType": "aws:memorydb:multi-region-cluster",
      "resourceTags": {
        "DisruptMemoryDb": "Allowed"
      },
      "selectionMode": "ALL"
    }
  },
  "actions": {
    "Disrupt-Transit-Gateway-Connectivity": {
      "actionId": "aws:network:transit-gateway-disrupt-cross-region-
connectivity",
      "parameters": {
        "duration": "PT3H",

```

```

        "region": "eu-west-1"
    },
    "targets": {
        "TransitGateways": "Transit-Gateway"
    }
},
"Disrupt-Subnet-Connectivity": {
    "actionId": "aws:network:route-table-disrupt-cross-region-
connectivity",
    "parameters": {
        "duration": "PT3H",
        "region": "eu-west-1"
    },
    "targets": {
        "Subnets": "Subnet"
    }
},
"Disrupt-Vpc-Endpoint": {
    "actionId": "aws:network:disrupt-vpc-endpoint",
    "parameters": {
        "duration": "PT3H"
    },
    "targets": {
        "VPCEndpoints": "VPC-Endpoint"
    }
},
"Pause-S3-Replication": {
    "actionId": "aws:s3:bucket-pause-replication",
    "parameters": {
        "duration": "PT3H",
        "region": "eu-west-1"
    },
    "targets": {
        "Buckets": "S3-Bucket"
    }
},
"Pause-DynamoDB-Replication": {
    "actionId": "aws:dynamodb:global-table-pause-replication",
    "parameters": {
        "duration": "PT3H"
    },
    "targets": {
        "Tables": "DynamoDB-Global-Table"
    }
}

```

```
    },
    "Pause-MemoryDB-Multi-Region-Cluster-Replication": {
      "actionId": "aws:memorydb:multi-region-cluster-pause-replication",
      "parameters": {
        "duration": "PT3H",
        "region": "eu-west-1"
      },
      "targets": {
        "MultiRegionClusters": "MemoryDB-Multi-Region-Cluster"
      }
    }
  },
  "stopConditions": [
    {
      "source": "none"
    }
  ],
  "roleArn": "",
  "logConfiguration": {
    "logSchemaVersion": 2
  },
  "tags": {
    "Name": "Cross-Region: Connectivity"
  },
  "experimentOptions": {
    "accountTargeting": "single-account",
    "emptyTargetResolutionMode": "skip"
  },
  "description": "Block application network traffic from experiment Region to
target Region and pause cross-Region replication"
}
```

# Working with multi-account experiments for AWS FIS

You can create and manage multi-account experiment templates using the AWS FIS console or the command line. You create a multi-account experiment by specifying the account targeting experiment option as "multi-account", and adding target account configurations. After you create a multi-account experiment template, you can use it to run an experiment.

With a multi-account experiment, you can set up and run real-world failure scenarios on an application that spans multiple AWS accounts within a Region. You run multi-account experiments from an *orchestrator account* that impacts resources in multiple *target accounts*.

When you run a multi-account experiment, target accounts with affected resources will be notified via their AWS Health dashboards, providing awareness to users in the target accounts. With multi-account experiments, you can:

- Run real world failure scenarios on applications that span multiple accounts with the central controls and guardrails that AWS FIS provides.
- Control the effects of a multi-account experiment using IAM roles with fine-grained permissions and tags to define the scope of each target.
- Centrally view the actions AWS FIS takes in each account from the AWS Management Console and through AWS FIS logs.
- Monitor and audit API calls AWS FIS makes in each account with AWS CloudTrail.

This section helps you get started with multi-account experiments.

## Topics

- [Concepts for multi-account experiments](#)
- [Best practices for multi-account experiments](#)
- [Prerequisites for multi-account experiments](#)
- [Create a multi-account experiment template](#)
- [Update a target account configuration](#)
- [Delete a target account configuration](#)

## Concepts for multi-account experiments

The following are the key concepts for multi-account experiments:

- **Orchestrator account** - The orchestrator account acts as a central account to configure and manage the experiment in the AWS FIS Console, as well as to centralize logging. The orchestrator account owns the AWS FIS experiment template and experiment.
- **Target accounts** - A target account is an individual AWS account with resources that can be affected by an AWS FIS multi-account experiment.
- **Target account configurations** - You define the target accounts that are part of an experiment by adding target account configurations to the experiment template. A target account configuration is an element of the experiment template that is required for multi-account experiments. You define one for each target account by setting an AWS account ID, IAM role, and an optional description.

## Best practices for multi-account experiments

The following are best practices to using multi-account experiments:

- When you configure targets for multi-account experiments, we recommend targeting with consistent resource tags across all target accounts. An AWS FIS experiment will resolve resources with consistent tags in each target account. An action must resolve at least one target resource in any target account or will fail, except for experiments with `emptyTargetResolutionMode` set to `skip`. Action quotas apply per account. If you want to target resources by resource ARNs, the same single-account limit per action applies.
- When you target resources in one or more availability zones using parameters or filters, you should specify an *AZ ID*, not an AZ name. The AZ ID is a unique and consistent identifier for an Availability Zone across accounts. To learn how to find the AZ ID for the availability zones in your account, see [Availability Zone IDs for your AWS resources](#).

## Prerequisites for multi-account experiments

To use stop conditions for a multi-account experiment, you must first configure cross-account alarms. IAM roles are defined when you create a multi-account experiment template. You can create the necessary IAM roles before you create the template.

## Content

- [Permissions for multi-account experiments](#)
- [Stop conditions for multi-account experiments \(optional\)](#)
- [Safety levers for multi-account experiments \(optional\)](#)

## Permissions for multi-account experiments

Multi-account experiments use IAM *role chaining* to grant permissions to AWS FIS to take actions on resources in target accounts. For multi-account experiments, you set up IAM roles in each target account and the orchestrator account. These IAM roles require a trust relationship between the target accounts and the orchestrator account, and between the orchestrator account and AWS FIS.

The IAM roles for the target accounts contain the permissions required to take action on resources and are created for an experiment template by adding target account configurations. You will create an IAM role for the orchestrator account with permission to assume the roles of target accounts and establish a trust relationship with AWS FIS. This IAM role is used as the `roleArn` for the experiment template.

To learn more about role chaining, see [Roles Terms and concepts](#) in IAM's User Guide

In the following example, you will set up permissions for an orchestrator account A to run an experiment with `aws:ebs:pause-volume-io` in target account B.

1. In account B, create an IAM role with the permissions required to run the action. For permissions required for each action, see [Actions reference](#). The following example shows the permissions a target account grants to run the EBS Pause Volume IO action [the section called "aws:ebs:pause-volume-io"](#).

JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "ec2:DescribeVolumes"
      ],
      "Resource": "*"
    }
  ]
}
```

```

    },
    {
      "Effect": "Allow",
      "Action": [
        "ec2:PauseVolumeIO"
      ],
      "Resource": "arn:aws:ec2:us-east-1:123456789012:volume/*"
    },
    {
      "Effect": "Allow",
      "Action": [
        "tag:GetResources"
      ],
      "Resource": "*"
    }
  ]
}

```

- Next, add a trust policy in account B that creates a trust relationship with account A. Choose a name for the IAM role for account A, which you will create in step 3.

JSON

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Principal": {
        "AWS": "AccountIdA"
      },
      "Action": "sts:AssumeRole",
      "Condition": {
        "StringLike": {
          "sts:ExternalId":
            "arn:aws:fis:region:accountId:experiment/*"
        },
        "ArnEquals": {
          "aws:PrincipalArn":
            "arn:aws:iam::111122223333:role/role_name"
        }
      }
    }
  ]
}

```

```
}

```

- In account A, create an IAM role. This role name must match the role you specified in the trust policy in step 2. To target multiple accounts, you grant the orchestrator permissions to assume each role. The following example shows the permissions for account A to assume account B. If you have additional target accounts, you will add additional role ARNs to this policy. You can only have one role ARN per target account.

JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": "sts:AssumeRole",
      "Resource": [
        "arn:aws:iam::111122223333:role/role_name"
      ]
    }
  ]
}
```

- This IAM role for account A is used as the `roleArn` for the experiment template. The following example shows the trust policy required in the IAM role that grants AWS FIS permissions to assume account A, the orchestrator account.

JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Principal": {
        "Service": [
          "fis.amazonaws.com"
        ]
      },
      "Action": "sts:AssumeRole"
    }
  ]
}
```

```
}
```

You can also use Stacksets to provision multiple IAM roles at one time. To use CloudFormation StackSets, you will need to set up the necessary StackSet permissions in your AWS accounts. To learn more, see [Working with AWS CloudFormation StackSets](#).

## Stop conditions for multi-account experiments (optional)

A *stop condition* is a mechanism to stop an experiment if it reaches a threshold that you define as an alarm. To set up a stop condition for your multi-account experiment, you can use cross-account alarms. You must enable sharing in each target account to make the alarm available to the orchestrator account using read-only permissions. Once shared, you can combine metrics from different target accounts using Metric Math. Then, you can add this alarm as a stop condition for the experiment.

To learn more about cross-account dashboards, see [Enabling cross-account functionality in CloudWatch](#).

## Safety levers for multi-account experiments (optional)

Safety levers are used to stop all running experiments and prevent new experiments from starting. You may want to use the safety lever to prevent FIS experiments during certain time periods or in response to application health alarms. Every AWS account has a safety lever per AWS Region. When a safety lever is engaged, it impacts all experiments running in the same account and region as the safety lever. To stop and prevent multi-account experiments, the safety lever must be engaged in the same account and region where the experiments are running.

## Create a multi-account experiment template

To learn how to create an experiment template via the AWS Management Console

See [Create an experiment template](#).

To create an experiment template using the CLI

1. Open the AWS Command Line Interface
2. To create an experiment from a saved JSON file with the account targeting experiment option set to "multi-account" (for example, `my-template.json`), replace the placeholder values

in *italics* with your own values, and then run the following [create-experiment-template](#) command.

```
aws fis create-experiment-template --cli-input-json file://my-template.json
```

This will return the experiment template in the response. Copy the `id` from the response, which is the ID of the experiment template.

3. Run the [create-target-account-configuration](#) command to add a target account configuration to the experiment template. Replace the placeholder values in *italics* with your own values, using the `id` from step 2 as the value for the `--experiment-template-id` parameter, and then run the following. The `--description` parameter is optional. Repeat this step for each target account.

```
aws fis create-target-account-configuration --experiment-template-id EXTxxxxxxxxx --account-id 111122223333 --role-arn arn:aws:iam::111122223333:role/role-name --description "my description"
```

4. Run the [get-target-account-configuration](#) command to retrieve the details for a specific target account configuration.

```
aws fis get-target-account-configuration --experiment-template-id EXTxxxxxxxxx --account-id 111122223333
```

5. Once you have added all your target account configurations, you can run the [list-target-account-configurations](#) command to see that your target account configurations have been created.

```
aws fis list-target-account-configurations --experiment-template-id EXTxxxxxxxxx
```

You can also verify that you have added target account configurations by running the [get-experiment-template](#) command. The template will return a read-only field `targetAccountConfigurationsCount` that is a count of all the target account configurations on the experiment template.

6. When you are ready, you can run the experiment template using the [start-experiment](#) command.

```
aws fis start-experiment --experiment-template-id EXTxxxxxxxxx
```

## Update a target account configuration

You can update an existing target account configuration if you want to change the role ARN or description for the the account. When you update a target account configuration, the changes do not affect any running experiments that use the template.

### To update a target account configuration using the AWS Management Console

1. Open the AWS FIS console at <https://console.aws.amazon.com/fis/>.
2. In the navigation pane, choose **Experiment templates**
3. Select the experiment template, and choose **Actions, Update experiment template**.
4. In the side panel, choose **Step 3, Configure service access**.
5. Modify the **target account configurations**, and choose **Update experiment template**.
6. Select **Step 5, Review and create**.

### To update a target account configuration using the CLI

Run the [update-target-account-configuration](#) command to command, replacing the placeholder values in *italics* with your own values. The `--role-arn` and `--description` parameters are optional, and will not be updated if not included.

```
aws fis update-target-account-configuration --experiment-template-id EXTxxxxxxxxx  
--account-id 111122223333 --role-arn arn:aws:iam::111122223333:role/role-name --  
description "my description"
```

## Delete a target account configuration

If you no longer need a target account configuration, you can delete it. When you delete a target account configuration, any running experiments that use the template are not affected. The experiment continues to run until completed or stopped.

### To delete a target account configuration using the AWS Management Console

1. Open the AWS FIS console at <https://console.aws.amazon.com/fis/>.
2. In the navigation pane, choose **Experiment templates**.
3. Select the experiment template, and choose **Actions, Update**.

4. In the side panel, choose **Step 3, Configure service access**.
5. Under **Target account configurations**, select **Remove** for the target account Role ARN you want to delete.
6. Select **Step 5, Review and create**.
7. Review the template and choose **Update experiment template**. When prompted for confirmation, enter update and choose **Update experiment template**.

### To delete a target account configuration using the CLI

Run the [delete-target-account-configuration](#) command, replacing the placeholder values in *italics* with your own values.

```
aws fis update-target-account-configuration --experiment-template-id EXTxxxxxxxxx --  
account-id 111122223333
```

# Scheduling experiments

With AWS Fault Injection Service (FIS), you can perform fault injection experiments on your AWS workloads. These experiments run on templates that contain one or more actions to run on specified targets. You can now schedule your experiments as a one-time task or recurring tasks natively from the FIS Console. In addition to [scheduled rules](#), FIS now offers a new scheduling capability. FIS now integrates with EventBridge Scheduler and creates rules on your behalf. EventBridge Scheduler is a serverless scheduler that allows you to create, run, and manage tasks from one central, managed service.

## Important

Experiment Scheduler with AWS Fault Injection Service is not available in AWS GovCloud (US-East) and AWS GovCloud (US-West).

## Topics

- [Create a scheduler role](#)
- [Create an experiment schedule](#)
- [Update an experiment schedule](#)
- [Disable or delete an experiment schedule](#)

## Create a scheduler role

An execution role is an IAM role that AWS FIS assumes in order to interact with EventBridge scheduler and for Event Bridge scheduler to Start FIS Experiment. You attach permission policies to this role to grant EventBridge Scheduler access to invoke FIS Experiment. The following steps describe how to create a new execution role and a policy to allow EventBridge to Start an Experiment.

### Create scheduler role using the AWS CLI

This is IAM role that is needed for Event Bridge to be able to schedule experiment on behalf of the customer.

1. Copy the following assume role JSON policy and save it locally as `fis-execution-role.json`. This trust policy allows EventBridge Scheduler to assume the role on your behalf.

JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Principal": {
        "Service": "scheduler.amazonaws.com"
      },
      "Action": "sts:AssumeRole"
    }
  ]
}
```

2. From the AWS Command Line Interface (AWS CLI), enter the following command to create a new role. Replace `FisSchedulerExecutionRole` with the name you want to give this role.

```
aws iam create-role --role-name FisSchedulerExecutionRole --assume-role-policy-document file://fis-execution-role.json
```

If successful, you'll see the following output:

```
{
  "Role": {
    "Path": "/",
    "RoleName": "FisSchedulerExecutionRole",
    "RoleId": "AROAZL22PDN5A6WKRBNUN",
    "Arn": "arn:aws:iam::123456789012:role/FisSchedulerExecutionRole",
    "CreateDate": "2023-08-24T17:23:05+00:00",
    "AssumeRolePolicyDocument": {
      "Version": "2012-10-17",
      "Statement": [
        {
          "Effect": "Allow",
          "Principal": {
            "Service": "scheduler.amazonaws.com"
          },
          "Action": "sts:AssumeRole"
        }
      ]
    }
  }
}
```

```

        "Action": "sts:AssumeRole"
      }
    ]
  }
}

```

- To create a new policy that allows EventBridge Scheduler to invoke the experiment, copy the following JSON and save it locally as `fis-start-experiment-permissions.json`. The following policy allows EventBridge Scheduler to call the `fis:StartExperiment` action on all experiment templates in your account. Replace the `*` at the end of `"arn:aws:fis:*:*:experiment-template/*"` with the ID of your experiment template if you want to limit the role to a single experiment template.

JSON

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": "fis:StartExperiment",
      "Resource": [
        "arn:aws:fis:*:*:experiment-template/*",
        "arn:aws:fis:*:*:experiment/*"
      ]
    }
  ]
}

```

- Run the following command to create the new permission policy. Replace `FisSchedulerPolicy` with the name you want to give this policy.

```
aws iam create-policy --policy-name FisSchedulerPolicy --policy-document file://fis-start-experiment-permissions.json
```

If successful, you'll see the following output. Note the policy ARN. You use this ARN in the next step to attach the policy to our execution role.

```
{
```

```

"Policy": {
  "PolicyName": "FisSchedulerPolicy",
  "PolicyId": "ANPAZL22PDN5ESVUWXLBD",
  "Arn": "arn:aws:iam::123456789012:policy/FisSchedulerPolicy",
  "Path": "/",
  "DefaultVersionId": "v1",
  "AttachmentCount": 0,
  "PermissionsBoundaryUsageCount": 0,
  "IsAttachable": true,
  "CreateDate": "2023-08-24T17:34:45+00:00",
  "UpdateDate": "2023-08-24T17:34:45+00:00"
}
}

```

5. Run the following command to attach the policy to your execution role. Replace `your-policy-arn` with the ARN of the policy you created in the previous step. Replace `FisSchedulerExecutionRole` with the name of your execution role.

```

aws iam attach-role-policy --policy-arn your-policy-arn --role-name
FisSchedulerExecutionRole

```

The `attach-role-policy` operation doesn't return a response on the command line.

6. You can restrict the scheduler to only run AWS FIS experiment templates that have a specific tag value. For example, the following policy grants the `fis:StartExperiment` permission for all AWS FIS experiments, but restricts the scheduler to only run experiment templates that are tagged `Purpose=Schedule`.

JSON

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": "fis:StartExperiment",
      "Resource": "arn:aws:fis:*:*:experiment/*"
    },
    {
      "Effect": "Allow",
      "Action": "fis:StartExperiment",
      "Resource": "arn:aws:fis:*:*:experiment-template/*",

```

```

    "Condition": {
      "StringEquals": {
        "aws:ResourceTag/Purpose": "Schedule"
      }
    }
  ]
}

```

## Create an experiment schedule

Before you schedule an experiment, you need one or more [Experiment template components](#) for your schedule to invoke. You can use an existing AWS resource, or create a new one.

Once experiment template is created, click on **Actions** and select **Schedule experiment**. You will be redirected to schedule experiment page. The name of the schedule will be filled in for you.

Follow to the schedule pattern section and choose either one-time schedule or recurring. Fill in required input fields and navigate to permissions.

The screenshot shows the AWS console interface for configuring a schedule pattern. The top navigation bar includes the AWS logo, a search bar, and a list of services: S3, Amazon EventBridge, AWS FIS, CloudFormation, CloudWatch, IAM, EC2, AWS Organizations, and Application Composer. The main content area is titled "Schedule pattern" and contains the following sections:

- Occurrence**: A sub-section with an "Info" link. It states "You can define an one-time or recurrent schedule." and offers two radio button options: "One-time schedule" (which is selected) and "Recurring schedule".
- Date and time**: A sub-section with the instruction "The date and time to invoke the target." It includes three input fields: a date field with a calendar icon (placeholder: YYYY/MM/DD), a time field (placeholder: hh:mm), and a timezone dropdown menu (current selection: (UTC -04:00) America/New...). Below these fields are small labels: "YYYY/MM/DD", "Use 24-hour format timestamp (hh:mm)", and "Timezone".
- Flexible time window**: A sub-section with the instruction "If you choose a flexible time window, Scheduler invokes your schedule within the time window you specify. For example, if you choose 15 minutes, your schedule runs within 15 minutes after the schedule start time." It features a dropdown menu with the text "Select".
- Schedule state**: A sub-section with the instruction "You can choose not to enable the schedule now. You will be able to enable the schedule after it has been created." It contains a radio button labeled "Enable", which is selected.

Schedule state will be enabled by default. Note: if you disable **schedule state**, the experiment will not be scheduled even if you create a schedule.

AWS FIS Experiment Scheduler is built on top of [EventBridge Scheduler](#). You can refer the documentation for the various [schedule types supported](#).

## To update schedule using the console

1. Open the [AWS FIS console](#).
2. In the left navigation pane, choose **Experiment Templates**.
3. Choose **Experiment Template** for which you want to create the schedule.
4. Click **Actions**, and select **Schedule Experiment** from the dropdown.
  - a. Under **Schedule name**, name is auto populated.
  - b. Under **Schedule pattern**, select **Recurring schedule**.
  - c. Under **Schedule type**, you can select a **Rate-based schedule**, see [schedule types](#).
  - d. Under **Rate expression**, choose a rate that is slower than the execution time of your experiment, e.g. **5 minutes**.
  - e. Under **Timeframe**, select your **Time Zone**.
  - f. Under **Start Date and Time**, specify a start date and time.
  - g. Under **End Date and Time**, specify an end date and time.
  - h. Under **Schedule State**, toggle the **Enable Schedule Option**.
  - i. Under **Permissions**, select **Use existing role**, and then search for `FisSchedulerExecutionRole`.
  - j. Choose **Next**.
5. Select **Review and create schedule**, review your scheduler details, and then choose **Create schedule**.

## Update an experiment schedule

You can update an experiment schedule so that it occurs at a specific date and time that suits you.

To update an experiment execution using the console

1. Open the [Amazon FIS console](#).
2. In the navigation pane, choose **Experiment Templates**.
3. Choose **Resource type: Experiment Template** for which a schedule is already created.
4. Click on the Experiment ID for the template. Then navigate to schedules Tab.

5. Check if there is a existing schedule associated with the experiment. Select the schedule associated and Click the button **Update Schedule**.

## Disable or delete an experiment schedule

To stop an experiment from executing or running on a schedule, you can delete or disable the rule. The following steps walk you through how to delete or disable an Experiment Execution using the AWS console.

To delete or disable a rule

1. Open the [Amazon FIS console](#).
2. In the navigation pane, choose **Experiment Templates**.
3. Choose **Resource type: Experiment Template** for which a schedule is already created.
4. Click on the Experiment ID for the template. Then navigate to schedules Tab.
5. Check if there is a existing schedule associated with the experiment. Select the schedule associated and Click the button **Update Schedule**.
6. Do one of the following:
  - a. To delete the schedule, select the button next to the rule **Delete Schedule**. Type delete and click the **Delete Schedule** button.
  - b. To disable the schedule, select the button next to the rule **Disable Schedule**. Type disable and click the **Disable Schedule** button.

# Safety Levels for AWS FIS

Safety levels are used to stop all running experiments and prevent new experiments from starting. You may want to use the safety lever to prevent FIS experiments during certain time periods or in response to application health alarms. Every AWS account has a safety lever per AWS Region.

For in-progress experiments that are stopped by the safety lever, you only pay for the action duration that ran before the experiment was stopped. Experiments that are prevented from starting will not incur any costs. The following sections provide information on how to get started using safety levers.

## Topics

- [Concepts for safety levers](#)
- [Working with safety levers](#)

## Concepts for safety levers

A safety lever can be engaged or disengaged.

- If disengaged, FIS experiments are allowed. By default, safety levers are disengaged.
- If engaged, in-progress experiments are stopped and no new experiments are allowed to start.

An experiment affected by a safety lever will end in one of the following statuses:

- **Stopped**, if the experiment was running when the safety lever was engaged.
- **Cancelled**, if the experiment was started when the safety lever was already engaged.

You cannot resume or rerun an experiment that has been stopped or cancelled. However, you can start a new experiment using the same experiment template once the safety lever is disengaged.

## Safety lever resource

Safety lever is a resource defined by the Amazon Resource Name (ARN). Safety levers include the following parameters:

- **Status**, which is either engaged or disengaged.

- **Reason**, which is a string input by the user to log why the safety lever status was changed.

## Working with safety levers

This section will detail how to view, engage, and disengage safety levers using the AWS FIS console or the command line.

### Viewing a safety lever

You can view the state of your safety lever for your account and region by following the steps below.

#### To view a safety lever using the console

1. [Open the AWS FIS console](#)
2. In the navigation pane, choose **Experiments**.
3. If the safety lever is engaged, you will see an alert banner across the top of the page. If there is no alert banner, the safety lever is disengaged.

#### To view a safety lever using the CLI

- Use the following command:

```
aws fis get-safety-lever --id "default"
```

A safety lever can be in one of the following states:

- **Disengaged** - The safety lever is not impacting any experiments. Experiments can run freely. Safety levers are disengaged by default.
- **Engaging** - The safety lever is changing from disengaged to engaged. There may still be experiments that have not been stopped yet. The safety lever cannot be changed while in this state.
- **Engaged** - The safety lever is active and no experiments are running. Any new experiments that attempt to start while the safety lever is engaged will be cancelled.

## Engaging a safety lever

### To engage a safety lever using the console

1. [Open the AWS FIS console](#)
2. In the navigation pane, choose **Experiments**.
3. Choose the **Stop all experiments** button.
4. Enter the reason for engaging the safety lever.
5. Choose **Confirm**.

### To engage a safety lever using the CLI

- Use the following command. Fill in the reason field with your own response.

```
aws fis update-safety-lever-state --id "default" --state  
"status=engaged,reason=xxxxx"
```

## Disengaging a safety lever

### To disengage a safety lever using the console

1. [Open the AWS FIS console](#)
2. In the navigation pane, choose **Experiments**.
3. Choose the **Disengage safety lever** button.
4. Enter the reason for disengaging the safety lever.
5. Choose **Confirm**.

### To disengage a safety lever using the CLI

- Use the following command:

```
aws fis update-safety-lever-state --id "default" --state  
"status=disengaged,reason=recovered"
```

# Monitoring AWS FIS experiments

You can use the following tools to monitor the progress and impact of your AWS Fault Injection Service (AWS FIS) experiments.

## AWS FIS console and AWS CLI

Use the AWS FIS console or the AWS CLI to monitor the progress of a running experiment. You can view the status of each action in the experiment, and the results of each action. For more information, see [the section called "View your experiments"](#).

## CloudWatch usage metrics and alarms

Use CloudWatch usage metrics to provide visibility into your account's usage of resources. AWS FIS usage metrics correspond to AWS service quotas. You can configure alarms that alert you when your usage approaches a service quota. For more information, see [Monitor using CloudWatch](#).

You can also create stop conditions for your AWS FIS experiments by creating CloudWatch alarms that define when an experiment goes out of bounds. When the alarm is triggered, the experiment stops. For more information, see [Stop conditions](#). For more information about creating CloudWatch alarms, see [Create a CloudWatch Alarm Based on a Static Threshold](#) and [Creating a CloudWatch Alarm Based on Anomaly Detection](#) in the *Amazon CloudWatch User Guide*.

## AWS FIS experiment logging

Enable experiment logging to capture detailed information about your experiment as it runs. For more information see [Experiment logging](#).

## Experiment state change events

Amazon EventBridge enables you to respond automatically to system events or resource changes. AWS FIS emits a notification when the state of an experiment changes. You can create rules for the events that you are interested in that specify the automated action to take when an event matches a rule. For example, sending a notification to an Amazon SNS topic or invoking a Lambda function. For more information, see [Monitor using EventBridge](#).

## CloudTrail logs

Use AWS CloudTrail to capture detailed information about the calls made to the AWS FIS API and store them as log files in Amazon S3. CloudTrail also logs calls made to service APIs for the

resources on which you're running experiments. You can use these CloudTrail logs to determine which calls were made, the source IP address where the call came from, who made the call, when the call was made, and so on.

## AWS Health Dashboard Notifications

AWS Health provides ongoing visibility into your resource performance and the availability of your AWS services and accounts. When you start an experiment, AWS FIS emits a notification to your AWS Health Dashboard. The notification is present for the duration of the experiment in each account that contains resources targeted in an experiment, including multi-account experiments. Multi-account experiments with only actions that do not include targets, such as `aws:ssm:start-automation-execution` and `aws:fis:wait`, will not emit a notification. Information about the role used to allow the experiment will be listed under **Affected resources**. To learn more about the AWS Health Dashboard, see [AWS Health Dashboard](#) in the AWS Health User Guide.

### Note

AWS Health delivers events on a *best effort* basis.

## Monitor AWS FIS usage metrics using Amazon CloudWatch

You can use Amazon CloudWatch to monitor the impact of AWS FIS experiments on targets. You can also monitor your AWS FIS usage.

For more information about viewing the state of an experiment, see [View your experiments](#).

## Monitor AWS FIS experiments

As you plan your AWS FIS experiments, identify the CloudWatch metrics that you can use to identify the baseline or "steady state" for the target resource types for the experiment. After you start an experiment, you can monitor those CloudWatch metrics for the targets selected through the experiment template.

For more information about the available CloudWatch metrics for a target resource type supported by AWS FIS, see the following:

- [Monitor your instances using CloudWatch](#)
- [Amazon ECS CloudWatch metrics](#)

- [Monitoring Amazon RDS metrics using CloudWatch](#)
- [Monitoring Run Command metrics using CloudWatch](#)

## AWS FIS usage metrics

You can use CloudWatch usage metrics to provide visibility into your account's usage of resources. Use these metrics to visualize your current service usage on CloudWatch graphs and dashboards.

AWS FIS usage metrics correspond to AWS service quotas. You can configure alarms that alert you when your usage approaches a service quota. For more information about CloudWatch alarms, see the [Amazon CloudWatch User Guide](#).

AWS FIS publishes the following metric in the **AWS/Usage** namespace.

Metric	Description
ResourceCount	The total number of the specified resource running on your account. The resource is defined by the dimensions associated with the metric.

The following dimensions are used to refine the usage metrics that are published by AWS FIS.

Dimension	Description
Service	The name of the AWS service containing the resource. For AWS FIS usage metrics, the value for this dimension is FIS.
Type	The type of entity that is being reported. Currently, the only valid value for AWS FIS usage metrics is Resource.
Resource	The type of resource that is running. The possible values are ExperimentTemplates for experiment templates, and ActiveExperiments for active experiments.

Dimension	Description
Class	This dimension is reserved for future use.

## Monitor AWS FIS experiments using Amazon EventBridge

When the state of an experiment changes, AWS FIS emits a notification. These notifications are made available as events through Amazon EventBridge (formerly called CloudWatch Events). AWS FIS emits these events on a best effort basis. Events are delivered to EventBridge in near real time.

With EventBridge, you can create rules that trigger programmatic actions in response to an event. For example, you can configure a rule that invokes an SNS topic to send an email notification or invokes a Lambda function to take some action.

For more information about EventBridge, see [Getting started with Amazon EventBridge](#) in the *Amazon EventBridge User Guide*.

The following is the syntax of an experiment state change event:

```
{
  "version": "0",
  "id": "12345678-1234-1234-1234-123456789012",
  "detail-type": "FIS Experiment State Change",
  "source": "aws.fis",
  "account": "123456789012",
  "time": "yyyy-mm-ddThh:mm:ssZ",
  "region": "region",
  "resources": [
    "arn:aws:fis:region:account_id:experiment/experiment-id"
  ],
  "detail": {
    "experiment-id": "EXPabcd1efg2HIJKL3",
    "experiment-template-id": "EXTa1b2c3de5f6g7h",
    "new-state": {
      "status": "new_value",
      "reason": "reason_string"
    },
    "old-state": {
      "status": "old_value",
      "reason": "reason_string"
    }
  }
}
```

```
    }  
  }  
}
```

`experiment-id`

The ID of the experiment whose state changed.

`experiment-template-id`

The ID of the experiment template used by the experiment.

`new_value`

The new state of the experiment. The possible values are:

- `completed`
- `failed`
- `initiating`
- `running`
- `stopped`
- `stopping`

`old_value`

The previous state of the experiment. The possible values are:

- `initiating`
- `pending`
- `running`
- `stopping`

## Experiment logging for AWS FIS

You can use experiment logging to capture detailed information about your experiment as it runs.

You are charged for experiment logging based on the costs associated with each log destination type. For more information, see [Amazon CloudWatch Pricing](#) (under **Paid Tier, Logs, Vended Logs**) and [Amazon S3 Pricing](#).

## Permissions

You must grant AWS FIS permissions to send logs to each log destination that you configure. For more information, see the following in the *Amazon CloudWatch Logs User Guide*:

- [Logs sent to CloudWatch Logs](#)
- [Logs sent to Amazon S3](#)

## Log schema

The following is the schema used in experiment logging. The current schema version is 2. The fields for details depend on the value of `log_type`. The fields for `resolved_targets` depend on the value of `target_type`. For more information, see [the section called "Example log records"](#).

```
{
  "id": "EXP123abc456def789",
  "log_type": "experiment-start | target-resolution-start | target-resolution-detail
| target-resolution-end | action-start | action-error | action-end | experiment-end",
  "event_timestamp": "yyyy-mm-ddThh:mm:ssZ",
  "version": "2",
  "details": {
    "account_id": "123456789012",
    "action_end_time": "yyyy-mm-ddThh:mm:ssZ",
    "action_id": "String",
    "action_name": "String",
    "action_start_time": "yyyy-mm-ddThh:mm:ssZ",
    "action_state": {
      "status": "pending | initiating | running | completed | cancelled |
stopping | stopped | failed",
      "reason": "String"
    },
    "action_targets": "String to string map",
    "error_information": "String",
    "experiment_end_time": "yyyy-mm-ddThh:mm:ssZ",
    "experiment_state": {
      "status": "pending | initiating | running | completed | stopping | stopped
| failed",
      "reason": "String"
    },
    "experiment_start_time": "yyyy-mm-ddThh:mm:ssZ",
    "experiment_template_id": "String",
```

```
    "page": Number,
    "parameters": "String to string map",
    "resolved_targets": [
      {
        "field": "value"
      }
    ],
    "resolved_targets_count": Number,
    "status": "failed | completed",
    "target_name": "String",
    "target_resolution_end_time": "yyyy-mm-ddThh:mm:ssZ",
    "target_resolution_start_time": "yyyy-mm-ddThh:mm:ssZ",
    "target_type": "String",
    "total_pages": Number,
    "total_resolved_targets_count": Number
  }
}
```

## Release notes

- Version 2 introduces:
  - The `target_type` field and changes the `resolved_targets` field from a list of ARNs to a list of objects. The valid fields for the `resolved_targets` object depend on the value of `target_type`, which is the [resource type](#) of the targets.
  - The `action-error` and `target-resolution-detail` event types which add the `account_id` field.
- Version 1 is the initial release.

## Log destinations

AWS FIS supports log delivery to the following destinations:

- An Amazon S3 bucket
- An Amazon CloudWatch Logs log group

### S3 log delivery

The logs are delivered to the following location.

```
bucket-and-optional-prefix/AWSLogs/account-id/fis/region/experiment-id/YYYY/MM/DD/account-id_awsfislogs_region_experiment-id_YYYYMMDDHHMMZ_hash.log
```

It can take several minutes before the logs are delivered to the bucket.

## CloudWatch Logs log delivery

The logs are delivered to a log stream named `/aws/fis/experiment-id`.

Logs are delivered to the log group in less than one minute.

## Example log records

The following are example log records for an experiment that runs the `aws:ec2:reboot-instances` action on an EC2 instance selected at random.

### Records

- [experiment-start](#)
- [target-resolution-start](#)
- [target-resolution-detail](#)
- [target-resolution-end](#)
- [action-start](#)
- [action-end](#)
- [action-error](#)
- [experiment-end](#)

### experiment-start

The following is an example record for the `experiment-start` event.

```
{
  "id": "EXPhjAXCGY78HV2a4A",
  "log_type": "experiment-start",
  "event_timestamp": "2023-05-31T18:50:45Z",
  "version": "2",
  "details": {
    "experiment_template_id": "EXTCDh1M8HHkhxoaQ",
```

```
    "experiment_start_time": "2023-05-31T18:50:43Z"  
  }  
}
```

## target-resolution-start

The following is an example record for the `target-resolution-start` event.

```
{  
  "id": "EXPhjAXCGY78HV2a4A",  
  "log_type": "target-resolution-start",  
  "event_timestamp": "2023-05-31T18:50:45Z",  
  "version": "2",  
  "details": {  
    "target_resolution_start_time": "2023-05-31T18:50:45Z",  
    "target_name": "EC2InstancesToReboot"  
  }  
}
```

## target-resolution-detail

The following is an example record for the `target-resolution-detail` event. If target resolution fails, the record also includes the `error_information` field.

```
{  
  "id": "EXPhjAXCGY78HV2a4A",  
  "log_type": "target-resolution-detail",  
  "event_timestamp": "2023-05-31T18:50:45Z",  
  "version": "2",  
  "details": {  
    "target_resolution_end_time": "2023-05-31T18:50:45Z",  
    "target_name": "EC2InstancesToReboot",  
    "target_type": "aws:ec2:instance",  
    "account_id": "123456789012",  
    "resolved_targets_count": 2,  
    "status": "completed"  
  }  
}
```

## target-resolution-end

If target resolution fails, the record also includes the `error_information` field. If `total_pages` is greater than 1, the number of resolved targets exceeded the size limit for one record. There are additional `target-resolution-end` records that contain the remaining resolved targets.

The following is example record for the `target-resolution-end` event for an EC2 action.

```
{
  "id": "EXPhjAXCGY78HV2a4A",
  "log_type": "target-resolution-end",
  "event_timestamp": "2023-05-31T18:50:45Z",
  "version": "2",
  "details": {
    "target_resolution_end_time": "2023-05-31T18:50:46Z",
    "target_name": "EC2InstanceToReboot",
    "target_type": "aws:ec2:instance",
    "resolved_targets": [
      {
        "arn": "arn:aws:ec2:us-east-1:123456789012:instance/
i-0f7ee2abffc330de5"
      }
    ],
    "page": 1,
    "total_pages": 1
  }
}
```

The following is example record for the `target-resolution-end` event for an EKS action.

```
{
  "id": "EXP24YfiucfyVPJpEJn",
  "log_type": "target-resolution-end",
  "event_timestamp": "2023-05-31T18:50:45Z",
  "version": "2",
  "details": {
    "target_resolution_end_time": "2023-05-31T18:50:46Z",
    "target_name": "myPods",
    "target_type": "aws:eks:pod",
    "resolved_targets": [
      {
        "pod_name": "example-696fb6498b-sxhw5",
        "namespace": "default",
        "cluster_arn": "arn:aws:eks:us-east-1:123456789012:cluster/fis-demo-
cluster",

```

```
        "target_container_name": "example"
      }
    ],
    "page": 1,
    "total_pages": 1
  }
}
```

## action-start

The following is an example record for the `action-start` event. If the experiment template specifies parameters for the action, the record also includes the `parameters` field.

```
{
  "id": "EXPhjAXCGY78HV2a4A",
  "log_type": "action-start",
  "event_timestamp": "2023-05-31T18:50:56Z",
  "version": "2",
  "details": {
    "action_name": "Reboot",
    "action_id": "aws:ec2:reboot-instances",
    "action_start_time": "2023-05-31T18:50:56Z",
    "action_targets": {"Instances": "EC2InstancesToReboot"}
  }
}
```

## action-error

The following is an example record for the `action-error` event. This event is only returned when an action fails. It is returned for each account where the action fails.

```
{
  "id": "EXPhjAXCGY78HV2a4A",
  "log_type": "action-error",
  "event_timestamp": "2023-05-31T18:50:56Z",
  "version": "2",
  "details": {
    "action_name": "pause-io",
    "action_id": "aws:ebs:pause-volume-io",
    "account_id": "123456789012",
    "action_state": {
      "status": "failed",
    }
  }
}
```

```
        "reason": "Unable to start Pause Volume IO. Target volumes must be attached
to an instance type based on the Nitro system. VolumeId(s): [vol-1234567890abcdef0]:"
    }
}
}
```

## action-end

The following is an example record for the action-end event.

```
{
  "id": "EXPhjAXCGY78HV2a4A",
  "log_type": "action-end",
  "event_timestamp": "2023-05-31T18:50:56Z",
  "version": "2",
  "details": {
    "action_name": "Reboot",
    "action_id": "aws:ec2:reboot-instances",
    "action_end_time": "2023-05-31T18:50:56Z",
    "action_state": {
      "status": "completed",
      "reason": "Action was completed."
    }
  }
}
```

## experiment-end

The following is an example record for the experiment-end event.

```
{
  "id": "EXPhjAXCGY78HV2a4A",
  "log_type": "experiment-end",
  "event_timestamp": "2023-05-31T18:50:57Z",
  "version": "2",
  "details": {
    "experiment_end_time": "2023-05-31T18:50:57Z",
    "experiment_state": {
      "status": "completed",
      "reason": "Experiment completed"
    }
  }
}
```

```
}
```

## Enable experiment logging

Experiment logging is disabled by default. To receive experiment logs for an experiment, you must create the experiment from an experiment template with logging enabled. The first time that you run an experiment that is configured to use a destination that hasn't been used previously for logging, we delay the experiment to configure log delivery to this destination, which takes about 15 seconds.

### To enable experiment logging using the console

1. Open the AWS FIS console at <https://console.aws.amazon.com/fis/>.
2. In the navigation pane, choose **Experiment templates**.
3. Select the experiment template, and choose **Actions, Update experiment template**.
4. For **Logs**, configure the destination options. To send logs to an S3 bucket, choose **Send to an Amazon S3 bucket** and enter the bucket name and prefix. To send logs to CloudWatch Logs, choose **Send to CloudWatch Logs** and enter the log group.
5. Choose **Update experiment template**.

### To enable experiment logging using the AWS CLI

Use the [update-experiment-template](#) command and specify a log configuration.

## Disable experiment logging

If you no longer want to receive logs for your experiments, you can disable experiment logging.

### To disable experiment logging using the console

1. Open the AWS FIS console at <https://console.aws.amazon.com/fis/>.
2. In the navigation pane, choose **Experiment templates**.
3. Select the experiment template, and choose **Actions, Update experiment template**.
4. For **Logs**, clear **Send to an Amazon S3 bucket** and **Send to CloudWatch Logs**.
5. Choose **Update experiment template**.

### To disable experiment logging using the AWS CLI

Use the [update-experiment-template](#) command and specify an empty log configuration.

## Log API calls with AWS CloudTrail

AWS Fault Injection Service (AWS FIS) is integrated with AWS CloudTrail, a service that provides a record of actions taken by a user, a role, or an AWS service in AWS FIS. CloudTrail captures all API calls for AWS FIS as events. The calls that are captured include calls from the AWS FIS console and code calls to the AWS FIS API operations. If you create a trail, you can enable continuous delivery of CloudTrail events to an Amazon S3 bucket, including events for AWS FIS. 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 FIS, 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](#).

## Use CloudTrail

CloudTrail is enabled on your AWS account when you create the account. When activity occurs in AWS FIS, 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 FIS, 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:

- [Create a Trail for Your AWS Account](#)
- [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 FIS actions are logged by CloudTrail and are documented in the [AWS Fault Injection Service API Reference](#). For the experiment actions that are carried out on a target resource, view

the API reference documentation for the service that owns the resource. For example, for actions that are carried out on an Amazon EC2 instance, see the [Amazon EC2 API Reference](#).

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 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](#).

## Understand AWS FIS 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 is an example CloudTrail log entry for a call to the AWS FIS StopExperiment action.

```
{
  "eventVersion": "1.08",
  "userIdentity": {
    "type": "AssumedRole",
    "principalId": "AKIAIOSFODNN7EXAMPLE:jdoh",
    "arn": "arn:aws:sts::111122223333:assumed-role/example/jdoh",
    "accountId": "111122223333",
    "accessKeyId": "AKIAI44QH8DHBEXAMPLE",
    "sessionContext": {
      "sessionIssuer": {
        "type": "Role",
        "principalId": "AKIAIOSFODNN7EXAMPLE",
        "arn": "arn:aws:iam::111122223333:role/example",
        "accountId": "111122223333",
        "userName": "example"
      },
      "webIdFederationData": {},
      "attributes": {
```

```
        "creationDate": "2020-12-03T09:40:42Z",
        "mfaAuthenticated": "false"
    }
},
"eventTime": "2020-12-03T09:44:20Z",
"eventSource": "fis.amazonaws.com",
"eventName": "StopExperiment",
"awsRegion": "us-east-1",
"sourceIPAddress": "192.51.100.25",
"userAgent": "Boto3/1.22.9 Python/3.8.13 Linux/5.4.186-113.361.amzn2int.x86_64
BotoCore/1.25.9",
"requestParameters": {
    "clientToken": "1234abc5-6def-789g-012h-ijklm34no56p",
    "experimentTemplateId": "ABCDE1fgHIJkLmNop",
    "tags": {}
},
"responseElements": {
    "experiment": {
        "actions": {
            "exampleAction1": {
                "actionId": "aws:ec2:stop-instances",
                "duration": "PT10M",
                "state": {
                    "reason": "Initial state",
                    "status": "pending"
                },
                "targets": {
                    "Instances": "exampleTag1"
                }
            },
            "exampleAction2": {
                "actionId": "aws:ec2:stop-instances",
                "duration": "PT10M",
                "state": {
                    "reason": "Initial state",
                    "status": "pending"
                },
                "targets": {
                    "Instances": "exampleTag2"
                }
            }
        }
    },
    "creationTime": 1605788649.95,
```

```
"endTime": 1606988660.846,
"experimentTemplateId": "ABCDE1fgHIJkLmNop",
"id": "ABCDE1fgHIJkLmNop",
"roleArn": "arn:aws:iam::111122223333:role/AllowFISActions",
"startTime": 1605788650.109,
"state": {
  "reason": "Experiment stopped",
  "status": "stopping"
},
"stopConditions": [
  {
    "source": "aws:cloudwatch:alarm",
    "value": "arn:aws:cloudwatch:us-east-1:111122223333:alarm:example"
  }
],
"tags": {},
"targets": {
  "ExampleTag1": {
    "resourceTags": {
      "Example": "tag1"
    },
    "resourceType": "aws:ec2:instance",
    "selectionMode": "RANDOM(1)"
  },
  "ExampleTag2": {
    "resourceTags": {
      "Example": "tag2"
    },
    "resourceType": "aws:ec2:instance",
    "selectionMode": "RANDOM(1)"
  }
}
},
"requestID": "1abcd23e-f4gh-567j-klm8-9np01q234r56",
"eventID": "1234a56b-c78d-9e0f-g1h2-34jk56m7n890",
"readOnly": false,
"eventType": "AwsApiCall",
"managementEvent": true,
"recipientAccountId": "111122223333",
"eventCategory": "Management"
}
```

The following is an example CloudTrail log entry for an API action that AWS FIS invoked as part of an experiment that includes the `aws:ssm:send-command` AWS FIS action. The `userIdentity` element reflects a request made with temporary credentials obtained by assuming a role. The name of the assumed role appears in `userName`. The ID of the experiment, `EXP21nT17WMzA6dnUgz`, appears in `principalId` and as part of the ARN of the assumed role.

```
{
  "eventVersion": "1.08",
  "userIdentity": {
    "type": "AssumedRole",
    "principalId": "AROATZZZ4JPIXUEXAMPLE:EXP21nT17WMzA6dnUgz",
    "arn": "arn:aws:sts::111122223333:assumed-role/AllowActions/
EXP21nT17WMzA6dnUgz",
    "accountId": "111122223333",
    "accessKeyId": "AKIAI44QH8DHBEXAMPLE",
    "sessionContext": {
      "sessionIssuer": {
        "type": "Role",
        "principalId": "AROATZZZ4JPIXUEXAMPLE",
        "arn": "arn:aws:iam::111122223333:role/AllowActions",
        "accountId": "111122223333",
        "userName": "AllowActions"
      },
      "webIdFederationData": {},
      "attributes": {
        "creationDate": "2022-05-30T13:23:19Z",
        "mfaAuthenticated": "false"
      }
    },
    "invokedBy": "fis.amazonaws.com"
  },
  "eventTime": "2022-05-30T13:23:19Z",
  "eventSource": "ssm.amazonaws.com",
  "eventName": "ListCommands",
  "awsRegion": "us-east-2",
  "sourceIPAddress": "fis.amazonaws.com",
  "userAgent": "fis.amazonaws.com",
  "requestParameters": {
    "commandId": "51dab97f-489b-41a8-a8a9-c9854955dc65"
  },
  "responseElements": null,
  "requestID": "23709ced-c19e-471a-9d95-cf1a06b50ee6",
  "eventID": "145fe5a6-e9d5-45cc-be25-b7923b950c83",
}
```

```
"readOnly": true,  
"eventType": "AwsApiCall",  
"managementEvent": true,  
"recipientAccountId": "111122223333",  
"eventCategory": "Management"  
}
```

# Troubleshooting AWS FIS

To troubleshoot errors, AWS FIS returns detailed errors from the `GetExperiment` API and FIS experiment logs. Errors are returned as part of the experiment state when the status of the experiment is **failed**. When multiple actions fail, the first failed action is returned as the experiment error. You can review your FIS experiment logs for any other errors. To learn how to log and monitor AWS FIS experiments, see [Monitoring AWS FIS experiments](#).

Depending on the type of failure, you may receive one of the following errors:

- **Reason:** Detailed description of the specific failure. Reason values should not be used for automation, as they are subject to change.
- **Code:** The type of failure. Code values should not be used for automation, as they are subject to change, unless otherwise specified in table below.
- **Location:** Context for the section of the experiment template that failed, such as the action or target.
- **Account ID:** The AWS Account where the failure occurred.

## Error codes

Error Code	Code Description
<code>ConfigurationFailure</code>	The action, target, experiment, or log is not configured correctly. Check the error location and ensure that the parameters and configurations are correct.
<code>DependentServiceFailure</code>	There was a failure from another AWS service. Try running the experiment again.
<code>InternalFailure</code>	An internal error occurred when running the experiment. You can automate based on this error code.
<code>InvalidTarget</code>	A target was unable to be resolved during target resolution or at the start of an action.

Error Code	Code Description
	<p>This may be due to one of the following reasons:</p> <ul style="list-style-type: none"><li>• The target does not exist, for example if it was deleted or if the ARN is incorrect.</li><li>• There is a tag for your target that does not resolve any resources.</li><li>• There is an action that is not linked to a target.</li></ul> <p>To troubleshoot, review your logs to identify which targets are not resolved. Check that all actions are linked to targets, and that your resource ID or tags exist and have not been misspelled.</p>
AuthorizationFailure	<p>There are two primary causes for experiment failure due to permissions errors:</p> <ul style="list-style-type: none"><li>• The IAM role you are targeting does not have the proper permissions to resolve targets or take action on your resources . To fix this error, review the permissions required for your actions in the <a href="#">FIS Actions Reference</a> and add them to your experiment IAM role.</li><li>• The <a href="#">AWS Service-linked role</a> (SLR) creation for FIS was denied by a <a href="#">service control policy</a> (SCP) in your organization. FIS uses the SLR to manage monitoring and resource selection for experiments. For more information, see <a href="#">Service-linked role permissions for AWS FIS</a>.</li></ul>

Error Code	Code Description
QuotaExceededFailure	The quota for the resource type has been exceeded. To determine if the quota can be increased, see <a href="#">Quotas and limitations for AWS Fault Injection Service</a> .

# Security in AWS Fault Injection Service

Cloud security at AWS is the highest priority. As an AWS customer, you benefit from data centers and network architectures that are 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 Fault Injection Service, 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.

This documentation helps you understand how to apply the shared responsibility model when using AWS FIS. The following topics show you how to configure AWS FIS 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 FIS resources.

## Contents

- [Data protection in AWS Fault Injection Service](#)
- [Identity and access management for AWS Fault Injection Service](#)
- [Infrastructure security in AWS Fault Injection Service](#)
- [Access AWS FIS using an interface VPC endpoint \(AWS PrivateLink\)](#)

## Data protection in AWS Fault Injection Service

The AWS [shared responsibility model](#) applies to data protection in AWS Fault Injection Service. 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 FIS 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.

## Encryption at rest

AWS FIS always encrypts your data at rest. Data in AWS FIS is encrypted at rest using transparent server-side encryption. This helps reduce the operational burden and complexity involved in protecting sensitive data. With encryption at rest, you can build security-sensitive applications that meet encryption compliance and regulatory requirements.

## Encryption in transit

AWS FIS encrypts data in transit between the service and other integrated AWS services. All data that passes between AWS FIS and integrated services is encrypted using Transport Layer Security (TLS). For more information about other integrated AWS services, see [Supported AWS services](#).

## Identity and access management for AWS Fault Injection Service

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 FIS resources. IAM is an AWS service that you can use with no additional charge.

### Contents

- [Audience](#)
- [Authenticating with identities](#)
- [Managing access using policies](#)
- [How AWS Fault Injection Service works with IAM](#)
- [AWS Fault Injection Service policy examples](#)
- [Use service-linked roles for AWS Fault Injection Service](#)
- [AWS managed policies for AWS Fault Injection Service](#)

## Audience

How you use AWS Identity and Access Management (IAM) differs, depending on the work that you do in AWS FIS.

**Service user** – If you use the AWS FIS service to do your job, then your administrator provides you with the credentials and permissions that you need. As you use more AWS FIS 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.

**Service administrator** – If you're in charge of AWS FIS resources at your company, you probably have full access to AWS FIS. It's your job to determine which AWS FIS 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.

**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 FIS.

## Authenticating with identities

Authentication is how you sign in to AWS using your identity credentials. You must be authenticated as the AWS account root user, an IAM user, or by assuming an IAM role.

You can sign in as a federated identity using credentials from an identity source like AWS IAM Identity Center (IAM Identity Center), single sign-on authentication, or Google/Facebook credentials. For more information about signing in, see [How to sign in to your AWS account](#) in the *AWS Sign-In User Guide*.

For programmatic access, AWS provides an SDK and CLI to cryptographically sign requests. For more information, see [AWS Signature Version 4 for API requests](#) in the *IAM User Guide*.

### AWS account root user

When you create an AWS account, you begin with one sign-in identity called the AWS account *root user* that has complete access to all AWS services and resources. We strongly recommend that you don't use the root user for everyday tasks. For tasks that require root user credentials, see [Tasks that require root user credentials](#) in the *IAM User Guide*.

### Federated identity

As a best practice, require human users to use federation with an identity provider to access AWS services using temporary credentials.

A *federated identity* is a user from your enterprise directory, web identity provider, or Directory Service that accesses AWS services using credentials from an identity source. Federated identities assume roles that provide temporary credentials.

For centralized access management, we recommend AWS IAM Identity Center. For more information, see [What is IAM Identity Center?](#) in the *AWS IAM Identity Center User Guide*.

### IAM users and groups

An [IAM user](#) is an identity with specific permissions for a single person or application. We recommend using temporary credentials instead of IAM users with long-term credentials. For more

information, see [Require human users to use federation with an identity provider to access AWS using temporary credentials](#) in the *IAM User Guide*.

An [IAM group](#) specifies a collection of IAM users and makes permissions easier to manage for large sets of users. For more information, see [Use cases for IAM users](#) in the *IAM User Guide*.

## IAM roles

An [IAM role](#) is an identity with specific permissions that provides temporary credentials. You can assume a role by [switching from a user to an IAM role \(console\)](#) or by calling an AWS CLI or AWS API operation. For more information, see [Methods to assume a role](#) in the *IAM User Guide*.

IAM roles are useful for federated user access, temporary IAM user permissions, cross-account access, cross-service access, and applications running on Amazon EC2. For more information, see [Cross account resource access in IAM](#) 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 defines permissions when associated with an identity or resource. AWS evaluates these policies when a principal makes a request. Most policies are stored in AWS as JSON documents. For more information about JSON policy documents, see [Overview of JSON policies](#) in the *IAM User Guide*.

Using policies, administrators specify who has access to what by defining which **principal** can perform **actions** on what **resources**, and under what **conditions**.

By default, users and roles have no permissions. An IAM administrator creates IAM policies and adds them to roles, which users can then assume. IAM policies define permissions regardless of the method used to perform the operation.

## Identity-based policies

Identity-based policies are JSON permissions policy documents that you attach to an identity (user, group, or role). These policies control what actions identities 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 *inline policies* (embedded directly into a single identity) or *managed policies* (standalone policies attached to multiple identities). To learn how to choose between

managed and inline policies, 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 include 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. You must [specify a principal](#) in a resource-based policy.

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.

## Other policy types

AWS supports additional policy types that can set the maximum permissions granted by more common policy types:

- **Permissions boundaries** – Set the maximum permissions that an identity-based policy can grant to an IAM entity. For more information, see [Permissions boundaries for IAM entities](#) in the *IAM User Guide*.
- **Service control policies (SCPs)** – Specify the maximum permissions for an organization or organizational unit in AWS Organizations. For more information, see [Service control policies](#) in the *AWS Organizations User Guide*.
- **Resource control policies (RCPs)** – Set the maximum available permissions for resources in your accounts. For more information, see [Resource control policies \(RCPs\)](#) in the *AWS Organizations User Guide*.
- **Session policies** – Advanced policies passed as a parameter when creating a temporary session for a role or federated user. 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*.

## How AWS Fault Injection Service works with IAM

Before you use IAM to manage access to AWS FIS, learn what IAM features are available to use with AWS FIS.

### IAM features you can use with AWS Fault Injection Service

IAM feature	AWS FIS support
<a href="#">Identity-based policies</a>	Yes
<a href="#">Resource-based policies</a>	No
<a href="#">Policy actions</a>	Yes
<a href="#">Policy resources</a>	Yes
<a href="#">Policy condition keys (service-specific)</a>	Yes
<a href="#">ACLs</a>	No
<a href="#">ABAC (tags in policies)</a>	Yes
<a href="#">Temporary credentials</a>	Yes
<a href="#">Principal permissions</a>	Yes
<a href="#">Service roles</a>	Yes
<a href="#">Service-linked roles</a>	Yes

To get a high-level view of how AWS FIS and other AWS services work with most IAM features, see [AWS services that work with IAM](#) in the *IAM User Guide*.

### Identity-based policies for AWS FIS

**Supports identity-based policies:** Yes

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*.

With IAM identity-based policies, you can specify allowed or denied actions and resources as well as the conditions under which actions are allowed or denied. To learn about all of the elements that you can use in a JSON policy, see [IAM JSON policy elements reference](#) in the *IAM User Guide*.

## Identity-based policy examples for AWS FIS

To view examples of AWS FIS identity-based policies, see [AWS Fault Injection Service policy examples](#).

## Resource-based policies within AWS FIS

### Supports resource-based policies: No

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.

To enable cross-account access, you can specify an entire account or IAM entities in another account as the principal in a resource-based policy. For more information, see [Cross account resource access in IAM](#) in the *IAM User Guide*.

## Policy actions for AWS FIS

### Supports policy actions: Yes

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. Include actions in a policy to grant permissions to perform the associated operation.

To see a list of AWS FIS actions, see [Actions defined by AWS Fault Injection Service](#) in the *Service Authorization Reference*.

Policy actions in AWS FIS use the following prefix before the action:

```
fis
```

To specify multiple actions in a single statement, separate them with commas.

```
"Action": [  
  "fis:action1",  
  "fis:action2"  
]
```

You can specify multiple actions using wildcards (\*). For example, to specify all actions that begin with the word List, include the following action:

```
"Action": "fis:List*"
```

## Policy resources for AWS FIS

**Supports policy resources:** Yes

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. As a best practice, specify a resource using its [Amazon Resource Name \(ARN\)](#). For actions that don't support resource-level permissions, use a wildcard (\*) to indicate that the statement applies to all resources.

```
"Resource": "*"
```

Some AWS FIS API actions support multiple resources. To specify multiple resources in a single statement, separate the ARNs with commas.

```
"Resource": [  
  "resource1",  
  "resource2"  
]
```

To see a list of AWS FIS resource types and their ARNs, see [Resource types defined by AWS Fault Injection Service](#) in the *Service Authorization Reference*. To learn with which actions you can specify the ARN of each resource, see [Actions defined by AWS Fault Injection Service](#).

## Policy condition keys for AWS FIS

**Supports service-specific policy condition keys:** Yes

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 specifies when statements execute based on defined criteria. 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. To see all AWS global condition keys, see [AWS global condition context keys](#) in the *IAM User Guide*.

To see a list of AWS FIS condition keys, see [Condition keys for AWS Fault Injection Service](#) in the *Service Authorization Reference*. To learn with which actions and resources you can use a condition key, see [Actions defined by AWS Fault Injection Service](#).

To view examples of AWS FIS identity-based policies, see [AWS Fault Injection Service policy examples](#).

## ACLs in AWS FIS

**Supports ACLs:** No

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.

## ABAC with AWS FIS

**Supports ABAC (tags in policies):** Yes

Attribute-based access control (ABAC) is an authorization strategy that defines permissions based on attributes called tags. You can attach tags to IAM entities and AWS resources, then design ABAC policies to allow operations when the principal's tag matches the tag on the resource.

To control access based on tags, you provide tag information in the [condition element](#) of a policy using the `aws:ResourceTag/key-name`, `aws:RequestTag/key-name`, or `aws:TagKeys` condition keys.

If a service supports all three condition keys for every resource type, then the value is **Yes** for the service. If a service supports all three condition keys for only some resource types, then the value is **Partial**.

For more information about ABAC, see [Define permissions with ABAC authorization](#) in the *IAM User Guide*. To view a tutorial with steps for setting up ABAC, see [Use attribute-based access control \(ABAC\)](#) in the *IAM User Guide*.

To view an example identity-based policy for limiting access to a resource based on the tags for that resource, see [Example: Use tags to control resource usage](#).

## Using temporary credentials with AWS FIS

**Supports temporary credentials:** Yes

Temporary credentials provide short-term access to AWS resources and are automatically created when you use federation or switch roles. AWS recommends that you dynamically generate temporary credentials instead of using long-term access keys. For more information, see [Temporary security credentials in IAM](#) and [AWS services that work with IAM](#) in the *IAM User Guide*.

## Cross-service principal permissions for AWS FIS

**Supports forward access sessions (FAS):** Yes

Forward access sessions (FAS) use the permissions of the principal calling an AWS service, combined with the requesting AWS service to make requests to downstream services. For policy details when making FAS requests, see [Forward access sessions](#).

## Service roles for AWS FIS

**Supports service roles:** Yes

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 roles for AWS FIS

**Supports service-linked roles:** Yes

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.

For details about creating or managing AWS FIS service-linked roles, see [Use service-linked roles for AWS Fault Injection Service](#).

## AWS Fault Injection Service policy examples

By default, users and roles don't have permission to create or modify AWS FIS resources. To grant users permission to perform actions on the resources that they need, an IAM administrator can create IAM policies.

To learn how to create an IAM identity-based policy by using these example JSON policy documents, see [Create IAM policies \(console\)](#) in the *IAM User Guide*.

For details about actions and resource types defined by AWS FIS, including the format of the ARNs for each of the resource types, see [Actions, resources, and condition keys for AWS Fault Injection Service](#) in the *Service Authorization Reference*.

### Contents

- [Policy best practices](#)
- [Example: Use the AWS FIS console](#)
- [Example: List available AWS FIS actions](#)
- [Example: Create an experiment template for a specific action](#)
- [Example: Start an experiment](#)
- [Example: Use tags to control resource usage](#)
- [Example: Delete an experiment template with a specific tag](#)
- [Example: Allow users to view their own permissions](#)
- [Example: Use condition keys for ec2:InjectApiError](#)
- [Example: Use condition keys for aws:s3:bucket-pause-replication](#)

### Policy best practices

Identity-based policies determine whether someone can create, access, or delete AWS FIS 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 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*.

## Example: Use the AWS FIS console

To access the AWS Fault Injection Service console, you must have a minimum set of permissions. These permissions must allow you to list and view details about the AWS FIS resources in your AWS account. If you create an identity-based policy that is more restrictive than the minimum required permissions, the console won't function as intended for entities (users or roles) with that policy.

You don't need to allow minimum console permissions for users that are making calls only to the AWS CLI or the AWS API. Instead, allow access to only the actions that match the API operation that they're trying to perform.

The following example policy grants permission to list and view all AWS FIS resources using AWS FIS console, but not to create, update, or delete them. It also grants permissions to view the available resources used by all AWS FIS actions that you could specify in an experiment template.

## JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "FISReadOnlyActions",
      "Effect": "Allow",
      "Action": [
        "fis:List*",
        "fis:Get*"
      ],
      "Resource": "*"
    },
    {
      "Sid": "AdditionalReadOnlyActions",
      "Effect": "Allow",
      "Action": [
        "ssm:Describe*",
        "ssm:Get*",
        "ssm:List*",
        "ec2:DescribeInstances",
        "rds:DescribeDBClusters",
        "ecs:DescribeClusters",
        "ecs:ListContainerInstances",
        "eks:DescribeNodegroup",
        "cloudwatch:DescribeAlarms",
        "iam:ListRoles"
      ],
      "Resource": "*"
    },
    {
      "Sid": "PermissionsToCreateServiceLinkedRole",
      "Effect": "Allow",
```

```

        "Action": "iam:CreateServiceLinkedRole",
        "Resource": "*",
        "Condition": {
            "StringEquals": {
                "iam:AWSServiceName": "fis.amazonaws.com"
            }
        }
    ]
}

```

## Example: List available AWS FIS actions

The following policy grants permission to list the available AWS FIS actions.

JSON

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "fis:ListActions"
      ],
      "Resource": "arn:aws:fis:*:*:action/*"
    }
  ]
}

```

## Example: Create an experiment template for a specific action

The following policy grants permission to create an experiment template for the action `aws:ec2:stop-instances`.

JSON

```

{
  "Version": "2012-10-17",

```

```

    "Statement": [
      {
        "Sid": "PolicyExample",
        "Effect": "Allow",
        "Action": [
          "fis:CreateExperimentTemplate"
        ],
        "Resource": [
          "arn:aws:fis:*:*:action/aws:ec2:stop-instances",
          "arn:aws:fis:*:*:experiment-template/*"
        ]
      },
      {
        "Sid": "PolicyPassRoleExample",
        "Effect": "Allow",
        "Action": [
          "iam:PassRole"
        ],
        "Resource": [
          "arn:aws:iam:*:111122223333:role/role-name"
        ]
      }
    ]
  }
}

```

## Example: Start an experiment

The following policy grants permission to start an experiment using the specified IAM role and experiment template. It also allows AWS FIS to create a service-linked role on the user's behalf. For more information, see [Use service-linked roles for AWS Fault Injection Service](#).

JSON

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "PolicyExample",
      "Effect": "Allow",
      "Action": [
        "fis:StartExperiment"
      ]
    }
  ]
}

```

```

    ],
    "Resource": [
      "arn:aws:fis:*:*:experiment-template/experiment-template-id",
      "arn:aws:fis:*:*:experiment/*"
    ]
  },
  {
    "Sid": "PolicyExampleforServiceLinkedRole",
    "Effect": "Allow",
    "Action": "iam:CreateServiceLinkedRole",
    "Resource": "*",
    "Condition": {
      "StringEquals": {
        "iam:AWSServiceName": "fis.amazonaws.com"
      }
    }
  }
]
}

```

## Example: Use tags to control resource usage

The following policy grants permission to run experiments from experiment templates that have the tag `Purpose=Test`. It does not grant permission to create or modify experiment templates, or run experiments using templates that do not have the specified tag.

JSON

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": "fis:StartExperiment",
      "Resource": "arn:aws:fis:*:*:experiment-template/*",
      "Condition": {
        "StringEquals": {
          "aws:ResourceTag/Purpose": "Test"
        }
      }
    }
  ]
}

```

```
]
}
```

## Example: Delete an experiment template with a specific tag

The following policy grants permission to delete an experiment template with tag Purpose=Test.

JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "fis:DeleteExperimentTemplate"
      ],
      "Resource": "*",
      "Condition": {
        "StringEquals": {
          "aws:ResourceTag/Purpose": "Test"
        }
      }
    }
  ]
}
```

## Example: 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": "*"
}
]
}

```

## Example: Use condition keys for ec2:InjectApiError

The following example policy uses the `ec2:FisTargetArns` condition key to scope target resources. This policy allows the AWS FIS actions `aws:ec2:api-insufficient-instance-capacity-error` and `aws:ec2:asg-insufficient-instance-capacity-error`.

JSON

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": "ec2:InjectApiError",
      "Resource": "*",
      "Condition": {

```

```

    "StringEquals": {
      "ec2:FisActionId": [
        "aws:ec2:api-insufficient-instance-capacity-error",
        "aws:ec2:asg-insufficient-instance-capacity-error"
      ]
    }
  },
  {
    "Effect": "Allow",
    "Action": "ec2:InjectApiError",
    "Resource": "*",
    "Condition": {
      "ForAllValues:ArnLike": {
        "ec2:FisTargetArns": [
          "arn:aws:autoscaling:*:*:autoScalingGroup:uuid:autoScalingGroupName/asg-name"
        ]
      }
    }
  },
  {
    "Effect": "Allow",
    "Action": "autoscaling:DescribeAutoScalingGroups",
    "Resource": "*"
  }
]
}

```

## Example: Use condition keys for aws:s3:bucket-pause-replication

The following example policy uses the `S3:IsReplicationPauseRequest` condition key to allow `PutReplicationConfiguration` and `GetReplicationConfiguration` only when used by AWS FIS in the context of the AWS FIS action `aws:s3:bucket-pause-replication`.

JSON

```

{
  "Version": "2012-10-17",
  "Statement": [
    {

```

```
    "Effect": "Allow",
    "Action": [
      "S3:PauseReplication"
    ],
    "Resource": "arn:aws:s3:::mybucket",
    "Condition": {
      "StringEquals": {
        "s3:DestinationRegion": "region"
      }
    }
  },
  {
    "Effect": "Allow",
    "Action": [
      "S3:PutReplicationConfiguration",
      "S3:GetReplicationConfiguration"
    ],
    "Resource": "arn:aws:s3:::mybucket",
    "Condition": {
      "BoolIfExists": {
        "s3:IsReplicationPauseRequest": "true"
      }
    }
  },
  {
    "Effect": "Allow",
    "Action": [
      "S3:ListBucket"
    ],
    "Resource": "arn:aws:s3:::*"
  },
  {
    "Effect": "Allow",
    "Action": [
      "tag:GetResources"
    ],
    "Resource": "*"
  }
]
```

## Use service-linked roles for AWS Fault Injection Service

AWS Fault Injection Service uses AWS Identity and Access Management (IAM) [service-linked roles](#). A service-linked role is a unique type of IAM role that is linked directly to AWS FIS. Service-linked roles are predefined by AWS FIS and include all of the permissions that the service requires to call other AWS services on your behalf.

A service-linked role makes setting up AWS FIS easier because you don't have to manually add the necessary permissions to manage monitoring and resource selection for experiments. AWS FIS defines the permissions of its service-linked roles, and unless defined otherwise, only AWS FIS can assume its roles. The defined permissions include the trust policy and the permissions policy, and that permissions policy cannot be attached to any other IAM entity.

In addition to the service-linked role, you must also specify an IAM role that grants permission to modify the resources that you specify as targets in an experiment template. For more information, see [IAM roles for AWS FIS experiments](#).

You can delete a service-linked role only after first deleting the related resources. This protects your AWS FIS resources because you can't inadvertently remove permission to access the resources.

### Service-linked role permissions for AWS FIS

AWS FIS uses the service-linked role named **AWSServiceRoleForFIS** to enable it to manage monitoring and resource selection for experiments.

The **AWSServiceRoleForFIS** service-linked role trusts the following services to assume the role:

- `fis.amazonaws.com`

The **AWSServiceRoleForFIS** service-linked role uses the managed policy **AmazonFISServiceRolePolicy**. This policy enables AWS FIS to manage monitoring and resource selection for experiments. For more information, see [AmazonFISServiceRolePolicy](#) in the *AWS Managed Policy Reference*.

You must configure permissions to allow an IAM entity (such as a user, group, or role) to create, edit, or delete a service-linked role. For the **AWSServiceRoleForFIS** service-linked role to be successfully created, the IAM identity that you use AWS FIS with must have the required permissions. To grant the required permissions, attach the following policy to the IAM identity.

## JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": "iam:CreateServiceLinkedRole",
      "Resource": "*",
      "Condition": {
        "StringLike": {
          "iam:AWSServiceName": "fis.amazonaws.com"
        }
      }
    }
  ]
}
```

For more information, see [Service-linked role permissions](#) in the *IAM User Guide*.

### Create a service-linked role for AWS FIS

You don't need to manually create a service-linked role. When you start an AWS FIS experiment in the AWS Management Console, the AWS CLI, or the AWS API, AWS FIS creates the service-linked role for you.

If you delete this service-linked role, and then need to create it again, you can use the same process to recreate the role in your account. When you start an AWS FIS experiment, AWS FIS creates the service-linked role for you again.

### Edit a service-linked role for AWS FIS

AWS FIS does not allow you to edit the **AWSServiceRoleForFIS** service-linked role. After you create a service-linked role, you cannot change the name of the role because various entities might reference the role. However, you can edit the description of the role using IAM. For more information, see [Editing a service-linked role](#) in the *IAM User Guide*.

## Delete a service-linked role for AWS FIS

If you no longer need to use a feature or service that requires a service-linked role, we recommend that you delete that role. That way you don't have an unused entity that is not actively monitored or maintained. However, you must clean up the resources for your service-linked role before you can manually delete it.

### Note

If the AWS FIS service is using the role when you try to clean up the resources, then the cleanup might fail. If that happens, wait for a few minutes and try the operation again.

### To clean up AWS FIS resources used by the `AWSServiceRoleForFIS`

Make sure that none of your experiments are currently running. If necessary, stop your experiments. For more information, see [Stop an experiment](#).

### To manually delete the service-linked role using IAM

Use the IAM console, the AWS CLI, or the AWS API to delete the `AWSServiceRoleForFIS` service-linked role. For more information, see [Deleting a service-linked role](#) in the *IAM User Guide*.

## Supported Regions for AWS FIS service-linked roles

AWS FIS supports using service-linked roles in all of the Regions where the service is available. For more information, see [AWS Fault Injection Service endpoints and quotas](#).

## AWS managed policies for AWS Fault Injection Service

An AWS managed policy is a standalone policy that is created and administered by AWS. AWS managed policies are designed to provide permissions for many common use cases so that you can start assigning permissions to users, groups, and roles.

Keep in mind that AWS managed policies might not grant least-privilege permissions for your specific use cases because they're available for all AWS customers to use. We recommend that you reduce permissions further by defining [customer managed policies](#) that are specific to your use cases.

You cannot change the permissions defined in AWS managed policies. If AWS updates the permissions defined in an AWS managed policy, the update affects all principal identities (users,

groups, and roles) that the policy is attached to. AWS is most likely to update an AWS managed policy when a new AWS service is launched or new API operations become available for existing services.

For more information, see [AWS managed policies](#) in the *IAM User Guide*.

### **AWS managed policy: AmazonFISServiceRolePolicy**

This policy is attached to the service-linked role named **AWSServiceRoleForFIS** to allow AWS FIS to manage monitoring and resource selection for experiments. For more information, see [Use service-linked roles for AWS Fault Injection Service](#).

### **AWS managed policy: AWSFaultInjectionSimulatorEC2Access**

Use this policy in an experiment role to grant AWS FIS permission to run experiments that use the [AWS FIS actions for Amazon EC2](#). For more information, see [the section called "Experiment role"](#).

To view the permissions for this policy, see [AWSFaultInjectionSimulatorEC2Access](#) in the *AWS Managed Policy Reference*.

### **AWS managed policy: AWSFaultInjectionSimulatorECSAccess**

Use this policy in an experiment role to grant AWS FIS permission to run experiments that use the [AWS FIS actions for Amazon ECS](#). For more information, see [the section called "Experiment role"](#).

To view the permissions for this policy, see [AWSFaultInjectionSimulatorECSAccess](#) in the *AWS Managed Policy Reference*.

### **AWS managed policy: AWSFaultInjectionSimulatorEKSAccess**

Use this policy in an experiment role to grant AWS FIS permission to run experiments that use the [AWS FIS actions for Amazon EKS](#). For more information, see [the section called "Experiment role"](#).

To view the permissions for this policy, see [AWSFaultInjectionSimulatorEKSAccess](#) in the *AWS Managed Policy Reference*.

### **AWS managed policy: AWSFaultInjectionSimulatorNetworkAccess**

Use this policy in an experiment role to grant AWS FIS permission to run experiments that use the [AWS FIS networking actions](#). For more information, see [the section called "Experiment role"](#).

To view the permissions for this policy, see [AWSFaultInjectionSimulatorNetworkAccess](#) in the *AWS Managed Policy Reference*.

## AWS managed policy: `AWSFaultInjectionSimulatorRDSAccess`

Use this policy in an experiment role to grant AWS FIS permission to run experiments that use the [AWS FIS actions for Amazon RDS](#). For more information, see [the section called "Experiment role"](#).

To view the permissions for this policy, see [AWSFaultInjectionSimulatorRDSAccess](#) in the *AWS Managed Policy Reference*.

## AWS managed policy: `AWSFaultInjectionSimulatorSSMAccess`

Use this policy in an experiment role to grant AWS FIS permission to run experiments that use the [AWS FIS actions for Systems Manager](#). For more information, see [the section called "Experiment role"](#).

To view the permissions for this policy, see [AWSFaultInjectionSimulatorSSMAccess](#) in the *AWS Managed Policy Reference*.

## AWS FIS updates to AWS managed policies

View details about updates to AWS managed policies for AWS FIS since this service began tracking these changes.

Change	Description	Date
<a href="#">AWSFaultInjectionSimulatorEC2Access</a> – Update to an existing policy	Added permission required for the "AZ: Application Slowdown" and "Cross-AZ: Traffic Slowdown" scenarios. The permissions are: <b>ec2:DescribeSubnets</b>	November 12, 2025
<a href="#">AWSFaultInjectionSimulatorECSAccess</a> – Update to an existing policy	Added permissions required for the "AZ: Application Slowdown" and "Cross-AZ: Traffic Slowdown" scenarios. The permissions are: <b>ecs:DescribeContainerInstances</b> , <b>ec2:DescribeSubnets</b> and <b>ec2:DescribeInstances</b>	November 12, 2025

Change	Description	Date
<a href="#">AWSFaultInjectionSimulatorECSAccess</a> – Update to an existing policy	Added permissions to allow AWS FIS to resolve ECS targets.	January 25, 2024
<a href="#">AWSFaultInjectionSimulatorNetworkAccess</a> – Update to an existing policy	Added permissions to allow AWS FIS to run experiments using the <b>aws:network:route-table-disrupt-cross-region-connectivity</b> and <b>aws:network:transit-gateway-disrupt-cross-region-connectivity</b> actions.	January 25, 2024
<a href="#">AWSFaultInjectionSimulatorEC2Access</a> – Update to an existing policy	Added permissions to allow AWS FIS to resolve EC2 instances.	November 13, 2023
<a href="#">AWSFaultInjectionSimulatorEKSAccess</a> – Update to an existing policy	Added permissions to allow AWS FIS to resolve EKS targets.	November 13, 2023
<a href="#">AWSFaultInjectionSimulatorRDSAccess</a> – Update to an existing policy	Added permissions to allow AWS FIS to resolve RDS targets.	November 13, 2023
<a href="#">AWSFaultInjectionSimulatorEC2Access</a> – Update to an existing policy	Added permissions to allow AWS FIS to run SSM documents on EC2 instances and to terminate EC2 instances.	June 2, 2023
<a href="#">AWSFaultInjectionSimulatorSSMAccess</a> – Update to an existing policy	Added permissions to allow AWS FIS to run SSM documents on EC2 instances.	June 2, 2023
<a href="#">AWSFaultInjectionSimulatorECSAccess</a> – Update to an existing policy	Added permissions to allow AWS FIS to run experiments using the new <b>aws:ecs:task</b> actions.	June 1, 2023

Change	Description	Date
<a href="#">AWSFaultInjectionSimulatorEKSAccess</a> – Update to an existing policy	Added permissions to allow AWS FIS to run experiments using the new <b>aws:eks:pod</b> actions.	June 1, 2023
<a href="#">AWSFaultInjectionSimulatorEC2Access</a> – New policy	Added a policy to allow AWS FIS to run an experiment that uses AWS FIS actions for Amazon EC2.	October 26, 2022
<a href="#">AWSFaultInjectionSimulatorECSAccess</a> – New policy	Added a policy to allow AWS FIS to run an experiment that uses AWS FIS actions for Amazon ECS.	October 26, 2022
<a href="#">AWSFaultInjectionSimulatorEKSAccess</a> – New policy	Added a policy to allow AWS FIS to run an experiment that uses AWS FIS actions for Amazon EKS.	October 26, 2022
<a href="#">AWSFaultInjectionSimulatorNetworkAccess</a> – New policy	Added a policy to allow AWS FIS to run an experiment that uses AWS FIS networking actions.	October 26, 2022
<a href="#">AWSFaultInjectionSimulatorRDSAccess</a> – New policy	Added a policy to allow AWS FIS to run an experiment that uses AWS FIS actions for Amazon RDS.	October 26, 2022
<a href="#">AWSFaultInjectionSimulatorSMSAccess</a> – New policy	Added a policy to allow AWS FIS to run an experiment that uses AWS FIS actions for Systems Manager.	October 26, 2022
<a href="#">AmazonFISServiceRolePolicy</a> – Update to an existing policy	Added permissions to allow AWS FIS to describe subnets.	October 26, 2022
<a href="#">AmazonFISServiceRolePolicy</a> – Update to an existing policy	Added permissions to allow AWS FIS to describe EKS clusters.	July 7, 2022
<a href="#">AmazonFISServiceRolePolicy</a> – Update to an existing policy	Added permissions to allow AWS FIS to list and describe the tasks in your clusters.	February 7, 2022

Change	Description	Date
<a href="#">AmazonFISServiceRolePolicy</a> – Update to an existing policy	Removed the events :ManagedBy condition for the events :DescribeRule action.	January 6, 2022
<a href="#">AmazonFISServiceRolePolicy</a> – Update to an existing policy	Added permissions to allow AWS FIS to retrieve history for the CloudWatch alarms used in stop conditions.	June 30, 2021
AWS FIS started tracking changes	AWS FIS started tracking changes to its AWS managed policies	March 1, 2021

## Infrastructure security in AWS Fault Injection Service

As a managed service, AWS Fault Injection Service 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 FIS 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.

## Access AWS FIS using an interface VPC endpoint (AWS PrivateLink)

You can establish a private connection between your VPC and AWS Fault Injection Service by creating an *interface VPC endpoint*. VPC endpoints are powered by [AWS PrivateLink](#), a technology that enables you to privately access AWS FIS APIs 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 FIS APIs.

Each interface endpoint is represented by one or more [elastic network interfaces](#) in your subnets.

For more information, see [Access AWS services through AWS PrivateLink](#) in the *AWS PrivateLink Guide*.

## Considerations for AWS FIS VPC endpoints

Before you set up an interface VPC endpoint for AWS FIS, review [Access an AWS service using an interface VPC endpoint](#) in the *AWS PrivateLink Guide*.

AWS FIS supports making calls to all of its API actions from your VPC.

## Create an interface VPC endpoint for AWS FIS

You can create a VPC endpoint for the AWS FIS service using either the Amazon VPC console or the AWS Command Line Interface (AWS CLI). For more information, see [Create a VPC endpoint](#) in the *AWS PrivateLink Guide*.

Create a VPC endpoint for AWS FIS using the following service name:

```
com.amazonaws.region.fis.
```

If you enable private DNS for the endpoint, you can make API requests to AWS FIS using its default DNS name for the Region, for example, `fis.us-east-1.amazonaws.com`.

## Create a VPC endpoint policy for AWS FIS

You can attach an endpoint policy to your VPC endpoint that controls access to AWS FIS. The policy specifies the following information:

- The principal that can perform actions.
- The actions that can be performed.
- The resources on which actions can be performed.

For more information, see [Control access to VPC endpoints using endpoint policies](#) in the *AWS PrivateLink Guide*.

### Example: VPC endpoint policy for specific AWS FIS actions

The following VPC endpoint policy grants access to the listed AWS FIS actions on all resources to all principals.

```
{
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "fis:ListExperimentTemplates",
        "fis:StartExperiment",
        "fis:StopExperiment",
        "fis:GetExperiment"
      ],
      "Resource": "*",
      "Principal": "*"
    }
  ]
}
```

### Example: VPC endpoint policy that denies access from a specific AWS account

The following VPC endpoint policy denies the specified AWS account access to all actions and resources, but grants all other AWS accounts access to all actions and resources.

```
{
  "Statement": [
    {
      "Effect": "Allow",
      "Action": "*",
      "Resource": "*",
      "Principal": "*"
    },
    {
      "Effect": "Deny",
      "Action": "*",
      "Resource": "*",
      "Principal": {
        "AWS": [ "123456789012" ]
      }
    }
  ]
}
```

# Tagging your AWS FIS resources

A *tag* is a metadata label that either you or AWS assigns to an AWS resource. Each tag consists of a *key* and a *value*. For tags that you assign, you define the key and the value. For example, you might define the key as `purpose` and the value as `test` for a resource.

Tags help you do the following:

- Identify and organize your AWS resources. Many AWS services support tagging, so you can assign the same tag to resources from different services to indicate that the resources are related.
- Control access to your AWS resources. For more information, see [Controlling Access Using Tags](#) in the *IAM User Guide*.

## Tagging restrictions

The following basic restrictions apply to tags on AWS FIS resources:

- Maximum number of tags that you can assign to a resource: 50
- Maximum key length: 128 Unicode characters
- Maximum value length: 256 Unicode characters
- Valid characters for keys and values: a-z, A-Z, 0-9, space, and the following characters: `_ . : / = + -` and `@`
- Keys and values are case sensitive
- You cannot use `aws :` as a prefix for keys, because it's reserved for AWS use

## Work with tags

The following AWS Fault Injection Service (AWS FIS) resources support tagging:

- Actions
- Experiments
- Experiment templates

You can use the console to work with tags for experiments and experiment templates. For more information, see the following:

- [Tag an experiment](#)
- [Tag experiment templates](#)

You can use the following AWS CLI commands to work with tags for actions, experiments, and experiment templates:

- [tag-resource](#) – Add tags to a resource.
- [untag-resource](#) – Remove tags from a resource.
- [list-tags-for-resource](#) – List the tags for a specific resource.

## Quotas and limitations for AWS Fault Injection Service

Your AWS account has default quotas, formerly referred to as limits, for each AWS service. Unless otherwise noted, each quota is region-specific. You can request increases for the quotas marked as adjustable in the table below.

To view the quotas for AWS FIS in your account, open the [Service Quotas console](#). In the navigation pane, choose **AWS services** and select **AWS Fault Injection Service**. Values up to and including the automatically approved quotas are instantly applied. The automatically approved quotas are outlined in the description column in the table below. If you require quotas exceeding the automatically approved limits, please submit a request. Values above the automatically approved limits are reviewed by customer support and approved where possible.

To request a quota increase, see [Requesting a quota increase](#) in the *Service Quotas User Guide*.

Your AWS account has the following quotas related to AWS FIS.

Name	Default	Adjustable	Description
Action duration in hours	Each supported Region: 12	No	The maximum number of hours allowed to run one action in this account in the current Region.
Action minutes per multi-Region strongly consistent (MRSC) global table for aws:dynamodb:global-table-pause-replication action	Each supported Region: 5,040	<a href="#">Yes</a>	The maximum cumulative action minutes per MRSC global table that aws:dynamodb:global-table-pause-replication can target in a 7-day rolling window.
Actions per experiment template	Each supported Region: 20	No	The maximum number of actions that you can create in an experiment template in this account in the current Region.

Name	Default	Adjustable	Description
Active experiments	Each supported Region: 5	No	The maximum number of active experiments that you can run simultaneously in this account in the current Region.
Completed experiment data retention in days	Each supported Region: 120	No	The maximum number of days allowed for AWS FIS to retain data about completed experiments in this account in the current Region.
Experiment duration in hours	Each supported Region: 12	No	The maximum number of hours allowed to run one experiment in this account in the current Region.
Experiment templates	Each supported Region: 500	No	The maximum number of experiment templates that you can create in this account in the current Region.
Maximum number of Managed Prefix Lists in <code>aws:network:route-table-disrupt-cross-region-connectivity</code>	Each supported Region: 15	No	The maximum number of Managed Prefix Lists that <code>aws:network:route-table-disrupt-cross-region-connectivity</code> will allow, per action.

Name	Default	Adjustable	Description
Maximum number of Route Tables in <code>aws:network:route-table-disrupt-cross-region-connectivity</code>	Each supported Region: 10	No	The maximum number of Route Tables that <code>aws:network:route-table-disrupt-cross-region-connectivity</code> will allow, per action.
Maximum number of routes in <code>aws:network:route-table-disrupt-cross-region-connectivity</code>	Each supported Region: 200	No	The maximum number of routes that <code>aws:network:route-table-disrupt-cross-region-connectivity</code> will allow, per action.
Parallel actions per experiment	Each supported Region: 10	No	The maximum number of actions that you can run in parallel in an experiment in this account in the current Region.
Stop conditions per experiment template	Each supported Region: 5	No	The maximum number of stop conditions that you can add to an experiment template in this account in the current Region.
Target Aurora DSQL clusters for <code>aws:dsql:cluster-connection-failure</code> action.	Each supported Region: 100	<a href="#">Yes</a>	The maximum number of Aurora DSQL clusters that <code>aws:dsql:cluster-connection-failure</code> can target when you identify targets using tags, per experiment.

Name	Default	Adjustable	Description
Target Auto Scaling groups for <code>aws:ec2:asg-insufficient-instance-capacity-error</code>	Each supported Region: 500	<a href="#">Yes</a>	The maximum number of Auto Scaling groups that <code>aws:ec2:asg-insufficient-instance-capacity-error</code> can target when you identify targets using tags, per experiment.
Target Buckets for <code>aws:s3:bucket-pause-replication</code>	Each supported Region: 25	<a href="#">Yes</a>	The maximum number of S3 Buckets that <code>aws:s3:bucket-pause-replication</code> can target when you identify targets using tags, per experiment.
Target Clusters for <code>aws:ecs:drain-container-instances</code>	Each supported Region: 100	<a href="#">Yes</a>	The maximum number of Clusters that <code>aws:ecs:drain-container-instances</code> can target when you identify targets using tags, per experiment.
Target Clusters for <code>aws:rds:failover-db-cluster</code>	Each supported Region: 160	<a href="#">Yes</a>	The maximum number of Clusters that <code>aws:rds:failover-db-cluster</code> can target when you identify targets using tags, per experiment.

Name	Default	Adjustable	Description
Target DBInstances for aws:rds:reboot-db-instances	Each supported Region: 130	<a href="#">Yes</a>	The maximum number of DBInstances that aws:rds:reboot-db-instances can target when you identify targets using tags, per experiment.
Target Instances for aws:ec2:reboot-instances	Each supported Region: 600	<a href="#">Yes</a>	The maximum number of Instances that aws:ec2:reboot-instances can target when you identify targets using tags, per experiment.
Target Instances for aws:ec2:stop-instances	Each supported Region: 400	<a href="#">Yes</a>	The maximum number of Instances that aws:ec2:stop-instances can target when you identify targets using tags, per experiment.
Target Instances for aws:ec2:terminate-instances	Each supported Region: 300	<a href="#">Yes</a>	The maximum number of Instances that aws:ec2:terminate-instances can target when you identify targets using tags, per experiment.

Name	Default	Adjustable	Description
Target Instances for <code>aws:ssm:send-command</code>	Each supported Region: 200	<a href="#">Yes</a>	The maximum number of Instances that <code>aws:ssm:send-command</code> can target when you identify targets using tags, per experiment.
Target Kinesis data streams for <code>aws:kinesis:stream-expired-iterator-exception</code> action.	Each supported Region: 280	<a href="#">Yes</a>	The maximum number of Kinesis data streams that <code>aws:kinesis:stream-expired-iterator-exception</code> can target when you identify targets using tags, per experiment.
Target Kinesis data streams for <code>aws:kinesis:stream-provisioned-throughput-exception</code> action.	Each supported Region: 280	<a href="#">Yes</a>	The maximum number of Kinesis data streams that <code>aws:kinesis:stream-provisioned-throughput-exception</code> can target when you identify targets using tags, per experiment.
Target ManagedResources for <code>aws:arc:start-zonal-autoshift</code> action.	Each supported Region: 200	<a href="#">Yes</a>	The maximum number of managed resources that <code>aws:arc:start-zonal-autoshift</code> can target when you identify targets using tags, per experiment.

Name	Default	Adjustable	Description
Target Nodegroups for aws:eks:terminate-nodegroup-instances	Each supported Region: 100	<a href="#">Yes</a>	The maximum number of Nodegroups that aws:eks:terminate-nodegroup-instances can target when you identify targets using tags, per experiment.
Target Pods for aws:eks:pod-cpu-stress	Each supported Region: 1,000	<a href="#">Yes</a>	The maximum number of Pods that aws:eks:pod-cpu-stress can target when you identify targets using parameters, per experiment.
Target Pods for aws:eks:pod-delete	Each supported Region: 1,000	<a href="#">Yes</a>	The maximum number of Pods that aws:eks:pod-delete can target when you identify targets using parameters, per experiment.
Target Pods for aws:eks:pod-io-stress	Each supported Region: 1,000	<a href="#">Yes</a>	The maximum number of Pods that aws:eks:pod-io-stress can target when you identify targets using parameters, per experiment.

Name	Default	Adjustable	Description
Target Pods for aws:eks:pod-memory-stress	Each supported Region: 1,000	<a href="#">Yes</a>	The maximum number of Pods that aws:eks:pod-memory-stress can target when you identify targets using parameters, per experiment.
Target Pods for aws:eks:pod-network-blackhole-port	Each supported Region: 1,000	<a href="#">Yes</a>	The maximum number of Pods that aws:eks:pod-network-blackhole-port can target when you identify targets using parameters, per experiment.
Target Pods for aws:eks:pod-network-latency	Each supported Region: 1,000	<a href="#">Yes</a>	The maximum number of Pods that aws:eks:pod-network-latency can target when you identify targets using parameters, per experiment.
Target Pods for aws:eks:pod-network-packet-loss	Each supported Region: 1,000	<a href="#">Yes</a>	The maximum number of Pods that aws:eks:pod-network-packet-loss can target when you identify targets using parameters, per experiment.

Name	Default	Adjustable	Description
Target ReplicationGroups for aws:elasticache:interrupt-cluster-az-power - Deprecation Planned	Each supported Region: 5	<a href="#">Yes</a>	The maximum number of ReplicationGroups that aws:elasticache:interrupt-cluster-az-power can target when you identify targets using tags/parameters, per experiment.
Target ReplicationGroups for aws:elasticache:replicationgroup-interrupt-az-power	Each supported Region: 20	<a href="#">Yes</a>	The maximum number of ReplicationGroups that aws:elasticache:replicationgroup-interrupt-az-power can target per experiment. A daily limit applies to targeting ReplicationGroups. For more information, visit: <a href="https://docs.aws.amazon.com/fis/latest/userguide/fis-quotas.html">https://docs.aws.amazon.com/fis/latest/userguide/fis-quotas.html</a> .
Target SpotInstances for aws:ec2:send-spot-instance-interruptions	Each supported Region: 100	<a href="#">Yes</a>	The maximum number of SpotInstances that aws:ec2:send-spot-instance-interruptions can target when you identify targets using tags, per experiment.

Name	Default	Adjustable	Description
Target Subnets for aws:network:disrupt-connectivity	Each supported Region: 100	<a href="#">Yes</a>	The maximum number of Subnets that aws:network:disrupt-connectivity can target when you identify targets using tags, per experiment. Quotas above 5 apply only to parameter scope:all. If you require a higher quota for another scope type, contact customer support.
Target Subnets for aws:network:route-table-disrupt-cross-region-connectivity	Each supported Region: 50	<a href="#">Yes</a>	The maximum number of Subnets that aws:network:route-table-disrupt-cross-region-connectivity can target when you identify targets using tags, per experiment.
Target Tasks for aws:ecs:stop-task	Each supported Region: 500	<a href="#">Yes</a>	The maximum number of Tasks that aws:ecs:stop-task can target when you identify targets using tags, per experiment.
Target Tasks for aws:ecs:task-cpu-stress	Each supported Region: 200	<a href="#">Yes</a>	The maximum number of Tasks that aws:ecs:task-cpu-stress can target when you identify targets using tags/parameters, per experiment.

Name	Default	Adjustable	Description
Target Tasks for aws:ecs:task-io-stress	Each supported Region: 200	<a href="#">Yes</a>	The maximum number of Tasks that aws:ecs:task-io-stress can target when you identify targets using tags/parameters, per experiment.
Target Tasks for aws:ecs:task-kill-process	Each supported Region: 200	<a href="#">Yes</a>	The maximum number of Tasks that aws:ecs:task-kill-process can target when you identify targets using tags/parameters, per experiment.
Target Tasks for aws:ecs:task-network-blackhole-port	Each supported Region: 200	<a href="#">Yes</a>	The maximum number of Tasks that aws:ecs:task-network-blackhole-port can target when you identify targets using tags/parameters, per experiment.
Target Tasks for aws:ecs:task-network-latency	Each supported Region: 200	<a href="#">Yes</a>	The maximum number of Tasks that aws:ecs:task-network-latency can target when you identify targets using tags/parameters, per experiment.

Name	Default	Adjustable	Description
Target Tasks for aws:ecs:task-network-packet-loss	Each supported Region: 200	<a href="#">Yes</a>	The maximum number of Tasks that aws:ecs:task-network-packet-loss can target when you identify targets using tags/parameters, per experiment.
Target TransitGateways for aws:network:transit-gateway-disrupt-cross-region-connectivity	Each supported Region: 50	<a href="#">Yes</a>	The maximum number of Transit Gateways that aws:network:transit-gateway-disrupt-cross-region-connectivity can target when you identify targets using tags, per experiment.
Target Virtual Interfaces for aws:directconnect:virtual-interface	Each supported Region: 5	<a href="#">Yes</a>	The maximum number of Virtual Interfaces that aws:directconnect:virtual-interface can target when you identify targets using tags, per experiment.
Target Volumes for aws:ebs:pause-volume-io	Each supported Region: 160	<a href="#">Yes</a>	The maximum number of Volumes that aws:ebs:pause-volume-io can target when you identify targets using tags, per experiment.

Name	Default	Adjustable	Description
Target account configurations per experiment template	Each supported Region: 40	<a href="#">Yes</a>	The maximum number of target account configurations that you can create for an experiment template in this account in the current Region.
Target functions for aws:lambda:invocation-add-delay action.	Each supported Region: 140	<a href="#">Yes</a>	The maximum number of Lambda functions that aws:lambda:invocation-add-delay can target when you identify targets using tags, per experiment.
Target functions for aws:lambda:invocation-error action.	Each supported Region: 140	<a href="#">Yes</a>	The maximum number of Lambda functions that aws:lambda:invocation-error can target when you identify targets using tags, per experiment.
Target functions for aws:lambda:invocation-http-integration-response action.	Each supported Region: 140	<a href="#">Yes</a>	The maximum number of Lambda functions that aws:lambda:invocation-http-integration-response can target when you identify targets using tags, per experiment.

Name	Default	Adjustable	Description
Target multi-Region clusters for aws:memorydb:multi-region-cluster-pause-replication action.	Each supported Region: 50	<a href="#">Yes</a>	The maximum number of MemoryDB multi-Region clusters that aws:memorydb:multi-region-cluster-pause-replication can target when you identify targets using tags, per experiment. If you require a higher quota, contact customer support.
Target tables for aws:dynamodb:global-table-pause-replication action	Each supported Region: 60	<a href="#">Yes</a>	The maximum number of global tables that aws:dynamodb:global-table-pause-replication can target, per experiment.

Your usage of AWS FIS is subject to the following additional limitations:

Name	Limitation
Daily limit on targets ReplicationGroups for aws:elasticache:replicationgroup-interrupt-az-power	The limit is 20 ReplicationGroups targeted per account per region per day. You can request an increase by creating a support case in the <a href="#">AWS Support Center Console</a> .

## Document history

The following table describes important documentation updates in the *AWS Fault Injection Service User Guide*.

Change	Description	Date
<a href="#">New action for AWS Direct Connect</a>	You can use the <b>aws:directconnect:virtual-interface</b> action to test the resilience of the AWS Direct Connect connection by temporarily disrupting the Border Gateway Protocol sessions between the on-premises networks and peers associated with target Virtual Interfaces.	December 18, 2025
<a href="#">New scenarios</a>	You can now use the new "AZ: Application Slowdown" and "Cross-AZ: Traffic Slowdown" scenarios.	November 12, 2025
<a href="#">New parameter for SSM documents</a>	The <b>AWSFIS-Run-Network-Latency-Sources</b> and <b>AWSFIS-Run-Network-Packet-Loss-Sources</b> SSM documents now support the <b>flowsPercent</b> parameter.	November 12, 2025
<a href="#">New parameter for ECS and EKS actions</a>	The <b>aws:ecs:task-network-latency</b> , <b>aws:ecs:task-network-packet-loss</b> , <b>aws:eks:pod-network-latency</b> and <b>aws:eks:pod-network-packet-loss</b> actions now	November 12, 2025

---

	support the <b>flowsPercent</b> parameter.	
<a href="#">AWS managed policy updates</a>	Updated managed policy <a href="#">AWSFaultInjectionSimulatorEC2Access</a> : Added <b>ec2:DescribeSubnets</b> permission.	November 12, 2025
<a href="#">AWS managed policy updates</a>	Updated managed policy <a href="#">AWSFaultInjectionSimulatorECSAccess</a> : Added <b>ecs:DescribeContainerInstances</b> , <b>ec2:DescribeSubnets</b> and <b>ec2:DescribeInstances</b> permissions.	November 12, 2025
<a href="#">New actions for Amazon Kinesis Data Streams</a>	You can use the <b>aws:kinesis:stream</b> actions.	October 15, 2025
<a href="#">New experiment for Amazon EBS</a>	You can use the <b>aws:eks:volume-io-latency</b> action to simulate elevated I/O latency on your Amazon EBS volumes.	September 16, 2025
<a href="#">New action scope AWS FIS</a>	You can use the <b>aws:network:disrupt-connectivity</b> action to disrupt connectivity to your S3 Express One Zone directory buckets. This scope is also now included in the AZ Availability: Power Interruption scenario.	August 18, 2025

---

<a href="#">MemoryDB support in AWS FIS</a>	You can use AWS FIS to test how your application with Amazon MemoryDB Multi-Region clusters responds to data replication being paused during a cross-Region network disruption.	July 15, 2025
<a href="#">ARC support in AWS FIS</a>	You can use AWS FIS to test how ARC zonal autoshift automatically recovers your application during an AZ power interruption.	March 26, 2025
<a href="#">New experiment report configuration</a>	You can now enable AWS FIS to generate reports for experiments that summarize experiment actions and responses from CloudWatch dashboards.	November 12, 2024
<a href="#">New Lambda actions</a>	You can now use <b>aws:lambda:action</b> actions to inject faults into invocations of your Lambda functions.	October 31, 2024
<a href="#">New safety lever feature</a>	AWS FIS now supports safety levers which allow you to quickly stop all running experiments and prevent new experiments from starting.	September 3, 2024
<a href="#">New Troubleshooting chapter</a>	AWS FIS added a troubleshooting guide that includes error codes and context for failed experiments.	August 13, 2024

---

<a href="#">New action</a>	You can now use the <code>aws:dynamodb:global-table-pause-replication</code> action to pause data replication between the target global table and its replica tables. The <code>aws:dynamodb:encrypted-global-table-pause-replication</code> action will no longer be supported.	April 24, 2024
<a href="#">New actions mode experiment option</a>	You can set actions mode to <code>skip-all</code> to generate a target preview before running an experiment.	March 13, 2024
<a href="#">AWS managed policy updates</a>	AWS FIS updated existing managed policies.	January 25, 2024
<a href="#">New scenarios and actions</a>	You can now use AWS FIS scenarios <code>Cross-Region:Connectivity</code> and <code>AZ Availability:Power Interruption</code> .	November 30, 2023
<a href="#">New action</a>	You can now use the <b><code>aws:ec2:asg-insufficient-in-stance-capacity-error</code></b> action.	November 30, 2023
<a href="#">New action</a>	You can now use the <b><code>aws:ec2:api-insufficient-in-stance-capacity-error</code></b> action.	November 30, 2023
<a href="#">New action</a>	You can now use the <b><code>aws:network:route-table-disrupt-cross-region-connectivity</code></b> action.	November 30, 2023

---

<a href="#">New action</a>	You can now use the <b>aws:network:transit-gateway-disrupt-cross-region-connectivity</b> action.	November 30, 2023
<a href="#">New action</a>	You can now use the <b>aws:dynamodb:encrypted-global-table-pause-replication</b> action.	November 30, 2023
<a href="#">New action</a>	You can now use the <b>aws:s3:bucket-pause-replication</b> action.	November 30, 2023
<a href="#">New action</a>	You can now use the <b>aws:elasticache:interrupt-cluster-az-power</b> action.	November 30, 2023
<a href="#">New experiment options</a>	You can now use AWS FIS experiment options for account targeting and empty target resolution.	November 27, 2023
<a href="#">Name change of AWS FIS</a>	Updated service name to AWS Fault Injection Service.	November 15, 2023
<a href="#">AWS managed policy updates</a>	AWS FIS updated existing managed policies.	November 13, 2023
<a href="#">New scenario library</a>	You can now use the AWS FIS scenario library feature.	November 7, 2023
<a href="#">New experiment scheduler</a>	You can now use the AWS FIS experiment scheduler feature.	November 7, 2023
<a href="#">AWS managed policy updates</a>	AWS FIS updated existing managed policies.	June 2, 2023

---

<a href="#">New actions</a>	You can use the new <b>aws:ecs:task</b> and <b>aws:eks:pod</b> actions.	June 1, 2023
<a href="#">AWS managed policy updates</a>	AWS FIS updated existing managed policies.	June 1, 2023
<a href="#">New pre-configured SSM document</a>	You can use the following pre-configured SSM document: AWSFIS-Run-Disk-Fill.	April 28, 2023
<a href="#">New action</a>	You can use the <b>aws:ebs:pause-volume-io</b> action to pause I/O between the target volumes and the instances they are attached to.	January 27, 2023
<a href="#">New action</a>	You can use the <b>aws:network:disrupt-connectivity</b> action to deny specific types of traffic to the target subnets.	October 26, 2022
<a href="#">New action</a>	You can use the <b>aws:eks:inject-kubernetes-custom-resource</b> action to run a ChaosMesh or Litmus experiment on a single target cluster.	July 7, 2022
<a href="#">Experiment logging</a>	You can configure your experiment templates to send experiment activity logs to CloudWatch Logs or to an S3 bucket.	February 28, 2022

---

<a href="#">New notifications</a>	When the state of an experiment changes, AWS FIS emits a notification. These notifications are made available as events through Amazon EventBridge.	February 24, 2022
<a href="#">New action</a>	You can use the <b>aws:ecs:stop-task</b> action to stop the specified task.	February 9, 2022
<a href="#">New action</a>	You can use the <b>aws:cloudwatch:assert-alarm-state</b> action to verify that the specified alarms are in one of the specified alarm states.	November 5, 2021
<a href="#">New pre-configured SSM documents</a>	You can use the following pre-configured SSM documents: AWSFIS-Run-IO-Stress, AWSFIS-Run-Network-Blackhold-Port, AWSFIS-Run-Network-Latency-Sources, AWSFIS-Run-Network-Packet-Loss, and AWSFIS-Run-Network-Packet-Loss-Sources.	November 4, 2021
<a href="#">New action</a>	You can use the <b>aws:ec2:stop-spot-instances</b> action to send a Spot Instance interruption notice to target Spot Instances and then interrupt the target Spot Instances.	October 20, 2021

[New action](#)

You can use the **aws:ssm:s  
tart-automation-execution**  
action to initiate the  
execution of an Automation  
runbook.

September 17, 2021

[Initial release](#)

The initial release of the AWS  
Fault Injection Service User  
Guide.

March 15, 2021