

Implementation Guide

Verifiable Controls Evidence Store



Verifiable Controls Evidence Store: Implementation Guide

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

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

Table of Contents

Solution overview	1
Cost	2
Architecture overview	4
Solution components	6
AWS Lambda	6
Logging and monitoring	6
CloudWatch dashboard	7
Table 1: Dashboard widgets	7
Security	9
Network configuration	9
Data protection	9
AWS Identity & Access Management (IAM)	9
Amazon S3 bucket configuration and policy	10
Amazon OpenSearch Service access policy	10
Data integrity in Amazon OpenSearch Service	10
Design considerations	11
Quotas	11
Data backup and restore	11
Regional deployments	12
Supported deployment Regions	12
Domain model	13
Deployment	14
Prerequisites	14
AWS credentials	14
Customizable configuration	14
Launch the stack	15
Post-deployment S3 evidence collector configuration	16
Troubleshooting common deployment issues	17
Using the solution's web application	18
Step 1: Sign in to the web interface	18
Step 2: Onboard an evidence provider	18
Step 3: Store evidence	21
Step 4: Search, retrieve, and inspect evidences	22
AWS Security Hub and AWS Config Evidence Collector	25

Using the solution via APIs	27
Create an IAM policy to access the API	27
Grant access to the entire API	27
Create the IAM policy	28
Create Evidence Provider API	28
Create Evidence API	28
Get Evidence API	29
Troubleshooting	30
Contact Support	30
Create case	30
How can we help?	30
Additional information	30
Help us resolve your case faster	31
Solve now or contact us	31
Additional resources	32
Uninstall the solution	33
Using the AWS Management Console	33
Using AWS Command Line Interface	34
Deleting the Amazon S3 buckets	34
Deleting DynamoDB tables	34
Deleting Amazon OpenSearch Service domain	35
Deleting Amazon QLDB ledger	35
Source code	36
Contributors	37
Revisions	38
Notices	39

Deploy a mechanism to store cloud security control findings and results

Publication date: *June 2022* ([last update](#): May 2023)

The **Verifiable Controls Evidence Store** solution provides a mechanism to centrally store the findings and results of cloud security controls governing AWS workloads, in the form of enduring evidence records that are safeguarded against tampering. The solution is useful where such controls, and other governance systems or processes, issue evidence for immutable storage, which can later be utilized in compliance evaluation, deployment decisions, or audit processes. For example, evidence of the findings from preventative controls run from an application deployment pipeline can be stored and retrieved, in near real-time, as part of a subsequent pipeline stage, to determine if the software release meets compliance requirements, before allowing deployment.

An *evidence record* is a system-generated (or human-generated) digital record of a historical fact, related to one or more target entities, and is issued by an evidence provider.

The solution automatically generates unique evidence record IDs associated with the evidences provided by evidence providers.

This implementation guide describes architectural considerations and configuration steps for deploying Verifiable Controls Evidence Store in the Amazon Web Services (AWS) Cloud.

The guide is intended for IT architects and developers who have practical experience architecting in the AWS Cloud and are looking to deploy the aforementioned capabilities in their AWS environment. The guide also covers the user interface and APIs that allow users to interact with the solution, such as members of application teams, control teams, and risk, assurance, and internal audit functions.

Note

AWS does not provide compliance or regulatory advice. You should independently evaluate the suitability of Verifiable Controls Evidence Store for your use case, including for the purposes of meeting any audit, compliance, and regulatory requirements that you may have.

Cost

You are responsible for the cost of the AWS services used while running this solution. As of the latest revision, the cost for running this solution with the default, production-ready settings of `r5.large.search` Amazon OpenSearch Service node instances in the Asia Pacific (Sydney) Region is approximately **\$882.00 per month**.

This includes storage for 10,000 10 KB evidences, 10,000 create evidence requests, and 20,000 query evidence requests. The cost varies depending on the solution settings, API invocation patterns, and data storage size.

AWS service	Cost
Amazon OpenSearch Service (<code>r5.large.search</code> instance type)	\$800.00
Amazon Kinesis	\$62.50
AWS Key Management Service	\$18.35
Amazon S3	\$0.56
Amazon QLDB	\$0.54
Amazon Simple Queue Service	\$0.10
Amazon Cognito	\$0
Amazon CloudFront	\$0
Amazon API Gateway	\$0
AWS Lambda	\$0
Total:	\$882.05

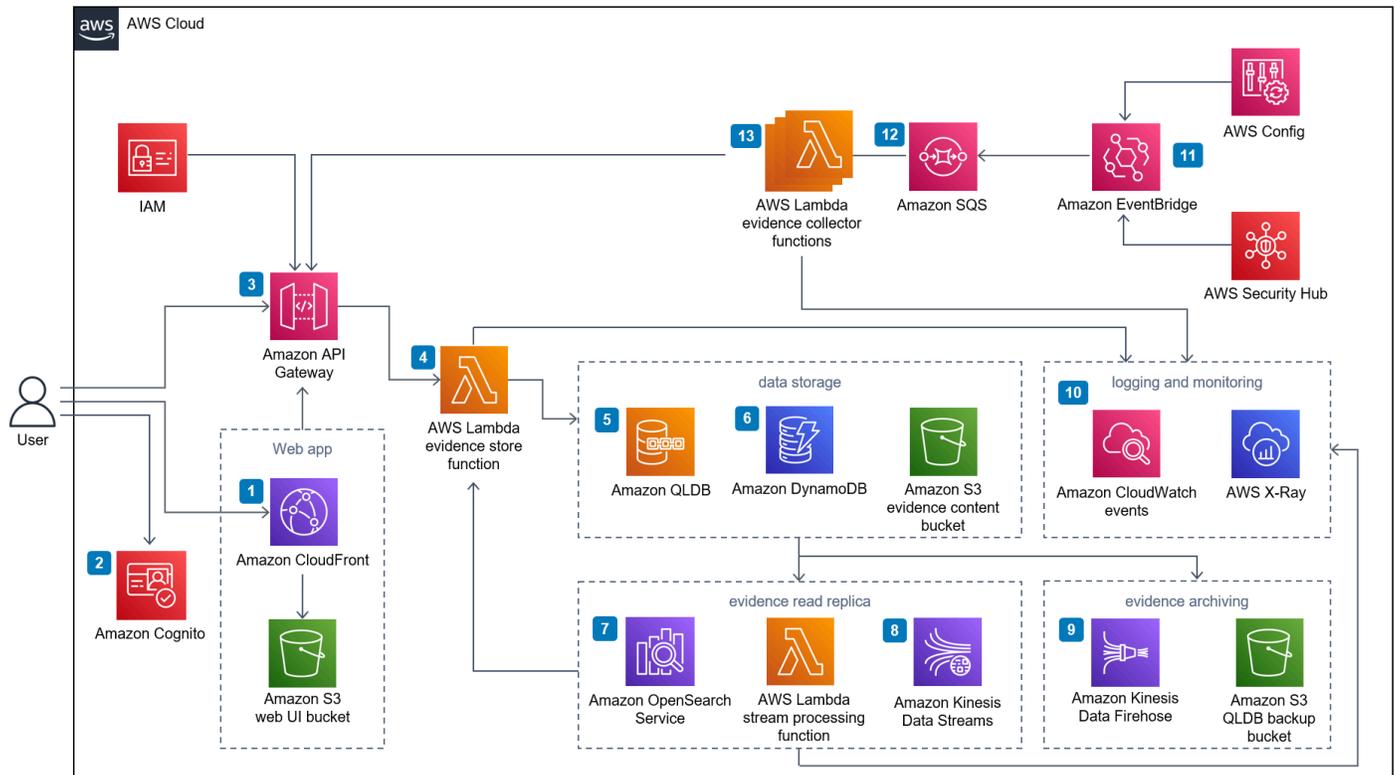
We recommend creating a [budget](#) through [AWS Cost Explorer](#) to help manage costs. Prices are subject to change. For full details, see the pricing webpage for each AWS service used in this solution.

Note

Depending on your performance requirements, a smaller instance size might be suitable. Total cost for the solution using OpenSearch Service `t3.small.search` instance type is **\$312.05** and total cost for OpenSearch Service `m3.medium.search` instance type is **\$482.05**. For more information about sizing, refer to [Sizing Amazon OpenSearch Service domains](#) in the *Amazon OpenSearch Service Developer Guide*.

Architecture overview

Deploying this solution with the default parameters builds the following environment in the target AWS account.



Verifiable Controls Evidence Store solution architecture diagram

Verifiable Controls Evidence Store follows a micro-service architecture, where the presentation service (UI) is a React single-page application. The business logic is powered by the combination of [Amazon API Gateway](#) and [AWS Lambda](#).

The solution provides a RESTful interface with CRUD (create, read, update, delete) APIs for managing evidence providers, as well as for storage and retrieval of evidence data.

At its core, the solution leverages [Amazon Quantum Ledger Database](#) (Amazon QLDB) to store evidence data. QLDB's inherent immutability, transparency, scalability, and cryptographic security help ensure data integrity, validity, full traceability, and auditability. To provide advanced and high-performance query capabilities, the solution also utilizes Amazon OpenSearch Service, which requires a dedicated instance.

By default, the solution builds the following infrastructure:

1. An [Amazon CloudFront](#) distribution to serve the optional UI. CloudFront delivers low latency, high performance, and secure static web hosting. An [Amazon Simple Storage Service](#) (Amazon S3) web UI bucket hosts the static web application artifacts.
2. An [Amazon Cognito](#) user pool to provide customers a quick and convenient authentication mechanism to explore the solution's functionalities without extensive configuration.
3. API Gateway to expose a set of RESTful APIs. API Gateway processes HTTP requests issued by the Evidence Store consumers. It orchestrates the authentication and authorization workflows by validating the request's credentials (signature and API key) against [AWS Identity and Access Management](#) (IAM) and its API usage plan.
4. An evidence store Lambda function to process the validated requests from API Gateway. This Lambda function encapsulates the solution's business logic, receiving rest requests from the user via API Gateway, validating them and storing, and retrieving data to and from the various databases.
5. Amazon QLDB to track and store evidence records. QLDB ensures evidence records' immutability and cryptographically verifiable nature. Evidence records are stored in Amazon S3.
6. [Amazon DynamoDB](#) to store evidence providers and their respective evidence content schemas. The request processing Lambda function relies on this data to validate evidence content before committing to its QLDB ledger.
7. A stream processing Lambda function to replicate evidence records to [Amazon OpenSearch Service](#), which offers advanced query capabilities (full text search) across the entire evidence record data structure.
8. [Amazon Kinesis Data Streams](#) to replicate records to Amazon OpenSearch Service to offer consumers a better query experience. [Amazon Kinesis](#) provides the channels for the solution to replicate and archive evidence records in near real-time.
9. [Amazon Data Firehose](#) to archive evidence records to an Amazon S3 bucket.
- 10 [AWS CloudWatch](#) and [AWS X-Ray](#) for logging and monitoring.
- 11 [AWS Config](#) and [AWS Security Hub](#) to publish findings to Amazon EventBridge.
- 12 [Amazon Simple Queue Service](#) (SQS) to provide rate limiting capabilities to AWS Config and the Security Hub Evidence Collector.
- 13 Evidence collector Lambda functions to invoke the Create Evidence API to record the finding. These include the Security Hub evidence collector and the S3 evidence collector.

Solution components

AWS Lambda

The solution deploys the following Lambda functions:

- **Evidence Store request processing** - This function performs all business logic around data persistence and retrieval for the Verifiable Controls Evidence Store. It receives requests from API Gateway, reads or writes data into its various databases, and returns the requested evidence records and associated artifacts to the consumers.
- **Stream processing** - This function replicates evidence data stored in QLDB to an Amazon OpenSearch Service read replica. As data records are persisted, QLDB streams them out to this Lambda function via Kinesis Data Streams. The Lambda function performs a number of data validation steps to ensure data integrity and authenticity prior to indexing it in the Amazon OpenSearch Service cluster.
- **Security Hub evidence collector** - This function provides out-of-the-box integration with AWS monitoring and control services including Security Hub and Config allowing their findings to be recorded as evidence.
- **S3 evidence collector** (Optional) - This function provides customers the ability to monitor a set of S3 buckets and create evidences when new objects are put into any of these buckets. The connector provides assurance around the integrity and traceability of these new objects.

Logging and monitoring

The solution outputs logs and instrumentation data to an AWS CloudWatch log group and AWS X-Ray respectively. The default log threshold is set to `info` and can be overridden by modifying the Lambda Functions' `LOG_LEVEL` environment variable. Possible log threshold values include:

- `error`
- `warn`
- `info`
- `verbose`
- `debug`
- `silly`

To facilitate filtering and querying log entries through CloudWatch Logs Insights, each log entry produced by the Verifiable Controls Evidence Store Lambda functions includes the following metadata:

- HTTP method
- HTTP resource
- API Gateway Request Id
- Lambda Request Id
- X-Ray TraceId

For addressing common errors, refer to the Troubleshooting guide in the solution's [README.md](#) file.

CloudWatch dashboard

The Verifiable Controls Evidence Store solution includes an Amazon CloudWatch dashboard that is configured to provide an overview of the health and operational status of all components.

Table 1: Dashboard widgets

Component	Value type
Synthetics Canary status	Pass/Fail
Invocation counts for each API	Count
Latency for each API	Milliseconds
Error counts for each API	Count
Lambda request duration	Milliseconds
Lambda success rate	Percentage
DynamoDB capacity	Count
DynamoDB latency	Milliseconds

Component	Value type
DynamoDB Throttled requests count	Count
QLDB latency	Count
Amazon OpenSearch Service domain status	Cluster status
Amazon OpenSearch Service counts for 2xx 3xx 4xx and 5xx responses	Count
Data replication delay between QLDB and Amazon OpenSearch Service	Milliseconds
Total number of evidence records indexed in Amazon OpenSearch Service	Count

Security

When you build systems on AWS infrastructure, security responsibilities are shared between you and AWS. This [shared model](#) reduces your operational burden because AWS operates, manages and controls the components from the host operating system, the virtualization layer, and the physical security of the facilities in which the services operate. For more information about AWS security, visit [AWS Cloud Security](#).

Network configuration

The Verifiable Controls Evidence Store solution is deployed in an [Amazon Virtual Private Cloud](#) (Amazon VPC), with its Lambda functions deployed within a private subnet and all databases in an isolated subnet. Traffic in and out of the isolated subnet is controlled by security groups applied to the Lambda functions. By default, the security group rules only allow inbound traffic from the Lambda functions' private subnet, to prevent unauthorized access to the data storage layer.

Data protection

All data committed to the solution is encrypted at rest using [AWS Key Management Service](#) (AWS KMS) customer managed keys. This includes data stored in the following services:

- Amazon S3
- Amazon QLDB
- Kinesis Data Streams
- DynamoDB
- Amazon OpenSearch Service
- Amazon SQS

Communications between the solution's different components are over HTTPS to ensure data encryption in transit.

AWS Identity & Access Management (IAM)

The solution leverages IAM to secure access to its APIs. It grants permissions to authenticated users to invoke APIs based on their specific persona.

Amazon S3 bucket configuration and policy

By default, all S3 buckets for Verifiable Controls Evidence Store have the following configuration:

- Block all public access
- Versioning enabled
- Access log enabled
- Encryption at rest by an AWS KMS customer managed key.

Additionally, they are also configured with a default resource policy that deny all non-HTTPS requests to ensure data in transit encryption.

Amazon OpenSearch Service access policy

In addition to the network control described above, the solution's Amazon OpenSearch Service domain is further secured by an identity-based access policy, which only grants access to the domain to the Lambda functions' run roles.

Data integrity in Amazon OpenSearch Service

To ensure the consistency and validity (including the integrity, authenticity, and immutability) of data replicated in Amazon OpenSearch Service, the solution comes with an API that validates a given evidence. The validation process involves comparing the evidence record stored in Amazon OpenSearch Service against its respective original in Amazon QLDB, in addition to cryptographically verifying the QLDB record's hash.

Design considerations

With the exception of its use of the Amazon OpenSearch Service, this solution is a serverless architecture. To enhance data query and retrieval performance, the solution replicates evidence records to an Amazon OpenSearch Service domain in near-real time, as they are committed to the QLDB ledger. This is achieved by leveraging QLDB's streaming capability with a combination of Kinesis Data Streams and Lambda.

The introduction of an Amazon SQS queue in front of the AWS Config and Security Hub Evidence Collector helps protect the Verifiable Controls Evidence Store and the services it depends on from the potentially large number of generated findings. In some scenarios, Security Hub and AWS Config can produce tens of thousands of findings in a short period of time, which, in the absence of an SQS queue, would overwhelm the store and severely impact its performance.

Amazon Cognito helps simplify the initial setup, allowing customers to quickly deploy and inspect the capabilities offered by the solution. When operating in a production environment, Amazon Cognito should be replaced by a customer's identity provider(s) to handle user authentication. For further instructions on how to set up identity federation with a SAML2 or OpenID Connect (OIDC) capable identity provide, refer to the solution's [README.md](#) file.

Quotas

The Verifiable Controls Evidence Store solution has the following limits:

- 500 evidence providers per AWS Region
- 10,000 requests per seconds per Region

Refer to the respective services' FAQ for detailed information on quota and limitation for each of the services used in the solution. Some service limitations can be increased by contacting [AWS Support](#), as needed.

Data backup and restore

All DynamoDB tables have Point-In-Time recovery activated by default. For QLDB, the entire ledger is continuously streamed to an S3 bucket and can be replayed onto a new ledger. Amazon OpenSearch Service data can be restored by restarting the QLDB stream, this initiates the data replication process from the QLDB source.

Regional deployments

This solution uses the Amazon QLDB service, which is not currently available in all AWS Regions. You must launch this solution in an AWS Region where Amazon QLDB is available. For the most current availability of AWS services by Region, refer to the [AWS Regional Services List](#).

Supported deployment Regions

As of June 2022, Verifiable Controls Evidence Store can be deployed in the following AWS Regions in accordance with the regional availability of its constituent services:

Region ID	Region name
us-east-2	US East (Ohio)
us-east-1	US East (N. Virginia)
us-west-2	US West (Oregon)
af-south-1	Africa (Cape Town)
ap-northeast-2	Asia Pacific (Seoul)
ap-southeast-2	Asia Pacific (Singapore)
ap-southeast-2	Asia Pacific (Sydney)
ap-northeast-1	Asia Pacific (Tokyo)
ca-central-1	Canada (Central)
eu-central-1	Europe (Frankfurt)
eu-west-1	Europe (Ireland)
eu-west-2	Europe (London)
me-south-1	Middle East (Bahrain)
us-gov-west-1	AWS GovCloud (US-West)

Domain model

This section introduces the key domain concepts used in the solution.

- **Evidence:** An *evidence* is an individual system-generated (or human-generated) digital record of a historical control outcome related to one or more target entities, such as an application release or a software deployment environment. The evidence is used to assess whether a control objective is met effectively, in addition to other purposes.
- **Evidence provider:** The certified and trusted originator of an evidence document is known as an *evidence provider*. The solution allows evidence originating from a variety of software services, systems, tools or processes, to be stored. An evidence provider must be onboarded to the solution.
- **Evidence schema:** The solution supports the storage of evidence with varying output and formats via *evidence schemas* that implicitly correspond to a custom evidence type. Evidence schemas define the data model for an evidence type, which allows consumers to interpret the evidence information as intended by the evidence provider. Evidence schemas effectively extend the common, unified data model of the solution to support evidence types for a variety of use cases.

Deployment

Before you launch the solution, review the cost, architecture, network security, and other considerations discussed in this guide. Follow the step-by-step instructions in this section to configure and deploy the solution into your account.

Time to deploy: Approximately 70 minutes.

Prerequisites

To deploy the solution, you must have the following:

- The latest version of the [AWS CLI](#), installed and configured.
- The latest version of the [AWS CDK](#).
- A CDK bootstrapped AWS account. For details, refer to [Bootstrapping](#) in the *CDK Developer Guide*.
- Node.js version 12 or newer.

AWS credentials

Verify your credentials to access the target AWS account are properly configured. These could be in environment variables or in the configuration files. If not, refer to [Configuration and credential file settings](#) in the *AWS CLI User Guide* to configure these first.

Customizable configuration

You can configure the Verifiable Controls Evidence Store solution to suit different customer requirements in terms of network environments, performance, or authentication methods.

You can also specify your own KMS encryption keys, data retention policy, Amazon OpenSearch Service nodes instance type, and proxy server if required by your infrastructure setup. Use the solution's `Default.json` configuration file found in the `configuration` directory as a sample for your reference. This is the same file used during default stack deployment. You can also create your own configuration files with different names and refer to them when deploying the stack.

By default, the solution's web front end uses Amazon Cognito to authenticate and authorize users. A unique domain prefix is required to provision a new Amazon Cognito user pool. Specify the

domain prefix under the `AGSSharedInfra.identityProvider.domainPrefix` section in the `Default.json` configuration file. For a code sample, refer to the solution's [README.md](#) file.

Launch the stack

To deploy this solution with an installation script, complete the following steps:

1. In your terminal, run the following command to clone the solution source code from the GitHub location:

```
git clone https://github.com/aws-solutions/verifiable-controls-evidence-store
```

2. Navigate to the source code folder created in step 1:

```
cd verifiable-controls-evidence-store/source
```

3. **(Optional)** To configure the S3 evidence collector:

- Navigate to the configuration folder:

```
cd configuration
```

- Under **AGSEvidenceStore.sourceBuckets**, specify the source S3 buckets in the **Default.json** file. These S3 buckets will be monitored by the S3 collector. For example:

```
"sourceBuckets": [  
  {  
    "account": "123456789",  
    "bucketArn": "arn:aws:s3:::bucket1"  
  },  
  {  
    "account": "123456789",  
    "bucketArn": "arn:aws:s3:::bucket2"  
  }  
]
```

- Navigate to the source directory, and set the deployment environment to **Sydney** (ap-southeast-2):

```
cd..  
export AWS_REGION=ap-southeast-2
```

4. To deploy the solution with the default configuration, run:

```
node install.js
```

If you would like to deploy with a customized configuration file, run:

```
node install.js <Configuration name>
```

Note

You are responsible for the cost of the AWS services used while running this solution. For more details, visit the [Cost](#) section in this guide, and refer to the pricing webpage for each AWS service used in this solution.

Post-deployment S3 evidence collector configuration

As the source S3 buckets are customer defined, you need to configure them to publish events to the S3 collector, and give it *read* permission. For more information on how to update bucket policies, refer to [Bucket owner granting cross-account bucket permissions](#) in the *Amazon S3 User Guide*.

The Lambda arn and Lambda role arn are presented as cdk output. For example,

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "Statement1",
      "Effect": "Allow",
      "Principal": {
        "AWS": "s3-connector-lambda-role-arn"
      },
      "Action": [
        "s3:ListBucket",
        "s3:GetObject*",
        "s3:PutObject*"
      ],
      "Resource": [
```

```
        "arn:aws:s3:::bucket1",  
        "arn:aws:s3:::bucket1/*"  
    ]  
  }  
]  
}
```

Once deployed, the S3 connector will continuously monitor the provided S3 buckets and will create evidences when a new object is uploaded.

To view the evidence records, log in to the Evidence Store UI using the CloudFront link (presented as cdk output), and use the credentials (the email and password provided and generated as part of the deployment process).

Troubleshooting common deployment issues

Error message	Solution
Failed to retrieve AGS Shared Infra version from target environment. Error: ConfigError: Missing region in config Failed to read Shared Infra version in the target account. Cannot proceed.	Ensure the environment variable <code>AWS_REGION</code> is set with your preferred deployment Region, for example, <code>ap-southeast-2</code> .
Current credentials could not be used to assume 'arn:aws:iam::[ACCOUNT_ID]:role/cdk-hnb659fds-file-publishing-role-[ACCOUNT_ID]-us-west-2', but are for the right account. Proceeding anyway.	The deployment account might not have been properly bootstrapped with CDK new style bootstrapping. Ensure the environment variable <code>CDK_NEW_BOOTSTRAP</code> is set to <code>1</code> prior to bootstrapping your account.

Using the solution's web application

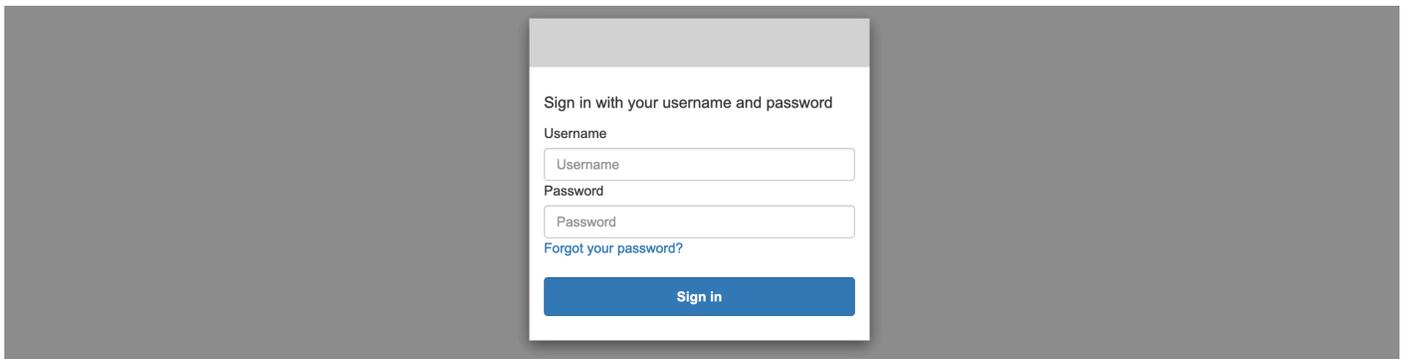
This section provides instructions on how to sign in to use the web application, onboard an evidence provider, store evidences, and how to search, retrieve and inspect evidences.

Step 1: Sign in to the web interface

After the solution CloudFormation stack has been deployed and launched, you can sign in to the web interface.

1. Sign in to the [AWS CloudFormation console](#) and select the **AGSSharedInfra-WebClient** stack.
2. Choose the **Outputs** tab.
3. Under the **Key** column, locate the **webClientOutput** key, and select the link.
4. Choose **Sign in**.

The Verifiable Controls Evidence Store redirects to the specific login mechanism that was configured as part of the solution's deployment steps under [Customizable Configuration](#). By default, the solution's web front end uses Amazon Cognito to authenticate and authorize users.



Front end Amazon Cognito login screen.

5. Follow the prompts to sign in.

Step 2: Onboard an evidence provider

Evidence providers must first be registered in order to create evidence records.

Note

The default user (admin) automatically includes the permissions to onboard evidence providers. If you were onboarded, either via Amazon Cognito or the customer's own IDP, verify that you have the custom:AGSRo1es attribute name permissions level required to complete this task.

1. [Sign in to the web interface](#), then from the left navigation menu, select **Evidence**, and then **Evidence Providers**.
2. Choose **Create an Evidence Provider**.

The screenshot displays the 'Evidence Providers' page in the Verifiable Evidence Store. The page includes a search bar, a table of providers, and a 'Create an Evidence Provider' button. The table contains the following data:

Provider name	Provider Description	Status	Onboard Date Time	Number of Evidence Types
canary-authority		Active	Mar 14, 2022, 5:21:00 AM	1

Create an Evidence Provider

3. On the **Onboard a new Evidence Provider** page:
 - a. Enter the evidence provider name.
 - b. (Optional) Enter an evidence provider ID.
 If not supplied, the system generates an ID.
 - c. Provide a description.
 - d. Choose **Add Evidence Schema** and specify an evidence schema with content that adheres to a valid JSON schema.

Note

You can register multiple evidence schemas pertaining to different implied evidence. Additional schemas can also be registered during subsequent updates to the evidence provider.

- e. Choose **Submit**.

Onboard a new Evidence Provider

Evidence provider details

Evidence Provider Name
Please provide the evidence provider's name

Evidence Provider Id
Please provide a unique evidence provider id - OPTIONAL

Description
Please provide a brief description for your evidence provider

Add Evidence Schema
Add an evidence schema

Cancel

Onboard a new Evidence Provider page

On submission of the form, an API key is generated for the evidence provider, which must be used when subsequently creating evidence against the provider.

Successfully onboarded a new evidence provider ✕

Your new evidence provider details are provided below. Please include the evidence provider id and API Key in your request when creating new evidences.

Evidence Provider Id
97c45c1a-4a9b-40bd-8f70-3dd9af18268f

Evidence Provider Name
Test Custom Control

API Key
4aJmkbkFMEXf6QecLTqj7JE7CCObb0X7nisQUFL1

Success message with API key for evidence provider

Step 3: Store evidence

Evidence records are created by onboarded evidence providers. The solution provides an API and GUI for evidence creation and storage.

1. [Sign in to the web interface](#), then from the left navigation menu, select **Evidence**, and then **Evidences**.
2. Choose **Create an evidence**.

The screenshot displays the 'Verifiable Evidence Store' web interface. The main heading is 'Evidence history' with a 'Create an evidence' button in the top right. Below the heading is a search bar labeled 'Evidence Search' with 'Reset' and 'Search' buttons. A table titled 'Evidence (13)' shows a list of evidence records. The table has columns for Id, Provider, Evidence type, Target Id, and Additional Tan. One record is visible with the following details:

Id	Provider	Evidence type	Target Id	Additional Tan
a9c6d9e1-238c-4485-b3a4-299d4985169e	Security Hub Evidence Collector	sec-hub-evidence	arn:aws:s3:::finding-test-bucket-two	1234, Software and Regulatory arn:aws:macie: /macie, dev, en

Create an evidence

3. On the **Record a new Evidence** page:
 - a. Select the evidence provider from the list.

If an evidence provider has not been onboarded, refer to [Step 2. Onboard an evidence provider](#).

- b. Provide the API key for the selected evidence provider.
- c. Select an evidence schema.
- d. Provide the evidence target ID.
- e. Provide the evidence content.

Note

You can record multiple evidence artifacts. Choose **Add new item** to provide additional target IDs and **Add New Artifact** to attach additional evidence artifacts.

f. Choose Submit.

Verifiable Evidence Store

Verifiable Evidence StoreX

Home > Evidences > Create

Record a new Evidence

Evidence details

Evidence Provider
Please select the evidence provider

API Key
Please provide the API Key for the selected evidence provider

Evidence Schema
Please select the evidence schema

Evidence Target Id
Please provide the evidence target id

Add new item
Please provide any additional target ids

Evidence Content
Please provide the evidence content

Add New Artifact
Add an attachment

Cancel Reset Submit

Record a new Evidence page

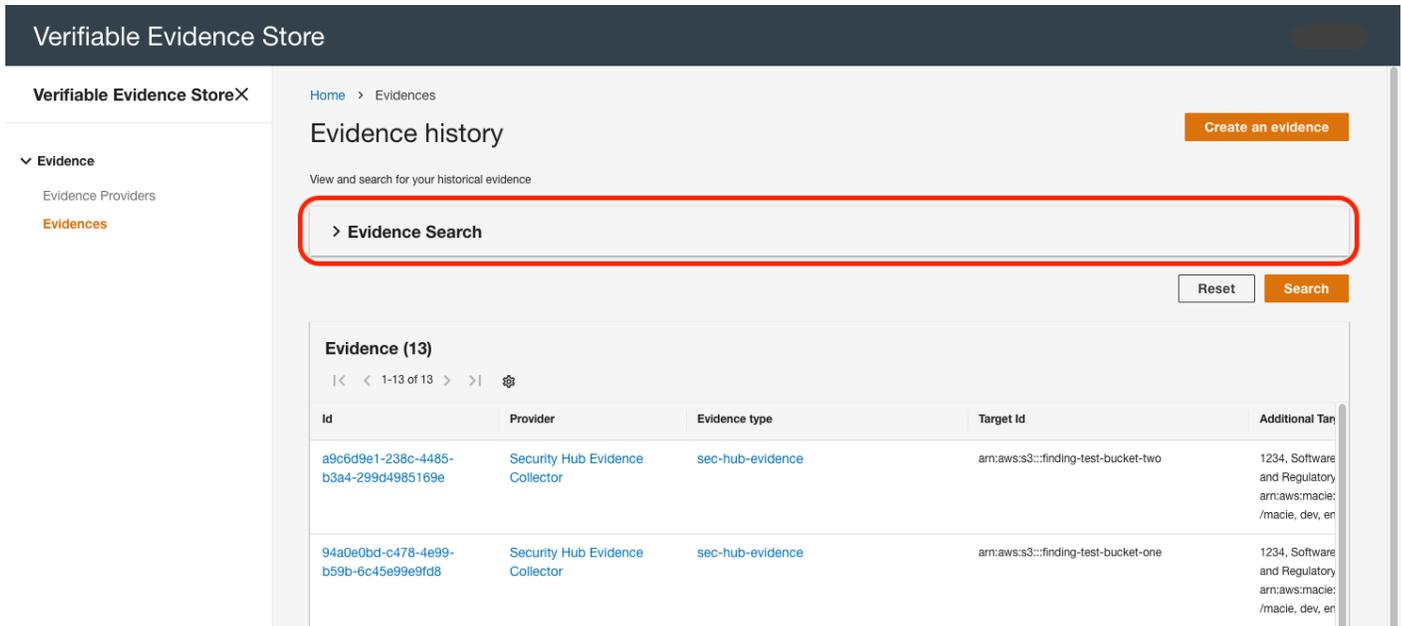
A success message confirms submission of the form.

Step 4: Search, retrieve, and inspect evidences

1. [Sign in to the web interface](#), then from the left navigation menu, select **Evidence**, and then **Evidences**.

Currently stored evidences appear in the list, and you can scroll and page through results.

2. Select the **Evidence Search** bar.



The screenshot shows the Verifiable Evidence Store interface. The main heading is "Verifiable Evidence Store". Below it, there's a navigation bar with "Home > Evidences". The main content area is titled "Evidence history" and includes a "Create an evidence" button. Below this, there's a search bar labeled "Evidence Search" which is highlighted with a red box. To the right of the search bar are "Reset" and "Search" buttons. Below the search bar, there's a table titled "Evidence (13)" with columns: "Id", "Provider", "Evidence type", "Target Id", and "Additional Tan". The table contains two rows of evidence data.

Id	Provider	Evidence type	Target Id	Additional Tan
a9c6d9e1-238c-4485-b3a4-299d4985169e	Security Hub Evidence Collector	sec-hub-evidence	arn:aws:s3:::finding-test-bucket-two	1234, Software and Regulatory arn:aws:macie:/macie_dev_en
94aDe0bd-c478-4e99-b59b-6c45e99e9fd8	Security Hub Evidence Collector	sec-hub-evidence	arn:aws:s3:::finding-test-bucket-one	1234, Software and Regulatory arn:aws:macie:/macie_dev_en

Evidence Search on the Evidence history page

3. Use the relevant fields in the **Evidence Search** box to specify your search criteria, and choose **Search**.
4. Review the results in the **Evidence** list below and select the evidence ID to inspect the evidence.

In addition to displaying information such as evidence details, issuing provider details, evidence schema, targets that the evidence was issued for, the web interface also displays the evidence's verified status. To provide the status, the solution performs a cryptographic verification of all of the evidence's information and returns a flag indicating whether or not the evidence content is still genuine, and has not been tampered with or altered in any way.

Attachments, additional metadata, and revisions are also presented in the **Evidence Details** page.

Evidence Details

▼ Evidence Details - a9c6d9e1-238c-4485-b3a4-299d4985169e

Provider
[Security Hub Evidence Collector](#)

Target Id
arn:aws:s3:::finding-test-bucket-two

Status
✔ Verified

Evidence Type
[sec-hub-evidence](#)

Additional Target Ids
1234, Software and Configuration Checks/Industry and Regulatory Standards/PCI-DSS, arn:aws:macie:ap-southeast-2::product/aws/macie, dev, env-1234, test-app

Created At
Mar 29, 2022, 12:36:51 PM

> Evidence content

Attachments - (1) | Metadata

File Name
7d76107e-26c1-4721-bdd5-05e93a85c9d8-1648582611355.json

Revisions - (10)

|< < 1-10 of 10 > >| ⚙

Revision Version	Created at
current	Feb 15, 2022, 3:38:28 PM

Evidence Details page

AWS Security Hub and AWS Config Evidence Collector

The solution's built-in Evidence Collector supports integration with AWS security and compliance services. This pre-configured evidence provider receives and records findings from AWS Security Hub and AWS Config.

Findings generated from these two services are pushed through EventBridge and Amazon SQS before being picked up by the Evidence Collector Lambda function, where they are processed and published to the Evidence Store for long-term storage.

Once persisted in the Evidence Store, these records can be queried and searched through the Verifiable Controls Evidence Store's API or web application, much like other evidence types.

By default, the Evidence Collector is configured to record findings from the following services via Security Hub integration:

- AWS Identity and Access Management Access Analyzer
- Amazon GuardDuty
- AWS Firewall Manager
- Amazon Inspector
- Amazon Macie
- Amazon Detective
- AWS Systems Manager Patch Manager

The AWS Security Hub and AWS Config Evidence Collector produces evidences that contain the following information:

- The original Security Hub or AWS Config finding's severity
- The finding's status
- The finding's summary
- The finding's created and updated timestamp
- The finding's account ID and AWS Region
- The finding's ID
- The finding's remediation recommendation
- The finding's source

- The finding's product

For further instructions on how to customize this list of finding sources, refer to the solution's [README.md](#) file.

Home > Evidenceproviders > Security-hub-evidence-collector

Security Hub Evidence Collector

Evidence Provider Details - Security Hub Evidence Collector

Details about the Security Hub Evidence Collector evidence provider.

Name Security Hub Evidence Collector	Status Active	Schemas Available 1
Description This provider collects evidences from Amazon Security Hub findings	Created At 12/22/2021, 3:25 PM	

Schema List

Search

Schema name

- [sec-hub-evidence-1.0](#)

Security Hub Evidence Collector registered as an evidence provider

Using the solution via APIs

Create an IAM policy to access the API

This solution does not automatically create an AWS Identity and Access Management (IAM) policy to invoke the created API. Follow this procedure to implement an IAM policy for access to the API.

Grant access to the entire API

Use the following procedure to grant access to the entire API to mask images and text in the JSON document. Note that using this procedure will allow a user to mask health data and view all information.

1. In the following JSON document, replace `us-east-1` with the AWS Region you are deploying in.
2. Replace `123456789012` with your account ID.
3. Replace `ab12cd3efg` with your API Gateway ID. You can find your ID in the Outputs tab of the AWS CloudFormation stack deployment.
4. Replace `prod` with the name of your staging environment. Note that if you did not change this in the mappings section of the AWS CloudFormation template when deploying, you can leave as is.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "execute-api:Invoke",
        "apigateway:PUT",
        "apigateway:POST",
        "apigateway:GET"
      ],
      "Resource": [
        "arn:aws:execute-api:<us-east-1>:<123456789012>:<ab12cd3efg>/<prod>/**",
        "arn:aws:apigateway:<us-east-1>::/restapis/<ab12cd3efg>/resources/**"
      ]
    }
  ]
}
```

```
}
```

Create the IAM policy

Use the following procedure to create the access policies:

1. Navigate to [AWS Identity and Access Management console](#).
2. In the navigation pane, select **Policies**, and then select the **Create policy** button.
3. Navigate to the **JSON** tab.
4. Copy and paste the modified JSON document you modified in the previous section for the access policy you want to create.

Create Evidence Provider API

As an alternative to using the web user interface to onboard an evidence provider, you can invoke the Create Evidence Provider API by issuing a POST request to `/providers`. The input name is required and you can specify multiple evidence schemas with the request. On invoking the API, the solution generates a unique provider ID and an API key for the name of the provider supplied, and returns them to the client.

You can register evidence schemas for a provider as part of the Create Evidence Provider API call, or as part of a separate API call using a particular provider ID in a POST request to `/providers/{provider-id}/schemas`. The required inputs to this request are `schema ID`, `provider ID` (generated earlier), and `content`, which must correspond to a valid JSON schema.

Create Evidence API

As an alternative to using the web user interface store evidence, you can invoke the Create Evidence API by issuing a POST request to `/evidences`. Like the web interface, evidence provider, evidence schema, a valid API key matching the onboarded evidence provider, target ID, and evidence content are required. On invoking the API, the solution generates a unique evidence record ID, and returns this to the client.

You can register evidence attachments for an evidence record as part of the Create Evidence API call.

Get Evidence API

As an alternative to using the web user interface to search, retrieve, and inspect evidences, you can invoke the Get Evidence API by issuing a GET request to `/evidences`. Like the web interface, you can provide any combination of search criteria.

Detailed information for a specific evidence record can be retrieved by invoking the `GET /evidences/{evidenceId}` API. The verified status of the evidence can be queried by issuing a GET request to the `/evidences/{evidenceId}/verificationstatus` API.

For APIs, refer to the solution's [API.md](#) file.

Troubleshooting

If you need help with this solution, contact Support to open a support case for this solution.

Contact Support

If you have [AWS Developer Support](#), [AWS Business Support](#), or [AWS Enterprise Support](#), you can use the Support Center to get expert assistance with this solution. The following sections provide instructions.

Create case

1. Sign in to [Support Center](#).
2. Choose **Create case**.

How can we help?

1. Choose **Technical**.
2. For **Service**, select **Solutions**.
3. For **Category**, select **Other Solutions**.
4. For **Severity**, select the option that best matches your use case.
5. When you enter the **Service**, **Category**, and **Severity**, the interface populates links to common troubleshooting questions. If you can't resolve your question with these links, choose **Next step: Additional information**.

Additional information

1. For **Subject**, enter text summarizing your question or issue.
2. For **Description**, describe the issue in detail.
3. Choose **Attach files**.
4. Attach the information that AWS Support needs to process the request.

Help us resolve your case faster

1. Enter the requested information.
2. Choose **Next step: Solve now or contact us**.

Solve now or contact us

1. Review the **Solve now** solutions.
2. If you can't resolve your issue with these solutions, choose **Contact us**, enter the requested information, and choose **Submit**.

Additional resources

AWS services

- [AWS Identity and Access Management](#)
- [Amazon Virtual Private Cloud](#) (Amazon VPC)
- [Amazon API Gateway](#)
- [AWS Lambda](#)
- [Amazon Simple Storage Service](#) (Amazon S3)
- [Amazon DynamoDB](#)
- [Amazon EventBridge](#)
- [Amazon CloudFront](#)
- [Amazon Quantum Ledger Database](#)
- [Amazon Kinesis Data Streams](#)
- [Amazon Data Firehose](#)
- [Amazon OpenSearch Service](#)
- [Amazon CloudWatch](#)
- [AWS X-Ray](#)
- [Amazon Simple Queue Service](#) (Amazon SQS)
- [Amazon Cognito](#)

Uninstall the solution

To delete the solution stacks with CDK, run `cdk destroy` from the `infra` folder.

Note

All resources deployed by the solution are removed when the stack is deleted.

You can also uninstall the Verifiable Controls Evidence Store solution via the AWS Management Console or by using the AWS Command Line Interface. AWS Solutions Implementations do not automatically delete data storage resources, in case there is stored data that may need to be retained. You must therefore manually delete the following resources to completely remove the solution from your account:

- OpenSearch Service domain
- Evidence content S3 bucket
- Evidence attachment S3 bucket
- Synthetic Canary log S3 bucket
- QLDB ledger journal export S3 bucket
- Evidence providers DynamoDB table
- Evidence schemas DynamoDB table
- Evidence QLDB ledger

Using the AWS Management Console

1. Sign in to the [AWS CloudFormation console](#).
2. On the **Stacks** page, select the following stacks in order, and choose **Delete** one by one:
 - AGSSecurityHubEvidenceCollector
 - AGSEvidenceStore
 - AGSSharedInfra-WebClient
 - AGSSharedInfra-BaseInfra

Important

Verify the first three stacks are completely deleted before starting to delete AGSSharedInfra-BaseInfra.

Using AWS Command Line Interface

Determine whether the AWS Command Line Interface (AWS CLI) is available in your environment. For installation instructions, refer to What Is the AWS Command Line Interface in the *AWS CLI User Guide*. After confirming that the AWS CLI is available, run the following command.

```
$ aws cloudformation delete-stack --stack-name AGSEvidenceStore
```

Please follow the list and order in section above when deleting stacks deployed by this solution.

Deleting the Amazon S3 buckets

This solution is configured to retain the solution-created Amazon S3 bucket (for deploying in an opt-in Region) if you decide to delete the AWS CloudFormation stack to prevent accidental data loss. After uninstalling the solution, you can manually delete this S3 bucket if you do not need to retain the data. Follow these steps to delete the Amazon S3 bucket.

1. Sign in to the [Amazon S3 console](#).
2. Choose **Buckets** from the left navigation pane.
3. Locate the S3 buckets with tags `agsService:AGSEvidenceStore`.
4. Select the S3 bucket and choose **Delete**.

To delete the S3 bucket using AWS CLI, run the following command:

```
$ aws s3 rb s3://<bucket-name> --force
```

Deleting DynamoDB tables

Follow these steps to delete the Evidence Store's DynamoDB tables using the AWS Management Console.

1. Sign in to the [Amazon DynamoDB console](#).
2. Navigate to the **Tables** section.
3. Select evidence-providers and evidence-schemas tables.
4. Choose **Delete**.

To delete the table using AWS CLI, run the following command:

```
$ aws dynamodb delete-table --table-name <tablename>
```

Deleting Amazon OpenSearch Service domain

Follow these steps to delete the Amazon OpenSearch Service domain using the AWS Management Console.

1. Sign in to the [Amazon OpenSearch Service console](#).
2. Navigate to the **Domains** section.
3. Select the evidence-read-replica domain and choose **Delete**.

To delete the domain using AWS CLI, run the following command:

```
$ aws opensearch delete-domain --domain-name evidence-read-replica
```

Deleting Amazon QLDB ledger

Follow these steps to delete the QLDB ledger using the AWS Management Console.

1. Sign in to the [Amazon QLDB console](#).
2. Select the Evidences ledger and **choose** Delete.

To delete the domain using AWS CLI, run the following command:

```
$ aws qldb delete-ledger --ledger-name Evidences
```

Source code

Visit our [GitHub repository](#) to download the source files for this solution and to share your customizations with others. The Verifiable Controls Evidence Store templates are generated using the [AWS Cloud Development Kit \(AWS CDK\)](#). Refer to the [README.md](#) file for additional information.

Contributors

- Van Vo Thanh
- Hu Jin
- Hafiz Saadullah
- Deenadayaalan Thirugnanasambandam
- Shaunak Joshi
- Swapnil Ogale
- Daniil Millwood

Revisions

Date	Change
June 2022	Initial release
May 2023	<ul style="list-style-type: none">• Updated the architecture diagram to display the multiple evidence collector functions available as part of the solution.• Added instructions on how to configure the S3 collector to specify the source S3 buckets before deploying the solution.• Added instructions on how to apply bucket policies for the S3 collector, after deploying the stack.• Mitigated impact caused by new default settings for S3 Object Ownership (ACLs disabled) for all new S3 buckets. For more information, refer to the CHANGELOG.md file in the GitHub repository.

Notices

Customers are responsible for making their own independent assessment of the information in this document. This document: (a) is for informational purposes only, (b) represents AWS current product offerings and practices, which are subject to change without notice, and (c) does not create any commitments or assurances from AWS and its affiliates, suppliers, or licensors. AWS products or services are provided “as is” without warranties, representations, or conditions of any kind, whether express or implied. AWS responsibilities and liabilities to its customers are controlled by AWS agreements, and this document is not part of, nor does it modify, any agreement between AWS and its customers.

Verifiable Controls Evidence Store is licensed under the terms of the [Apache License, Version 2.0](#).