
Navigating GDPR Compliance on AWS

AWS Whitepaper



Navigating GDPR Compliance on AWS: AWS Whitepaper

Copyright © 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

Abstract	1
Abstract	1
General Data Protection Regulation Overview	2
Changes the GDPR Introduces to Organizations Operating in the EU	2
AWS Preparation for the GDPR	2
AWS Data Processing Addendum (DPA)	2
The Role of AWS Under the GDPR	3
AWS as a Data Processor	3
AWS as a Data Controller	3
Shared Security Responsibility Model	3
Strong Compliance Framework and Security Standards	5
AWS Compliance Program	5
Cloud Computing Compliance Controls Catalog	5
The CISPE Code of Conduct	6
Data Access Controls	7
AWS Identity and Access Management	7
Temporary Access Tokens Through AWS STS	8
Multi-Factor-Authentication	8
Access to AWS Resources	9
Defining Boundaries for Regional Services Access	9
Control Access to Web Applications and Mobile Apps	10
Monitoring and Logging	12
Manage and Configure Assets with AWS Config	12
Compliance Auditing and Security Analytics	13
Collecting and Processing Logs	14
Discovering and Protecting Data at Scale	15
Centralized Security Management	16
Protecting your Data on AWS	18
Encrypt Data at Rest	18
Encrypt Data in Transit	19
Encryption Tools	19
AWS Key Management Service	19
AWS Cryptographic Services and Tools	22
Data Protection by Design and by Default	22
How AWS Can Help	23
Contributors	25
Document history	26
Notices	27

Navigating GDPR Compliance on AWS

Publication date: **April 2022** ([Document history](#) (p. 26))

Abstract

This document provides information about services and resources that Amazon Web Services (AWS) offers customers to help them align with the requirements of the General Data Protection Regulation (GDPR) that might apply to their activities. These include adherence to IT security standards, the AWS Cloud Computing Compliance Controls Catalog (C5) attestation, adherence to the Cloud Infrastructure Services Providers in Europe (CISPE) Code of Conduct, data access controls, monitoring and logging tools, encryption, and key management.

General Data Protection Regulation Overview

The [General Data Protection Regulation \(GDPR\)](#) is a European privacy law ([Regulation 2016/679 of the European Parliament and of the Council of April 27, 2016](#)) that became enforceable on May 25, 2018. The GDPR replaces the EU Data Protection Directive ([Directive 95/46/EC](#)), and is intended to harmonize data protection laws throughout the European Union (EU) by applying a single data protection law that is binding throughout each EU member state.

The GDPR applies to all processing of *personal data* either by organizations that have an establishment in the EU, or to organizations that process personal data of EU residents when offering goods or services to individuals in the EU or monitoring the behavior of EU residents in the EU. Personal data is any information relating to an identified or identifiable natural person.

Changes the GDPR Introduces to Organizations Operating in the EU

One of the key aspects of the GDPR is that it creates consistency across EU member states on how personal data can be processed, used, and exchanged securely. Organizations must demonstrate the security of the data they are processing and their compliance with the GDPR on a continual basis, by implementing and regularly reviewing technical and organizational measures, as well as compliance policies applicable to the processing of personal data. EU supervisory authorities can issue fines of up to EUR 20 million, or 4% of annual worldwide turnover, whichever is higher, for a breach of the GDPR.

AWS Preparation for the GDPR

AWS compliance, data protection, and security experts work with customers around the world to answer their questions and help them prepare to run workloads in the cloud under the GDPR. These teams also review the readiness of AWS against the requirements of the GDPR.

Note

We can confirm that all AWS services can be used in compliance with the GDPR.

AWS Data Processing Addendum (DPA)

AWS offers a GDPR-compliant Data Processing Addendum (GDPR DPA), which enables customers to comply with GDPR contractual obligations. The [AWS GDPR DPA is incorporated into the AWS Service Terms](#) and applies automatically to all customers globally who require it to comply with the GDPR.

On 16 July 2020, the Court of Justice of the European Union (CJEU) issued a ruling regarding the EU-US Privacy Shield and Standard Contractual Clauses (SCCs), also known as “model clauses.” The CJEU ruled that the EU-US Privacy Shield is no longer valid for the transfer of personal data from the European Union (EU) to the United States (US). However, in the same ruling, the CJEU validated that companies can continue to use SCCs as a mechanism for transferring data outside of the EU.

Following this ruling, AWS customers and partners can continue to use AWS to transfer their content from Europe to the US and other countries, in compliance with EU data protection laws – including the General Data Protection Regulation (GDPR). AWS customers can rely on the SCCs included in the AWS Data Processing Addendum (DPA) if they choose to transfer their data outside the European Union in compliance with GDPR. As the regulatory and legislative landscape evolves, we will work to ensure that our customers and partners can continue to enjoy the benefits of AWS everywhere they operate. For additional information, see the [EU-US Privacy Shield FAQ](#).

The Role of AWS Under the GDPR

Under the GDPR, AWS acts as both a data processor and a data controller.

Under Article 32, controllers and processors are required to “...implement appropriate technical and organizational measures” that consider “the state of the art and the costs of implementation and the nature, scope, context and purposes of processing as well as the risk of varying likelihood and severity for the rights and freedoms of natural persons”. The GDPR provides specific suggestions for what types of security actions may be required, including:

- The [pseudonymization](#) and encryption of personal data.
- The ability to ensure the ongoing confidentiality, integrity, availability and resilience of processing systems and services.
- The ability to restore the availability and access to personal data in a timely manner, in the event of a physical or technical incident.
- A process to regularly test, assess, and evaluate the effectiveness of technical and organizational measures to ensure the security of the processing.

AWS as a Data Processor

When customers and AWS Partner Network (APN) Partners use AWS Services to process personal data in their content, AWS acts as a data processor. Customers and APN Partners can use the controls available in AWS services, including security configuration controls, to process personal data. Under these circumstances, the customer or APN Partners may act as a data controller or a data processor, and AWS acts as a data processor or sub-processor. The AWS GDPR-compliant Data Processing Addendum (DPA) incorporates the commitments of AWS as a data processor.

AWS as a Data Controller

When AWS collects personal data and determines the purposes and means of processing that personal data, it acts as a data controller. For example, when AWS processes account information for account registration, administration, services access, or contact information for the AWS account to provide assistance through customer support activities, it acts as a data controller.

Shared Security Responsibility Model

Security and Compliance is a shared responsibility between AWS and the customer. When customers move their computer systems and data to the cloud, security responsibilities are shared between the customer and the cloud service provider. When customers move to the AWS Cloud, AWS is responsible for protecting the global infrastructure that runs all of the services offered in the AWS Cloud. For abstracted service, such as Amazon S3 and Amazon DynamoDB AWS is also responsible for the security of the operating system and platform. Customers and APN Partners, acting either as data controllers

or data processors, are responsible for anything they put in the cloud or connect to the cloud. This differentiation of responsibility is commonly referred to as security *of* the cloud versus security *in* the cloud. This shared model can help reduce customers' operational burden, and provide them with the necessary flexibility and control to deploy their infrastructure in the AWS Cloud. For more information, see the [AWS Shared Responsibility Model](#).

The GDPR does not change the AWS shared responsibility model, which continues to be relevant for customers and APN Partners who are focused on using cloud computing services. The shared responsibility model is a useful approach to illustrate the different responsibilities of AWS (as a data processor or sub-processor) and customers or APN Partners (as either data controllers or data processors) under the GDPR.

Strong Compliance Framework and Security Standards

According to the GDPR, appropriate technical and organizational measures might need to include "... the ability to ensure the ongoing confidentiality, integrity, availability and resilience of the processing systems and services," as well as reliable restore, testing, and overall risk management processes.

AWS Compliance Program

AWS continually maintains a high bar for security and compliance across all of our global operations. Security has always been our highest priority – truly "job zero." AWS regularly undergoes independent third-party attestation audits to provide assurance that control activities are operating as intended. More specifically, AWS is audited against a variety of global and regional security frameworks dependent on region and industry. Currently, AWS participates in over 50 different audit programs.

The results of these audits are documented by the assessing body and made available for all AWS customers through [AWS Artifact](#). AWS Artifact is a no-cost, self-service portal for on-demand access to AWS compliance reports. When new reports are released, they are made available in AWS Artifact, allowing customers to continuously monitor the security and compliance of AWS with immediate access to new reports.

Customers can take advantage of internationally recognized certifications and accreditations, demonstrating compliance with rigorous international standards, such as ISO 27017 for cloud security, ISO 27018 for cloud privacy, SOC 1, SOC 2 and SOC 3, PCI DSS Level 1 and others. AWS also helps customers meet local security standards such as BSI's Common Cloud Computing Controls Catalogue (C5), a German government-backed attestation.

For more detailed information about the AWS certification programs, reports, and third-party attestations, see [AWS Compliance Programs](#). For service-specific information, see [AWS Services in Scope](#).

Cloud Computing Compliance Controls Catalog

[Cloud Computing Compliance Controls Catalog \(C5\)](#) is a German government-backed attestation scheme that was introduced in Germany by the Federal Office for Information Security (BSI). It was created to help organizations demonstrate operational security against common cyberattacks within the context of the German government's [Security Recommendations for Cloud Providers](#).

The technical and organizational measures of data protection and the measures for information security target data security to ensure confidentiality, integrity and availability. C5 defines security requirements that can be also relevant for data protection. AWS customers and their compliance advisors can use the C5 attestation as a resource to understand the range of IT-Security assurance services that AWS offers them as they move their workloads to the cloud. C5 adds the regulatory-defined IT-Security level equivalent to the IT-Grundschutz, with the addition of cloud-specific controls.

C5 adds more controls that provide information pertaining to data location, service provisioning, place of jurisdiction, existing certification, information disclosure obligations, and a full-service description. Using this information, you can evaluate how legal regulations (such as data privacy), your own policies, or the threat environment relate to your use of cloud computing services.

The CISPE Code of Conduct

[CISPE](#) (Cloud Infrastructure Services Providers in Europe) is a coalition of cloud computing leaders serving millions of European customers. The [CISPE Data Protection Code of Conduct](#) (CISPE Code) is the first pan-European data protection code of conduct for cloud infrastructure service providers under Article 40 of the European Union's General Data Protection Regulation (GDPR). It was approved by the [European Data Protection Board](#) (EDPB) in May 2021 and formally adopted by the [French Data Protection Authority](#) (CNIL), acting as the competent supervisory authority, in June 2021.

The CISPE Code assures organizations that their cloud infrastructure service provider meets the requirements applicable to a data processor under the GDPR. This gives cloud customers additional confidence that they can choose services that have been independently verified for their compliance with the GDPR.

The CISPE Code goes beyond GDPR compliance by requiring cloud infrastructure service providers to give customers the choice to select services that store and process customer data exclusively within the European Economic Area. Cloud infrastructure service providers must also commit that they will not access or use any customer data, except as necessary to provide and maintain the declared services. In particular, the cloud infrastructure service providers must commit to not use customer data for their own purposes, including for data mining, profiling or direct marketing. Ernst and Young CertifyPoint (EYCP) independently certified AWS services listed on the [CISPE Public Register](#) complying with the CISPE Code. EYCP was the first "monitoring body" accredited by CNIL to verify cloud infrastructure provider's compliance with the CISPE Code.

AWS supports more security standards and compliance certifications than any other cloud provider, and we are continuously reviewing the needs of our customers as the regulatory environment evolves. The CISPE Code provides an added level of assurance to our customers that AWS Cloud services can be used in compliance with the GDPR and addresses our customers' compliance requirements today.

Data Access Controls

Article 25 of the GDPR states that the controller “shall implement appropriate technical and organisational measures for ensuring that, by default, only personal data which are necessary for each specific purpose of the processing are processed”. The following AWS access control mechanisms can help customers comply with this requirement by allowing only authorized administrators, users, and applications to get access to AWS resources and customer data.

AWS Identity and Access Management

When you create an AWS account, a *root* user account is automatically created for your AWS account. This user account has complete access to all your AWS services and resources in your AWS account. Instead of using this account for everyday tasks, you should only use it to initially create additional roles and user accounts, and for administrative activities that require it. AWS recommends that you apply the principle of least privilege from the start: define different user accounts and roles for different tasks, and specify the minimum set of permissions required to complete each task. This approach is a mechanism for tuning a key concept introduced in GDPR: data protection by design. [AWS Identity and Access Management](#) (IAM) is a web services that you can use to securely control access to your AWS resources.

Users and roles define IAM identities with specific permissions. An authorized user can assume an IAM role to perform specific tasks. Temporary credentials are created when the role is assumed. For example, you can use IAM roles to securely provide applications that run in [Amazon Elastic Compute Cloud](#) (Amazon EC2) with temporary credentials required to access other AWS resources, such as Amazon S3 buckets, and [Amazon Relational Database Service](#) (Amazon RDS) or [Amazon DynamoDB](#) databases. Similarly, [execution roles](#) provide [AWS Lambda](#) functions with the required permissions to access other AWS Services and resources, such as [Amazon CloudWatch Logs](#) for log streaming or reading a message from an [Amazon Simple Queue Service](#) (Amazon SQS) queue. When you create a role, you add policies to it to define authorizations.

To help customers monitor resources policies and identify resources with public or cross-account access they may not intend, [IAM Access Analyzer](#) can be enabled to generate comprehensive findings that identify resources that can be accessed from outside an AWS account. IAM Access Analyzer evaluates resource policies using mathematical logic and inference to determine the possible access paths allowed by the policies. IAM Access Analyzer continuously monitors for new or updated policies, and it analyzes permissions granted using policies for IAM roles--but also for services resources like Amazon S3 buckets, [AWS Key Management Service](#) (AWS KMS) keys, Amazon SQS queues, and Lambda functions.

[Access Analyzer for S3](#) alerts you when buckets are configured to allow access to anyone on the internet or other AWS accounts, including AWS accounts outside of your organization. When reviewing an at-risk bucket in Access Analyzer for Amazon S3, you can block all public access to the bucket with a single click. AWS recommends that you block all access to your buckets unless you require public access to support a specific use case. Before you block all public access, ensure that your applications will continue to work correctly without public access. For more information, see [Using Amazon S3 to Block Public Access](#).

IAM also provides last accessed information to help you identify unused permissions so that you can remove them from the associated principals. Using last accessed information, it is possible to refine your policies and allow access to only those services and actions that are needed. This helps to better adhere to and apply the [best practice of least privilege](#). You can view last accessed information for entities or policies that exist in IAM, or across an entire [AWS Organizations](#) environment.

Temporary Access Tokens Through AWS STS

You can use the [AWS Security Token Service \(AWS STS\)](#) to create and provide trusted users with temporary security credentials that grant access to your AWS resources. Temporary security credentials work almost identically to the long-term access key credentials that you provide for your IAM users, with the following differences:

- Temporary security credentials are for short-term use. You can configure the amount of time that they are valid, from 15 minutes up to a maximum of 12 hours. After temporary credentials expire, AWS does not recognize them or allow any kind of access from API requests made with them.
- Temporary security credentials are not stored with the user account. Instead, they are generated dynamically and provided to the user when requested. When (or before) temporary security credentials expire, a user can request new credentials, if that user has permissions to do so.

These differences provide the following advantages when you use temporary credentials:

- You do not have to distribute or embed long-term AWS security credentials with an application.
- Temporary credentials are the basis for roles and identity federation. You can provide access to your AWS resources to users by defining a temporary AWS identity for them.
- Temporary security credentials have a limited customizable lifespan. Because of this, you do not have to rotate them or explicitly revoke them when they're no longer needed. After temporary security credentials expire, they cannot be reused. You can specify the maximum amount of time the credentials are valid.

Multi-Factor-Authentication

For extra security, you can add two-factor authentication to your AWS account and to IAM users. With multi-factor authentication (MFA) enabled, when you sign into the [AWS Management Console](#), you are prompted for your user name and password (the first factor), as well as an authentication response from your AWS MFA device (the second factor). You can enable MFA for your AWS account and for individual IAM users you have created in your account. You can also use MFA to control access to AWS service APIs.

For example, you can define a policy that allows full access to all AWS API operations in Amazon EC2, but explicitly denies access to specific API operations—such as `StopInstances` and `TerminateInstances`—if the user is not authenticated with MFA.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "AllowAllActionsForEC2",
      "Effect": "Allow",
      "Action": "ec2:*",
      "Resource": "*"
    },
    {
      "Sid": "DenyStopAndTerminateWhenMFAIsNotPresent",
      "Effect": "Deny",
      "Action": [
        "ec2:StopInstances",
        "ec2:TerminateInstances"
      ],
      "Resource": "*",
      "Conditions": {
        "BoolIfExists": {"aws:MultiFactorAuthPresent": false}
      }
    }
  ]
}
```

```
}  
  }  
}
```

To add an extra layer of security to your Amazon S3 buckets, you can configure [MFA Delete](#), which requires additional authentication to change the versioning state of a bucket and permanently delete an object version. MFA Delete provides added security in the event that your security credentials are compromised.

To use MFA Delete, you can use either a hardware or virtual MFA device to generate an authentication code. See the [Multi-factor Authentication page](#) for a list of supported hardware or virtual MFA devices.

Access to AWS Resources

To implement granular access to your AWS resources, you can grant different levels of permissions to different people for different resources. For example, you can allow only some users complete access to Amazon EC2, Amazon S3, DynamoDB, [Amazon Redshift](#), and other AWS Services.

For other users, you can allow read-only access to only some Amazon S3 buckets; permission to administer only some Amazon EC2 instances, or access to only your billing information.

The following policy is an example of one method you can use to allow all actions on a specific Amazon S3 bucket and explicitly deny access to every AWS service that is not Amazon S3.

```
{  
  "Version": "2012-10-17",  
  "Statement": [  
    {  
      "Effect": "Allow",  
      "Action": "s3:*",  
      "Resource": [  
        "arn:aws:s3:::bucket-name",  
        "arn:aws:s3:::bucket-name/*"  
      ],  
    },  
    {  
      "Effect": "Deny",  
      "NotAction": "s3:*",  
      "NotResource": [  
        "arn:aws:s3:::bucket-name",  
        "arn:aws:s3:::bucket-name/*"  
      ]  
    }  
  ]  
}
```

You can attach a policy to a user account or to a role. For other examples of IAM policies, see [Example IAM Identity-Based Policies](#).

Defining Boundaries for Regional Services Access

As a customer, you maintain ownership of your content, and you select which AWS services can process, store, and host your content. AWS does not access or use your content for any purpose without your consent. Based on the Shared Responsibility Model, you choose the AWS Regions in which your content

is stored, allowing you to deploy AWS services in the locations of your choice, in accordance with your specific geographic requirements. For example, if you want to ensure your content is located only in Europe, you can choose to deploy AWS services exclusively in one of the European AWS Regions.

IAM policies provide a simple mechanism to limit access to services in specific Regions. You can add a global condition ([aws:RequestedRegion](#)) to the IAM policies attached to your IAM Principals to enforce this for all AWS services. For example, [the following policy](#) uses the `NotAction` element with the `Deny` effect, which explicitly denies access to all of the actions not listed in the statement if the requested Region is not European. Actions in the CloudFront, IAM, [Amazon Route 53](#), and [AWS Support](#) services should not be denied because these are popular AWS global services.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "DenyAllOutsideRequestedRegions",
      "Effect": "Deny",
      "NotAction": [
        "cloudfront:*",
        "iam:*",
        "route53:*",
        "support:*"
      ],
      "Resource": "*",
      "Condition": {
        "StringNotLike": {
          "aws:RequestedRegion": [
            "eu-*"
          ]
        }
      }
    }
  ]
}
```

This sample IAM policy can also be implemented as a Service Control Policy (SCP) in AWS Organizations, which defines the permission boundaries applied to specific AWS accounts or Organizational Units (OUs) within an organization. This enables you to control user access to regional services in complex multi-account environments.

Geo-limiting capabilities exist for newly launched Regions. [Regions introduced after March 20, 2019](#) are disabled by default. You must enable these Regions before you can use them. If an AWS Region is disabled by default, you can use the AWS Management Console to enable and disable the Region. Enabling and disabling AWS Regions enables you to control whether users in your AWS account can access resources in that Region. For more information, see [Managing AWS Regions](#).

Control Access to Web Applications and Mobile Apps

AWS provides services for managing data access control within customer applications. If you need to add user login and access control features to your web applications and mobile apps, you can use [Amazon Cognito](#). [Amazon Cognito user pools](#) provide a secure user directory that scales to hundreds of millions of users. To protect the identity of the users, you can add multi-factor authentication (MFA) to your user pools. You can also use adaptive authentication, which uses a risk-based model to predict when you might need another authentication factor.

With [Amazon Cognito Identity Pools](#) (Federated Identities), you can see who accessed your resources and where the access originated (mobile app or web application). You can use this information to create IAM roles and policies that allow or deny access to a resource based on the type of access origin (mobile app or web application) and Identity Provider.

Monitoring and Logging

Article 30 of the GDPR states that “...each controller and, where applicable, the controller’s representative, shall maintain a record of processing activities under its responsibility”. This article also includes details about which information must be recorded when you monitor the processing of all personal data. Controllers and processors are also required to send breach notifications in a timely manner, so detecting incidents quickly is important. To help enable customers to comply with these obligations, AWS offers the following monitoring and logging services.

Manage and Configure Assets with AWS Config

[AWS Config](#) provides a detailed view of the configuration of many types of AWS resources in your AWS account. This includes how the resources are related to one another, and how they were previously configured, so you can see how the configurations and relationships change over time.

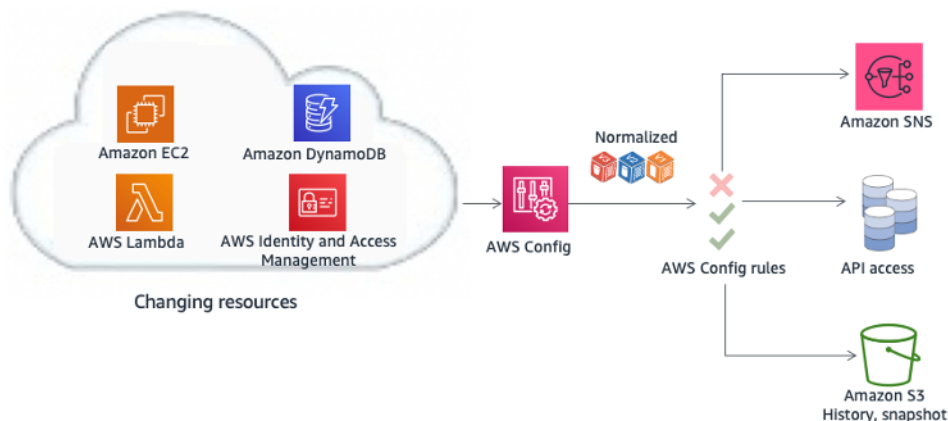


Figure 1 – Monitor configuration changes over time with AWS Config

An AWS resource is an entity that you can work with in AWS, such as an EC2 instance, an [Amazon Elastic Block Store](#) (Amazon EBS) volume, a security group, or an [Amazon Virtual Private Cloud](#) (Amazon VPC). For a complete list of AWS resources supported by AWS Config, see [Supported AWS Resource Types](#).

With AWS Config, you can do the following:

- Evaluate your AWS resource configurations for to verify the settings are correct.
- Get a snapshot of the current configurations of the supported resources that are associated with your AWS Account.
- Get configurations of one or more resources that exist in your account.
- Get historical configurations of one or more resources.
- Get a notification when a resource is created, modified, or deleted.
- See relationships between resources. For example, find all resources that use a particular security group.

Compliance Auditing and Security Analytics

With [AWS CloudTrail](#), you can continuously monitor AWS account activity. A history of the AWS API calls for your account is captured, including API calls made through the AWS Management Console, the AWS SDKs, the command line tools, and higher-level AWS services. You can identify which users and accounts called AWS APIs [for services that support CloudTrail](#), the source IP address the calls were made from, and when the calls occurred. You can integrate CloudTrail into applications using the API, automate trail creation for your organization, check the status of your trails, and control how administrators enable and disable CloudTrail logging.

CloudTrail logs can be aggregated from [multiple Regions](#) and [multiple AWS accounts](#) into a single Amazon S3 bucket. AWS recommends that you write logs--especially AWS CloudTrail logs--to an Amazon S3 bucket with restricted access in an AWS account designated for logging (Log Archive). The permissions on the bucket should prevent deletion of the logs, and they should also be encrypted at rest using Server-Side Encryption with Amazon S3-managed encryption keys (SSE-S3) or AWS KMS-managed keys (SSE-KMS). CloudTrail log file integrity validation can be used to determine whether a log file was modified, deleted, or unchanged after CloudTrail delivered it. This feature is built using industry standard algorithms: SHA-256 for hashing and SHA-256 with RSA for digital signing. This makes it computationally hard to modify, delete, or forge CT; log files without detection. You can use the AWS command line interface (AWS CLI) to validate the files in the location where CloudTrail delivered them.

CloudTrail logs aggregated in an Amazon S3 bucket can be analyzed for auditing purposes or for troubleshooting activities. Once the logs are centralized, you can integrate with Security Information and Event Management (SIEM) solutions or use AWS services, such as [Amazon Athena](#) or [CloudTrail Insights](#), to analyze them and [visualize them using Amazon QuickSight Dashboards](#). Once you have CloudTrail logs centralized, you can also use the same Log Archive account to centralize logs from other sources, such as CloudWatch Logs and AWS load balancers.



Figure 2 – Example architecture for compliance auditing and security analytics with AWS CloudTrail

AWS CloudTrail logs can also trigger preconfigured Amazon CloudWatch events. You can use these events to notify users or systems that an event has occurred, or for remediation actions. For example, if you want to monitor activities on your Amazon EC2 instances, you can create a [CloudWatch Event rule](#). When a specific activity happens on the Amazon EC2 instance and the event is captured in the logs, the rule triggers an AWS Lambda function, which sends a notification email about the event to the administrator. (See Figure 3.) The email includes details such as when the event happened, which user performed the action, Amazon EC2 details, and more. The following diagram shows the architecture of the event notification.

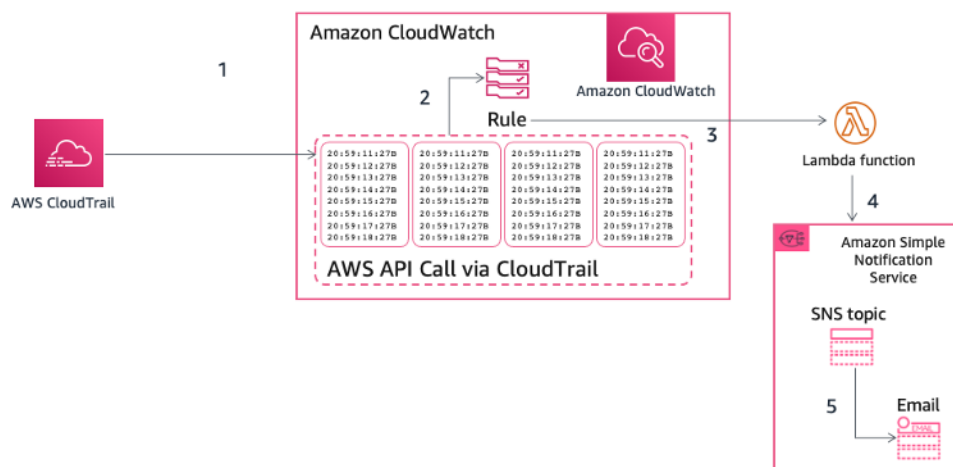


Figure 3 – Example of AWS CloudTrail event notification

Collecting and Processing Logs

CloudWatch Logs can be used to monitor, store, and access your log files from Amazon EC2 instances, AWS CloudTrail, Route 53, and other sources. See the [AWS Services That Publish Logs to CloudWatch Logs](#) documentation page.

Logs information includes, for example:

- Granular logging of access to Amazon S3 objects
- Detailed information about flows in the network through VPC-Flow Logs
- Rule-based configuration verification and actions with AWS Config rules
- Filtering and monitoring of HTTP access to applications with web application firewall (WAF) functions in CloudFront

Custom application metrics and logs can also be published to CloudWatch Logs by installing the [CloudWatch Agent](#) on Amazon EC2 instances or on-premises servers.

Logs can be analyzed interactively using CloudWatch Logs Insights, performing queries to help you respond more efficiently and effectively to operational issues.

CloudWatch Logs can be processed in near real-time by configuring subscription filters and delivered to other services such as an [Amazon OpenSearch Service](#) (OpenSearch Service) cluster, an [Amazon Kinesis](#) stream, an Amazon Kinesis Data Firehose stream, or Lambda for custom processing, analysis, or loading to other systems.

[CloudWatch metric filters](#) can be used to define patterns to look for in log data, transform them into numerical CloudWatch metrics, and set up alarms based on your business requirements. For example, following the AWS recommendation not to use the root user for everyday tasks, it is possible to [set up a specific CloudWatch metric filter](#) on a CloudTrail log (delivered to CloudWatch Logs) to create a Custom metric and configure an alarm to notify the relevant stakeholders when root credentials are used to access your AWS account.

Logs such as Amazon S3 server access logs, Elastic Load Balancing access logs, VPC flow logs, and AWS Global Accelerator flow logs can be delivered directly to an Amazon S3 bucket. For example, when you enable [Amazon Simple Storage Service server access logs](#), you can get detailed information regarding

the requests that are made to your Amazon S3 bucket. An access log record contains details about the request, such as the request type, the resources specified in the request, and the time and date the request was processed. For more information about the contents of a log message, see [Amazon Simple Storage Service Server Access Log Format](#) in the *Amazon Simple Storage Service Developer Guide*. Server access logs are useful for many applications because they give bucket owners insight into the nature of requests made by clients that are not under their control. By default, Amazon S3 does not collect service access logs, but when you enable logging, Amazon S3 usually delivers access logs to your bucket within a few hours. If you require a faster delivery or need to deliver logs to multiple destinations, [consider using CloudTrail logs](#) or a combination of both CloudTrail logs and Amazon S3. Logs can be encrypted at rest by configuring default object encryption in the destination bucket. The objects are encrypted using server-side encryption with either Amazon S3-managed keys (SSE-S3) or KMS keys (formerly CMKs) stored in [AWS Key Management Service](#) (AWS KMS).

Logs stored in an Amazon S3 bucket can be queried and analyzed using [Amazon Athena](#). Amazon Athena is an interactive query service that enables you to analyze data in S3 using standard SQL. You can use Athena to run ad-hoc queries using ANSI SQL, without the need to aggregate or load the data into Athena. Athena can process unstructured, semi-structured, and structured data sets and integrates with [Amazon QuickSight](#) for easy visualization.

Logs are also a useful source of information for automated threat detection. [Amazon GuardDuty](#) is a continuous security monitoring service that analyzes and processes events from several sources, such as VPC Flow Logs, CloudTrail management event logs, CloudTrail Amazon S3 data event logs, and DNS logs. It uses threat intelligence feeds, such as lists of malicious IP addresses and domains, and machine learning to identify unexpected and potentially unauthorized and malicious activity within your AWS environment. When you enable GuardDuty in a Region, it immediately starts analyzing your CloudTrail event logs. It consumes CloudTrail management and Amazon S3 data events directly from CloudTrail through an independent and duplicative stream of events.

Discovering and Protecting Data at Scale with Amazon Macie

Article 32 of the GDPR states that "...the controller and the processor shall implement appropriate technical and organisational measures to ensure a level of security appropriate to the risk, including inter alia as appropriate: [...]

(b) the ability to ensure the ongoing confidentiality, integrity, availability and resilience of processing systems and services;

[...]

(d) a process for regularly testing, assessing and evaluating the effectiveness of technical and organisational measures for ensuring the security of the processing."

Having an ongoing data classification process is critical for adjusting security data processing to the nature of data. If your organization manages sensitive data, monitor where it resides, protect it properly, and provide evidence that you are enforcing data security and privacy as required to meet regulatory compliance requirements. To help the customer identify and protect their sensitive data at scale, AWS offers [Amazon Macie](#), a fully managed data security and data privacy service that uses pattern matching and machine learning models for detection of Personally Identifiable Information (PII) to discover and protect sensitive data stored in S3 buckets. Amazon Macie scans these buckets and provides a data categorization of them using managed data identifiers that are designed to detect several categories of sensitive data. Macie can [detect PII](#) such as full name, email address, birth date, national identification number, taxpayer identification or reference number, and more. The customer can define custom data identifiers that reflect their organization's particular scenarios (for example, customer account numbers or internal data classification).

Amazon Macie continually evaluates the object inside the buckets and automatically provides a summary of findings (Figure 4) for any unencrypted or publicly accessible data discovered that match with the defined data category. This data can include alerts for any unencrypted, publicly accessible objects or buckets shared with AWS accounts outside those you have defined in AWS Organizations. Amazon Macie is integrated with other AWS services, such as [AWS Security Hub](#), to generate actionable security findings and provide an automatic and reactive action to the finding (Figure 5).

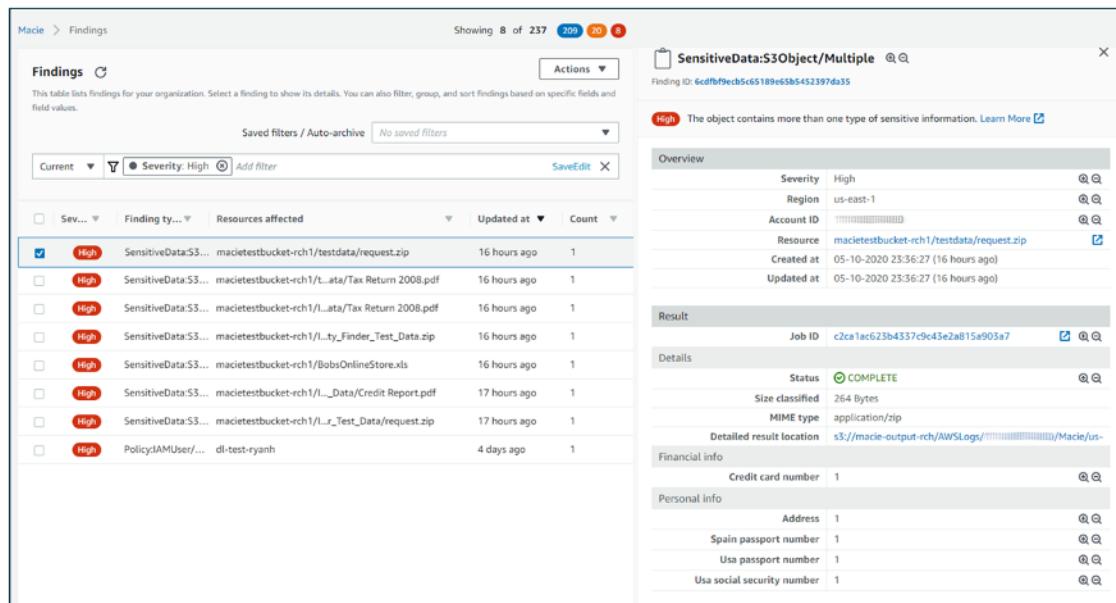


Figure 4 – Data inspections and finding example

Centralized Security Management

Many organizations have challenges related to visibility and centralized management of their environments. As your operational footprint grows, this challenge can be compounded unless you carefully consider your security designs. Lack of knowledge, combined with decentralized and uneven management of governance and security processes, can make your environment vulnerable.

AWS provides tools that help you to address some of the most challenging requirements for IT management and governance, and tools for supporting a data protection by design approach.

AWS Control Tower provides a method to set up and govern a new, secure, multi-account AWS environment. It automates the setup of a [landing zone](#), which is a multi-account environment that is based on best-practices blueprints, and enables governance using guardrails that you can choose from a pre-packaged list. Guardrails implement governance rules for security, compliance, and operations. AWS Control Tower provides identity management using AWS IAM Identity Center (successor to AWS Single Sign-On) (IAM Identity Center) default directory and enables cross-account audit using IAM Identity Center and IAM. It also centralizes logs coming from CloudTrail and AWS Config logs, which are stored in Amazon S3.

AWS Security Hub is another service that supports centralization and can improve visibility into an organization. Security Hub centralizes and prioritizes security and compliance findings from across AWS accounts and services, such as Amazon GuardDuty and [Amazon Inspector](#), and can be integrated with security software from third-party partners to help you analyze security trends and identify the highest priority security issues.

Amazon GuardDuty is an intelligent threat detection service that can help customers more accurately and easily monitor and protect their AWS accounts, workloads, and data stored in Amazon S3.

GuardDuty analyzes billions of events across your AWS accounts from several sources, including AWS CloudTrail Management Events, CloudTrail Amazon S3 Data Events, Amazon Virtual Private Cloud Flow Logs, and DNS logs. For example, it detects unusual API calls, suspicious outbound communications to known malicious IP addresses, or possible data theft using DNS queries as the transport mechanism. GuardDuty is able to provide more accurate findings by leveraging machine learning-powered threat intelligence and third-party security partners.

Amazon Inspector is an automated security assessment service that helps improve the security and compliance of applications deployed on Amazon EC2 instances. Amazon Inspector automatically assesses applications for exposure, vulnerabilities, and deviations from best practices. After performing an assessment, Amazon Inspector produces a detailed list of security findings prioritized by level of severity.

Amazon CloudWatch Events enables you to set up your AWS account to send events to other AWS accounts, or to become a receiver for events from other accounts or organizations. This mechanism can be very useful for implementing cross-account incident response scenarios, by taking timely corrective actions (for example, by calling a Lambda function, or running a command on Amazon EC2 instance) as necessary any time a security incident event occurs.

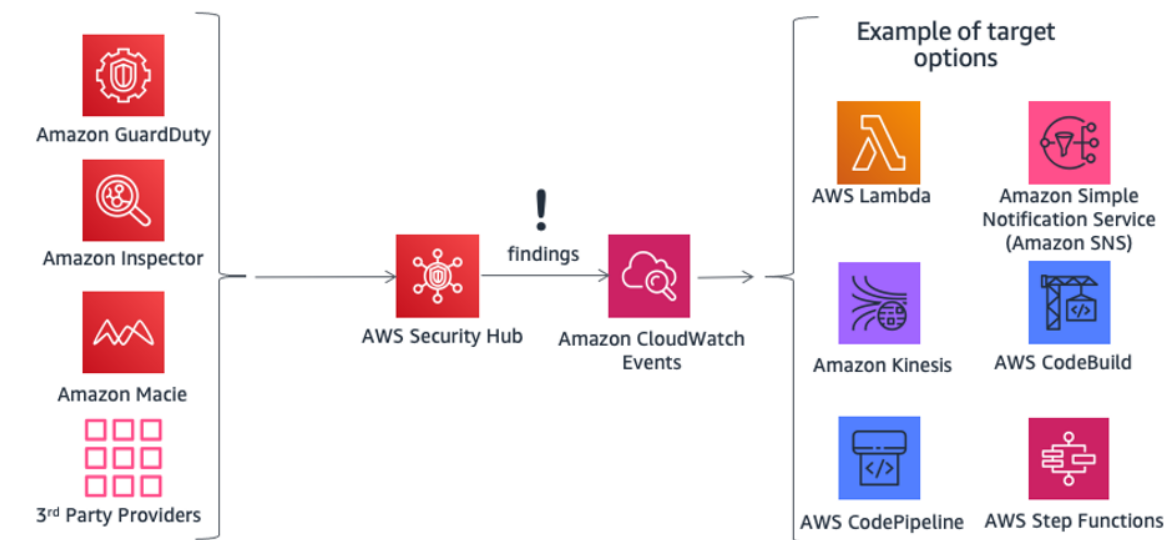


Figure 5 – Taking action with AWS Security Hub and Amazon CloudWatch Events

AWS Organizations helps you centrally manage and govern complex environments. It enables you to control access, compliance, and security in a multi-account environment. AWS Organizations supports **Service Control Policies (SCPs)**, which define the AWS service actions available to use with specific accounts or Organizational Units (OUs) within an organization.

AWS Systems Manager provides you visibility and control of your infrastructure on AWS. You can view operational data from multiple AWS services from a unified console and automate operational tasks across them. You can have information about recent API activities, resource configuration changes, operational alerts, software inventory, and patch compliance status. Using the integration with other AWS services, you can also take action on resources depending on your operational needs, to help make your environment in a compliance status.

For example, by integrating Amazon Inspector with AWS Systems Manager, security assessments are simplified and automated, because you can install Amazon Inspector agent automatically using Amazon Elastic Compute Cloud Systems Manager when an Amazon EC2 instance is launched. You can also perform automatic remediations for Amazon Inspector findings by using Amazon EC2 System Manager and Lambda functions.

Protecting your Data on AWS

Article 32 of the GDPR requires that organizations must "...implement appropriate technical and organisational measures to ensure a level of security appropriate to the risk, including ...the pseudonymisation and encryption of personal data[...]". In addition, organizations must safeguard against the unauthorized disclosure of or access to personal data."

Encryption reduces the risks associated with the storage of personal data because data is unreadable without the correct key. A thorough encryption strategy can help mitigate the impact of various security events, including some security breaches.

Encrypt Data at Rest

Encrypting data at rest is vital for regulatory compliance and data protection. It helps to ensure that sensitive data saved on disks is not readable by any user or application without a valid key. AWS provides multiple options for encryption at rest and encryption key management. For example, you can use the AWS Encryption SDK with a CMK created and managed in AWS KMS to encrypt arbitrary data.

Encrypted data can be securely stored at rest and can be decrypted only by a party with authorized access to the CMK. As a result, you get confidential envelope-encrypted data, policy mechanisms for authorization and authenticated encryption, and audit logging through AWS CloudTrail. Some of the AWS foundation services have built-in encryption at rest features, providing the option to encrypt data before it is written to non-volatile storage. For example, you can encrypt Amazon EBS volumes and configure Amazon S3 buckets for Server-Side Encryption (SSE) using AES-256 encryption. Amazon S3 also supports *client-side encryption*, which allows you to encrypt data before sending it to Amazon S3. AWS SDKs support client-side encryption to facilitate encryption and decryption operations of objects. Amazon RDS also supports Transparent Data Encryption (TDE).

It is possible to encrypt data on Linux Amazon EC2 instance stores by using built-in Linux libraries. This method encrypts files transparently, which protects confidential data. As a result, applications that process the data are unaware of the disk-level encryption.

You can use two methods to encrypt files on instance stores:

- **Disk-level encryption** — With this method, the entire disk, or a block within the disk, is encrypted using one or more encryption keys. Disk encryption operates below the file system level, is operating-system agnostic, and hides directory and file information, such as name and size. Encrypting File System, for example, is a Microsoft extension to the Windows NT operating system's New Technology File System (NTFS) that provides disk encryption.
- **File system-level encryption** — With this method, files and directories are encrypted, but not the entire disk or partition. File-system-level encryption operates on top of the file system and is portable across operating systems.

For Non-Volatile Memory express (NVMe) [SSD instance store volumes](#), *disk-level encryption* is the default option. Data in an NVMe instance storage is encrypted using an XTS-AES-256 block cipher implemented in a hardware module on the instance. The encryption keys are generated using the hardware module and are unique to each NVMe instance storage device. All encryption keys are destroyed when the instance is stopped or terminated and cannot be recovered. You cannot use your own encryption keys.

Encrypt Data in Transit

AWS strongly recommends encrypting data in transit from one system to another, including resources within and outside of AWS.

When you create an AWS account, a logically isolated section of the AWS Cloud—the Amazon Virtual Private Cloud (Amazon VPC)—is provisioned to it. There, you can launch AWS resources in a virtual network that you define. You have complete control over your virtual networking environment, including selecting your own IP address range, creation of subnets, and configuration of route tables and network gateways. You can also create a hardware Virtual Private Network (VPN) connection between your corporate datacenter and your Amazon VPC, so you can use the AWS Cloud as an extension of your corporate datacenter.

For protecting communication between your Amazon VPC and your corporate datacenter, you can select from [several VPN connectivity options](#), and choose one that best matches your needs. You can use the AWS Client VPN to enable secure access to your AWS resources using client-based VPN services. You can also use a third-party software VPN appliance available in the AWS Marketplace, which you can install on an Amazon EC2 instance in your Amazon VPC. Alternatively, you can create an IPsec VPN connection to protect the communication between your VPC and your remote network. To create a dedicated private connection from a remote network to your Amazon VPC, you can use [AWS Direct Connect](#). You can combine this connection with an AWS Site-to-Site VPN to create an IPsec-encrypted private connection.

AWS provides HTTPS endpoints using the TLS protocol for communication, which provides encryption in transit when you use AWS APIs. You can use the [AWS Certificate Manager \(ACM\)](#) service to generate, manage, and deploy the private and public certificates you use to establish encrypted transport between systems for your workloads. Elastic Load Balancing is integrated with ACM and is used to support HTTPS protocols. If your content is distributed through Amazon CloudFront, it supports encrypted endpoints.

Encryption Tools

AWS offers various highly scalable data encryption services, tools, and mechanisms to help protect your data stored and processed on AWS. For information about AWS Service functionality and privacy, refer to [Privacy Features of AWS Services](#).

Cryptographic services from AWS use a wide range of encryption and storage technologies that are designed to maintain integrity of your data at rest or in transit. AWS offers four primary tools for cryptographic operations.

- **AWS Key Management Service** (AWS KMS) is an AWS managed service that generates and manages both [root keys](#) and [data keys](#). AWS KMS is integrated [with many AWS services](#) to provide server-side encryption of data using AWS KMS keys from customer accounts. AWS KMS Hardware Security Modules (HSMs) are FIPS 140-2 Level 2 validated.
- **AWS CloudHSM** provides [HSMs](#) that are FIPS 140-2 Level 3 validated. They securely store a variety of your self-managed cryptographic keys, including KMS keys and data keys.
- **AWS Cryptographic Services and Tools**
 - **AWS Encryption SDK** provides a client-side encryption library for implementing encryption and decryption operations on *all* types of data.
 - **Amazon DynamoDB Encryption Client** provides a client-side encryption library for encrypting data tables before sending them to a database service, such as [Amazon DynamoDB](#).

AWS Key Management Service

AWS Key Management Service is a managed service that makes it easy for you to create and control the encryption keys used to encrypt your data, and uses Hardware Security Modules (HSMs) to protect

the security of your keys. AWS KMS is integrated with several other AWS services to help you protect the data you store with these services. AWS KMS is also integrated with AWS CloudTrail to provide you with logs of all your key usage for your regulatory and compliance needs.

You can easily create, import, and rotate keys, as well as define usage policies and audit usage from the AWS Management Console or by using the AWS SDK or AWS CLI.

The CMKs in AWS KMS, whether imported by you or created on your behalf by KMS, are stored in highly durable storage in an encrypted format to help ensure that they can be used when needed. You can choose to have KMS automatically rotate CMKs created in KMS once per year without having to re-encrypt data that has already been encrypted with your KMS key. You don't need to keep track of older versions of your CMKs because KMS keeps them available to automatically decrypt previously encrypted data.

For any CMK in AWS KMS, you can control who has access to those keys and which services they can be used with through a number of access controls, including grants, and key policy conditions within key policies or IAM policies. You can also import keys from your own key management infrastructure and use them in KMS.

For example, the following policy uses the `kms:ViaService` condition to allow a customer managed CMK to be used for the specified actions only when the request comes from Amazon EC2 or Amazon RDS in a specific Region (`us-west-2`) on behalf of a specific user (`ExampleUser`).

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Principal": {
        "AWS": "arn:aws:iam::111122223333:user/ExampleUser"
      }
      "Action": [
        "kms:Encrypt*",
        "kms:Decrypt",
        "kms:ReEncrypt*",
        "kms:GenerateDataKey*",
        "kms:CreateGrant",
        "kms:ListGrants",
        "kms:DescribeKey"
      ],
      "Resource": "*",
      "Condition": {
        "ForAnyValue:StringEquals": {
          "kms:ViaService": [
            "ec2.us-west-2.amazonaws.com",
            "rds.us-west-2.amazonaws.com"
          ]
        }
      }
    }
  ]
}
```

AWS Service Integration

AWS KMS has integrated with a number of AWS services— refer to the [KMS website](#) for a full list of integrated services. These integrations enable you to easily use AWS KMS CMKs to encrypt the data you store with these services. In addition to using a customer managed CMK, a number of the integrated services enable you to use an AWS-managed CMK that is created and managed for you automatically, but is only usable within the specific service that created it.

Audit Capabilities

[AWS CloudTrail](#) records each use of a key that you store in AWS KMS in a log file that is delivered to the Amazon S3 bucket that you specified in your configuration of CloudTrail. The information recorded includes details of the user, time, date, operation performed, and the key used.

Security

AWS KMS is designed to make sure that no one has access to your KMS keys. The service is built on systems that are designed to protect your KMS keys with extensive hardening techniques, such as never storing plaintext KMS keys on disk, not persisting them in memory, and limiting which systems can access hosts that use keys. All access to update software on the service is controlled by a multi-party access control that is audited and reviewed by an independent group within AWS.

For more information about AWS KMS, see the [AWS Key Management Service](#) whitepaper.

AWS CloudHSM

The [AWS CloudHSM](#) is a cloud-based hardware security module (HSM) that helps you meet corporate, contractual, and regulatory compliance requirements for data security by enabling you to generate and use your encryption keys on a FIPS 140-2 Level 3 validated hardware.

With AWS CloudHSM, you control the encryption keys and cryptographic operations performed by HSM.

AWS and AWS Marketplace partners offer a variety of solutions for protecting sensitive data within the AWS platform, but for applications and data subject to rigorous contractual or regulatory requirements for managing cryptographic keys, additional protection is sometimes necessary. Previously, the only option to store sensitive data (or the encryption keys protecting the sensitive data) may have been in on-premises datacenters. This might have prevented you from migrating these applications to the cloud, or significantly slowed their performance. With AWS CloudHSM, you can protect your encryption keys within HSM's designed and validated to government standards for secure key management. You can securely generate, store, and manage the cryptographic keys used for data encryption to make sure that only you can get access to them. AWS CloudHSM helps you comply with strict key management requirements without sacrificing application performance.

The AWS CloudHSM service works with Amazon VPC. AWS CloudHSM instances are provisioned inside your Amazon VPC with an IP address that you specify, which provides simple and private network connectivity to your Amazon EC2 instances. When you locate your HSM instances near your Amazon EC2 instances, you decrease network latency, which can improve application performance. AWS provides dedicated and exclusive (single tenant) access to HSM instances, which are isolated from other AWS customers. Available in multiple Regions and Availability Zones, AWS CloudHSM enables you to add secure and durable key storage to your applications.

Integration with AWS Services and Third-Party Applications

You can use CloudHSM with Amazon Redshift, Amazon RDS for Oracle, or third-party applications (such as SafeNet Virtual KeySecure) as your Root of Trust, Apache (SSL termination), or Microsoft SQL Server (transparent data encryption). You can also use AWS CloudHSM when you write your own applications and continue to use the standard cryptographic libraries, including PKCS#11, Java JCA/JCE, and Microsoft CAPI and CNG.

Audit Activities

If you need to track resource changes, or audit activities for security and compliance purposes, you can review the management API calls over the AWS CloudHSM made from your account using AWS CloudTrail. Additionally, you can audit operations on the HSM appliance using syslog or send syslog log messages to your own log collector.

AWS Cryptographic Services and Tools

AWS offers mechanisms that comply with a wide range of cryptographic security standards that you can use to implement best-practice encryption. The [AWS Encryption SDK](#) is a client-side encryption library, available in Java, Python, C, JavaScript, and a command line interface that supports Linux, macOS, and Windows. It offers advanced data protection features including secure, authenticated, symmetric key algorithm suites, such as 256-bit AES-GCM with key derivation and signing. Because it was specifically designed for applications that use Amazon DynamoDB, the [DynamoDB Encryption Client](#) enables users to protect their table data before it is sent to the database. It also verifies and decrypts data when it is retrieved. The client is available in Java and Python.

Linux DM-Crypt Infrastructure

Dm-crypt is a Linux kernel-level encryption mechanism that allows users to mount an encrypted file system. Mounting a file system is the process in which a file system is attached to a directory (mount point), which makes it available to the operating system. After mounting, all files in the file system are available to applications without any additional interaction. These files are, however, encrypted when stored on disk.

Device mapper is an infrastructure in the Linux 2.6 and 3.x kernel that provides a generic method to create virtual layers of block devices. The device mapper crypt target provides transparent encryption of block devices using the kernel crypto API. The [solution in this post](#) uses dm-crypt in conjunction with a disk-backed file system mapped to a logical volume by the Logical Volume Manager (LVM). LVM provides logical volume management for the Linux kernel.

Data Protection by Design and by Default

Any time a user or an application tries to use the AWS Management Console, the AWS API, or the AWS CLI, a request is sent to AWS. The AWS service receives the request and executes a set of several steps to determine whether to allow or deny the request, according to a specific [policy evaluation logic](#). Except for root credential requests, all requests on AWS are denied by default (the default *deny* policy is applied). This means that everything that is not explicitly allowed by the policy is denied. In the definition of policies and as a best practice, AWS suggests that you apply the [least privilege principle](#), which means that every component (such as users, modules, or services) must be able to access only the resources required to complete its tasks.

This approach aligns with Article 25 of the GDPR, which states that “the controller shall implement appropriate technical and organizational measures for ensuring that, by default, only personal data which are necessary for each specific purpose of the processing are processed”.

AWS also provides tools to implement *infrastructure as code*, which is a powerful mechanism for including security from the beginning of the design of an architecture. AWS CloudFormation provides a common language to describe and provision all infrastructure resources, including security policies and processes. With these tools and practices, security becomes part of your code and can be versioned, monitored, and modified (with a versioning system) according to the requirements of your organization. This enables *data protection by design*, because security processes and policies can be included in the definition of your architecture, and can also be continuously monitored by security measures in your organization.

How AWS Can Help

Table 1 – How AWS can help you navigate GDPR compliance

Area	Description	AWS Services and Tools
Strong Compliance Framework	Appropriate technical and organizational measures may need to include “the ability to ensure the ongoing confidentiality, integrity, availability, and resilience of the processing systems and services.”	<p>SOC 1 / SSAE 16 / ISAE 3402 (formerly SAS 70) / SOC 2 / SOC 3</p> <p>PCI DSS Level 1</p> <p>ISO 9001 / ISO 27001 / ISO 27017 / ISO 27018</p> <p>NIST FIPS 140-2</p> <p>Common Cloud Computing Controls Catalog (C5)</p>
Data Access Control	The controller “... shall implement appropriate technical and organizational measures for ensuring that, by default, only personal data that are necessary for each specific purpose of the processing are processed.”	<p>AWS Identity and Access Management (IAM)</p> <p>Amazon Cognito</p> <p>AWS Shield and AWS WAF</p> <p>AWS Resource Access Manager</p> <p>Amazon CloudFront</p> <p>AWS Organizations</p> <p>AWS CloudTrail</p>
Monitoring and Logging	<p>“Each controller and, where applicable, the controller’s representative, shall maintain a record of processing activities under its responsibility.”</p> <p>“...the controller and the processor shall implement appropriate technical and organizational</p>	<p>AWS Config</p> <p>Amazon CloudWatch</p> <p>AWS Control Tower</p> <p>Amazon GuardDuty</p> <p>Amazon Inspector</p> <p>Amazon Macie</p> <p>AWS Systems Manager</p> <p>AWS Security Hub</p> <p>AWS Tools and SDKs</p>

Area	Description	AWS Services and Tools
	measures to ensure a level of security appropriate to the risk [...]"	
Protecting your Data on AWS	Organizations must "implement appropriate technical and organizational measures to ensure a level of security appropriate to the risk, including the pseudonymisation and encryption of personal data."	AWS Certificate Manager
		AWS CloudHSM
		AWS Key Management Service

Contributors

Contributors to this document include:

- Tim Anderson, Technical Industry Specialist, Amazon Web Services
- Carmela Gambardella, Public Sector Solutions Architect, Amazon Web Services
- Giuseppe Russo, Security Assurance Manager, Amazon Web Services
- Marta Taggart, Senior Program Manager, Amazon Web Services
- Luca Iannario, Public Sector Solutions Architect, Amazon Web Services

Document history

To be notified about updates to this whitepaper, subscribe to the RSS feed.

update-history-change	update-history-description	update-history-date
Minor update (p. 26)	Fix non-inclusive language.	April 6, 2022
Minor update (p. 26)	Updated to include the addition of new AWS services and functionalities.	December 1, 2020
Initial publication (p. 26)	Whitepaper first published.	November 1, 2017

Note

To subscribe to RSS updates, you must have an RSS plug-in enabled for the browser that you are using.

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 current AWS 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. The responsibilities and liabilities of AWS 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.

© 2022 Amazon Web Services, Inc. or its affiliates. All rights reserved.