Navigating GDPR Compliance on AWS

October 2019

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.

Contents

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[1] (Regulation 2016/679 of the European Parliament and of the Council of April 27, 2016[2]) 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

The GDPR tries to create consistency across EU member states for 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 and Security experts work with customers across the world to answer their questions and help them run workloads in the cloud under the GDPR. These teams also review the responsibilities of AWS against the requirements of the GDPR.

> 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.

The Role of AWS Under the GDPR

Under the GDPR, AWS can be both a data processor and a data controller.

AWS as a Data Processor

When customers and AWS Solution Providers use AWS services to process personal data in their content, AWS acts as a data processor. Customers and AWS Solution Providers can use the controls available in AWS services, including security configuration controls, to process personal data. Under these circumstances, the customer or AWS Solution Providers 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, AWS stores account information as a data controller for account registration, administration, services access, customer contact, and support.

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.

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 securing the underlying infrastructure that supports the cloud, and customers 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. AWS operates, manages, and controls the infrastructure components, from the host operating system and virtualization layer, to the physical security of the facilities in which the service operates. Customers assume responsibility and management of the guest operating system (including updates and security patches), other associated application software, and the configuration of the security group firewall provided by AWS. For more information, see AWS Shared Responsibility Model.

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 Compliance enables customers to understand the robust controls in place at AWS to maintain security and data protection in the AWS Cloud. When systems are built in the AWS Cloud, compliance responsibilities are shared. By tying together governance-focused, audit-friendly service features with applicable compliance or audit standards, AWS Compliance enablers, such as AWS Config, AWS CloudTrail, AWS Identity and Access Management, Amazon GuardDuty, and AWS Security Hub, build on traditional programs, which helps customers to establish and operate in an AWS security-controlled environment. The IT infrastructure that AWS provides to its customers is designed and managed in alignment with security best practices and a variety of IT security standards, including:

- SOC 1/SSAE 16/ISAE 3402 (formerly SAS 70)
- SOC 2
- SOC 3
- FISMA, DIACAP, and FedRAMP
- DoD SRG
- PCI DSS Level 1
- ISO 9001 / ISO 27001
- ITAR
- FIPS 140-2
- MTCS Tier 3

In addition, the flexibility and control that the AWS platform provides enables customers to deploy solutions that meet several industry-specific standards[3].

AWS provides a wide range of information regarding its IT control environment to customers through whitepapers, reports, certifications, accreditations, and other third-party attestations. For more information, see the Amazon Web Services: Risk and Compliance whitepaper.

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. The C5 attestation can be used by AWS customers and their compliance advisors to understand the range of IT-Security assurance services that AWS offers, 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 that pertains 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

The GDPR contemplates the approval of codes of conduct to help controllers and processors demonstrate compliance under the regulation. One such code that is awaiting official approval from EU data protection authorities is the *CISPE Code of Conduct for Cloud Infrastructure Service Providers* (the *Code*)[4]. The Code gives customers comfort that their cloud provider uses appropriate data protection standards, which are consistent with the GDPR.

The following are a few key benefits of the Code:

- **Clarifies who is responsible for which aspects of data protection** – The Code explains the role of both the cloud provider and the customer under the GDPR, specifically within the context of cloud infrastructure services.
- **Defines the principles providers must follow** – The Code develops key principles in the GDPR about clear actions and commitments that providers should undertake to demonstrate their compliance with GDPR and help customers comply. Customers can use these concrete benefits in their own compliance and data protection strategies.
- **Gives customers the privacy and security information necessary to help them achieve their compliance goals** – The Code requires providers to be transparent about the steps they are taking to deliver on their privacy and security commitments. A few of these steps include the implementation of privacy and security safeguards, notification of data breaches, data deletion, and transparency of third-party sub-processing. All of these commitments are verified by third party, independent monitoring bodies. Customers can use this information to fully understand the high levels of security provided.

At the time of publication, AWS has registered Amazon EC2, Amazon Simple Storage Service (Amazon S3), Amazon Relational Database Service (Amazon RDS), AWS Identity and Access Management (IAM), AWS CloudTrail, and Amazon Elastic Block Store (Amazon EBS) as fully compliant with the Code. For more information, see CISPE Public Register. This provides AWS customers with additional assurances that they control their data in a safe, secure, and compliant environment when they use AWS. AWS compliance with the Code adds to the list of internationally recognized certifications and accreditations that AWS has achieved. This includes ISO 27001, ISO 27018, ISO 9001, SOC 1, SOC 2, SOC 3, PCI DSS Level 1, among others.

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. With IAM Roles, you can allow any users to perform specific tasks to assume it and leveraging on temporary credentials for the role session. You can use IAM roles to securely give applications that run in Amazon EC2 the credentials required to get access to other AWS resources, such as Amazon S3 buckets, and Amazon RDS or DynamoDB databases.

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 a few minutes to several 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 account and to individual user accounts. With multi-factor authentication (MFA) enabled, when you sign into an AWS website, 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.

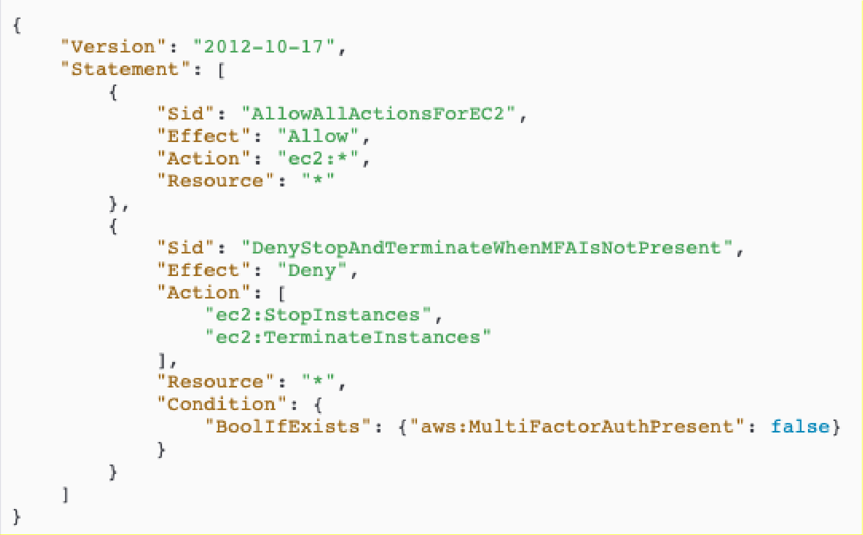

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Sid": "AllowAllActionsForEC2",
            "Effect": "Allow",
            "Action": "ec2:*",
            "Resource": "*"
        },
        {
            "Sid": "DenyStopAndTerminateWhenMFAIsNotPresent",
            "Effect":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

Figure 1 – Require MFA for specific Amazon EC2 API operations

Access to AWS Objects Resources

To implement granular access to your AWS objects, 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 Elastic Compute Cloud (Amazon EC2), Amazon Simple Storage Service (Amazon S3), Amazon 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 to access 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.

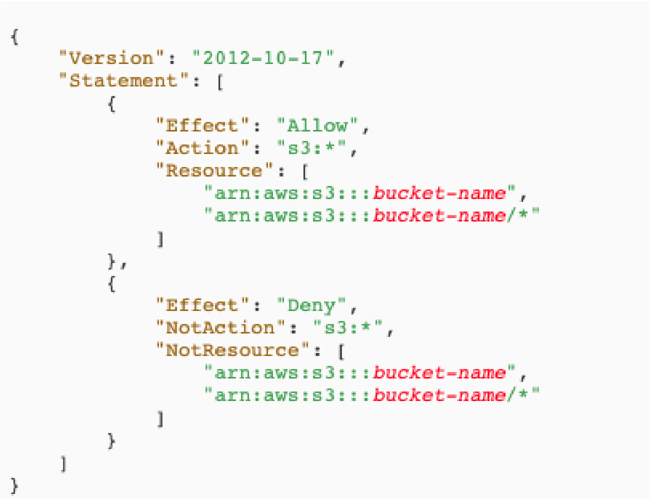

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Effect": "Allow",
            "Action": "s3:*",
            "Resource": [
                "arn:aws:s3:::bucket-name",
                "arn:aws:s3:::bucket-name/*"
            ]
        },
        {
            "Effect": "Deny",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

Figure 2 – Limit management to a specific Amazon S3 bucket

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.

Access to Operational & Configuration Data

You can use AWS Systems Manager to see and manage the operations of your AWS infrastructure. You can audit and enforce compliance to defined states. AWS Systems Manager Parameter Store can centrally manage data defining parameters. This enables you to implement granular access to parameter data, whether it is plain-text data (such as database strings) or secrets (such as passwords). You can provide this access control through customized permissions to users and resources (such as instances) for parameter access and to use the integration with IAM. For example, in a development environment, credentials are often hardcoded. Instead of hardcoding your credentials, you can use Parameter Store to save passwords and allow your developers to get access to the credentials with the AWS API get-parameter .

The following API snippet example shows the password retrieval *get-parameter*:

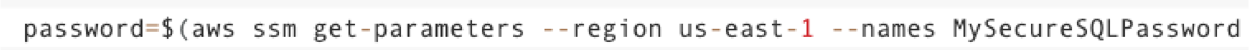

```
password=$(aws ssm get-parameters --region us-east-1 --names MySecureSQLPassword
```

Another available option for protecting secrets needed to access your applications, services, and IT resources is AWS Secrets Manager. The service enables you to easily rotate, manage, and retrieve database credentials, API keys, and other secrets throughout their lifecycle. Users and applications retrieve secrets with a call to Secrets Manager APIs, eliminating the need to hardcode sensitive information in plain text. Secrets Manager offers secret rotation with built-in integration for Amazon RDS, Amazon Redshift, and Amazon DocumentDB.

Geo-Restrictions

You can use geo-restrictions—also known as geoblocking—to prevent users in specific geographic locations from accessing content that you're distributing through an Amazon CloudFront web distribution.

There are two options for using geo-restrictions:

- **CloudFront geo-restriction feature** – Select this option to restrict access to all of the files that are associated with a CloudFront distribution, and to restrict access at the country level.
- **Third-party geolocation service** – Select this option to restrict access to a subset of the files that are associated with a distribution, or to restrict access at a finer level of granularity than the country level.

Beyond these two options, geo-limiting capabilities exist for newly launched Regions. While AWS Regions introduced before March 20, 2019 are enabled by default. Regions introduced after March 20, 2019, such as Asia Pacific (Hong Kong) and Middle East (Bahrain), 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 allows you to control whether users in your AWS account can access resources in that Region.[5]

Control Access to Web Applications and Mobile Apps

AWS provides service for managing data access control within their 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, you can see who accessed your resources and where the access originated (mobile app or web application). You can use this information to create security policies that allow or deny access to a resource based on the type of access origin (mobile app or web application).

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, as required by the GDPR. 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 the AWS resources in your AWS account. This includes how the resources are related to one another, and how they were previously configured, so that you can see how the configurations and relationships change over time.

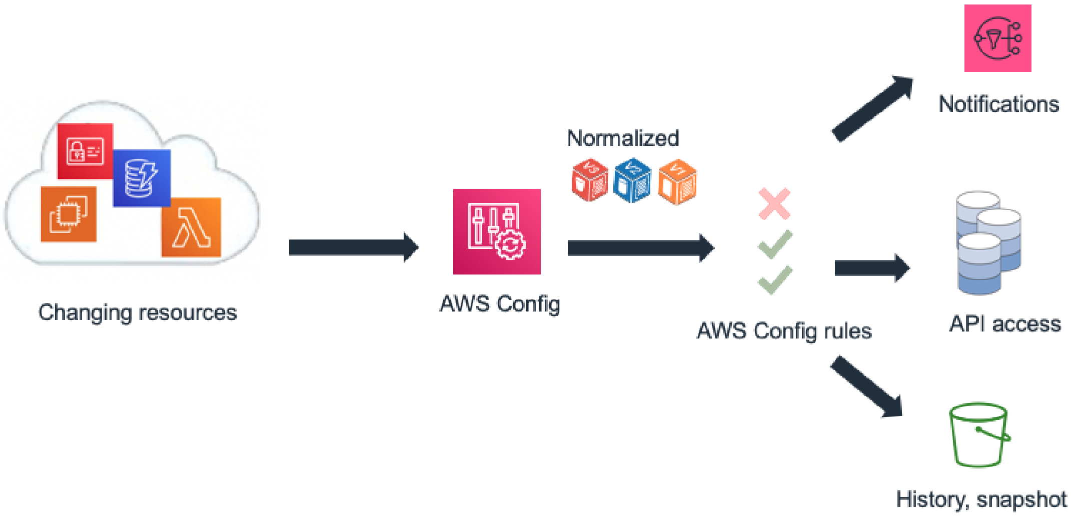

Figure 3 – Monitor configuration changes over time with AWS Config

An AWS resource is an entity that you can work with in AWS, such as an Amazon Elastic Compute Cloud (EC2) instance, an Amazon Elastic Block Store (EBS) volume, a security group, or an Amazon Virtual Private Cloud (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, you might want to find all resources that use a particular security group.

Compliance Auditing & Security Analytics with AWS CloudTrail

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. You can organize and store CloudTrail logs in an Amazon S3 bucket for auditing purposes or for troubleshooting activities.

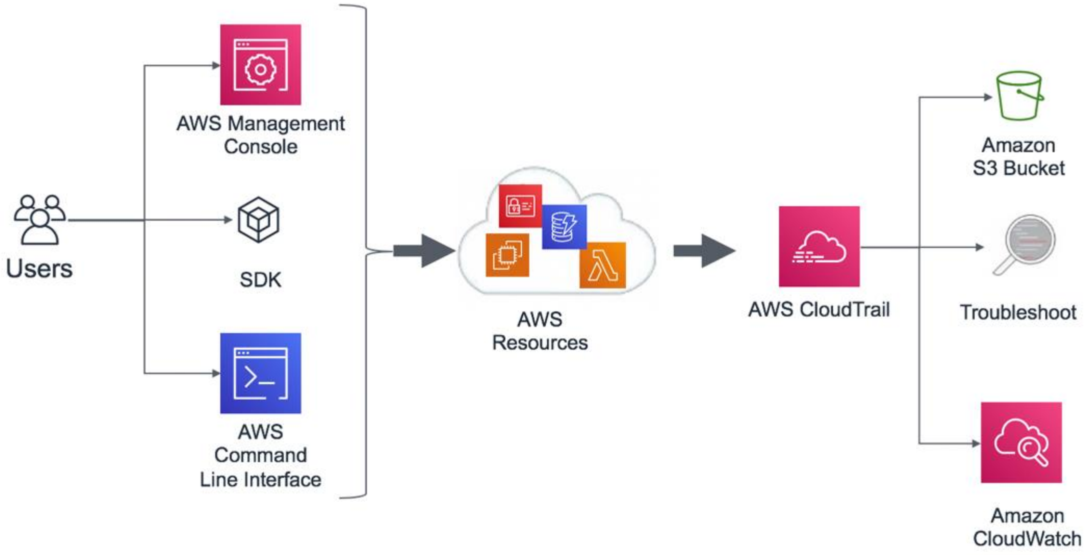

Figure 4 – Example architecture for compliance auditing and security analytics with AWS CloudTrail

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 (when it happened, which user performed the action, Amazon EC2 details, etc.) to the administrator. The following diagram shows the architecture of the event notification.

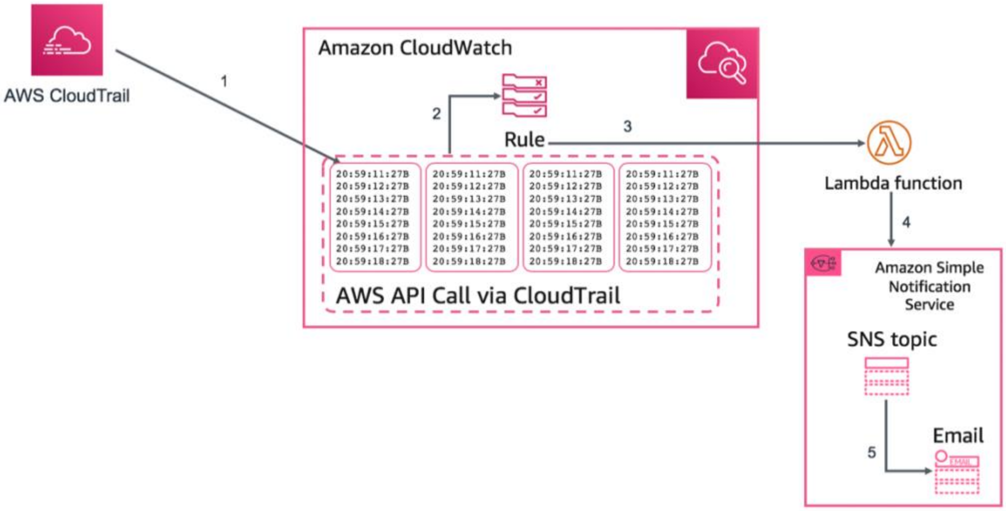

```
20:59:11:27B  20:59:11:27B  20:59:11:27B  20:59:11:27B
20:59:12:27B  20:59:12:27B  20:59:12:27B  20:59:12:27B
20:59:13:27B  20:59:13:27B  20:59:13:27B  20:59:13:27B
20:59:14:27B  20:59:14:27B  20:59:14:27B  20:59:14:27B
20:59:15:27B  20:59:15:27B  20:59:15:27B  20:59:15:27B
20:59:16:27B  20:59:16:27B  20:59:16:27B  20:59:16:27B
20:59:17:27B  20:59:17:27B  20:59:17:27B  20:59:17:27B
20:59:18:27B  20:59:18:27B  20:59:18:27B  20:59:18:27B
```

Figure 5 – Example of AWS CloudTrail event notification

Log Formats

When you enable logging, you can get detailed access logs for 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 S3 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 delivers access logs to your bucket on an hourly basis.

This information includes:

- 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 WAF functions in CloudFront

Logs are also a useful source of information for threat detection. Amazon GuardDuty analyzes logs from AWS CloudTrail, VPC Flow Logs, and AWS DNS, which enables you to continuously monitor your AWS accounts and workloads. This service uses machine learning, threat intelligence, and anomaly detection to deliver detailed and actionable alerts any time a malicious activity or an unauthorized behavior is recorded.

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, 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 an easy method to set up and govern a new, secure, multi-account AWS environment. It automates the setup of a landing zone[6] 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 Single Sign-On (SSO) default directory and enables cross-account audit using AWS SSO and AWS IAM. It also centralizes logs coming from Amazon 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, 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 CloudWatch Events enables you to set up your AWS account to send events to other AWS accounts, or 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 EC2 instance) as necessary any time a security incident event occurs.

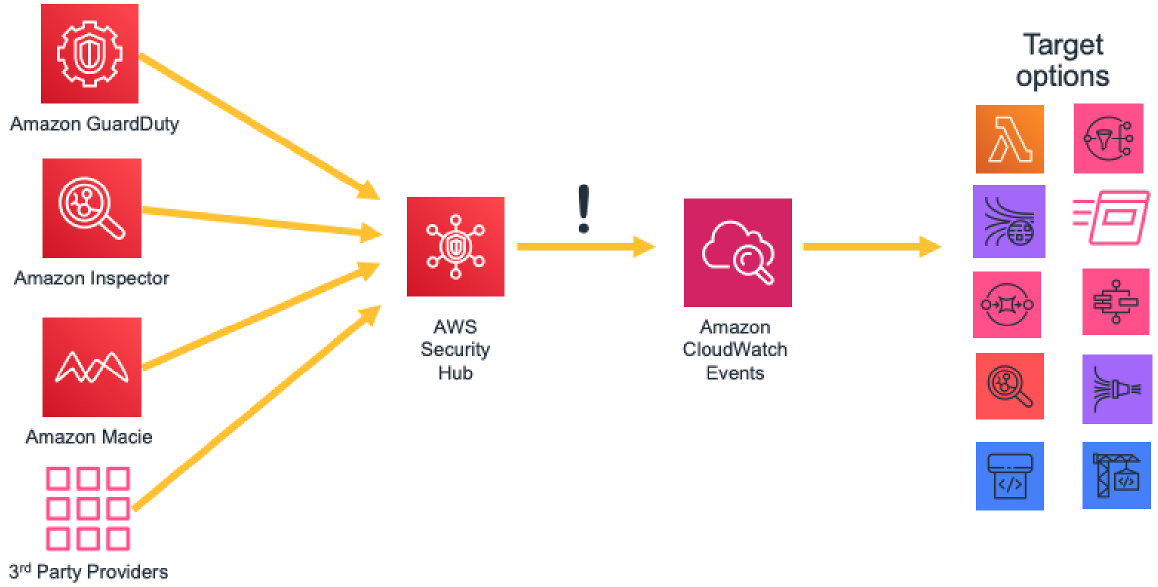

Figure 6 – Taking action with AWS Security Hub and Amazon CloudWatch Events

AWS Organizations helps you centrally manage and govern very complex environments. It enables you to control access, compliance, and security in a multi-account environment. AWS Organizations supports the Service control Policy (SCP), which defines the AWS service actions available to use with different accounts in an organization.

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 customer master key (CMK) created and managed in AWS Key Management Service (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 Elastic Block Store (Amazon EBS) volumes and configure Amazon Simple Storage Service (Amazon S3) buckets for server-side encryption (SSE) using AES-256 encryption. Amazon Relational Database Service (Amazon RDS) also supports Transparent Data Encryption (TDE).

Another method for encrypting data on Linux EC2 instance stores is 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. The first method is *disk encryption*, in which the entire disk, or 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.

The second method is *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, 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 is provisioned to it, the Amazon Virtual Private Cloud (Amazon VPC). 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, which you can install on an Amazon EC2 instance in your Amazon VPC. Or, 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 connection.

AWS provides HTTPS endpoints using the TLS (Transport Layer Security) 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. Amazon 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, see AWS Service Capabilities for Privacy Considerations[7].

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 master keys and data keys. AWS KMS is integrated with many AWS services to provide server-side encryption of data using KMS keys from customer accounts. 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 master 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 (AWS KMS) 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 Command Line Interface (AWS CLI).

The master keys in AWS KMS, whether imported by you or created on your behalf by AWS KMS and known as customer master keys (CMKs), 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 AWS KMS automatically rotate CMKs created in AWS KMS once per year without having to re-encrypt data that has already been encrypted with your master key. You don’t need to keep track of older versions of your CMKs because AWS KMS keeps them available to automatically decrypt previously encrypted data.

For any CMK in 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*).

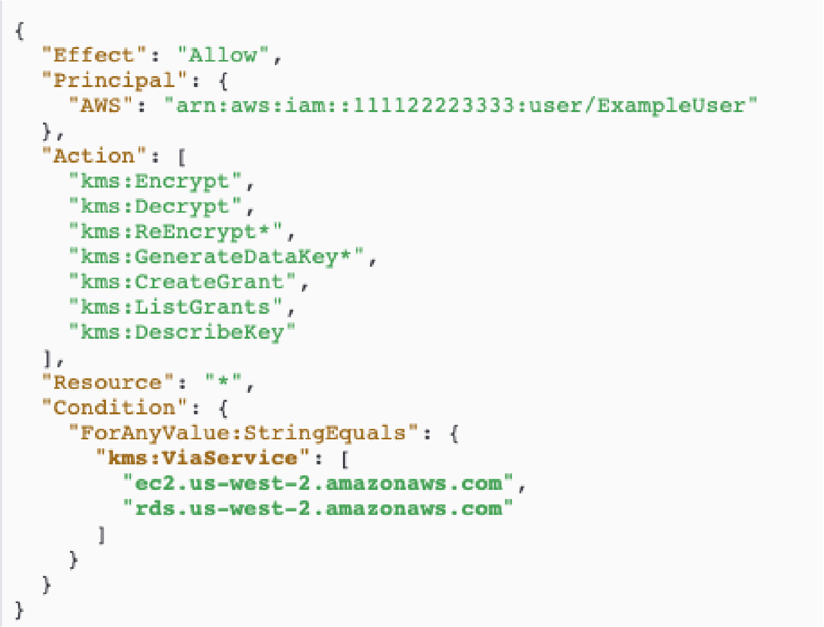

```
{
  "Effect": "Allow",
  "Principal": {
    "AWS": "arn:aws:iam::111122223333:user/ExampleUser"
  },
  "Action": [
    "kms:Encrypt",
    "kms:Decrypt",
    "kms:ReEncrypt*",
    "kms:GenerateDataKey*",
    "kms:CreateGrant",
    "kms:ListGrants",
    "kms:DescribeKey"
  ],
  "Resource": "*",
  "Condition":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

Figure 7 – Example of a policy for Amazon KMS

AWS Service Integration

AWS KMS has integrated with a number of AWS services (over fifty at the time of this writing). These integrations allow 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 allow 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

If AWS CloudTrail is enabled for your AWS account, each use of a key that you store in KMS is recorded in a log file that is delivered to the Amazon S3 bucket that you specified when you enabled AWS CloudTrail. The information recorded includes details of the user, time, date, and the key used.

Security

AWS KMS is designed to make sure that no one has access to your master keys. The service is built on systems that are designed to protect your master keys with extensive hardening techniques, such as never storing plaintext master 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 Amazon.

For more information about AWS KMS, see the AWS Key Management Service whitepaper.

AWS CloudHSM

The AWS CloudHSM service helps you meet corporate, contractual, and regulatory compliance requirements for data security by using dedicated Hardware Security Module (HSM) appliances in the AWS Cloud. With 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 HSMs 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 Virtual Private Cloud (Amazon VPC). 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 Elastic Compute Cloud (Amazon EC2) instances. When you locate your CloudHSM instances near your Amazon EC2 instances, you decrease network latency, which can improve application performance. AWS provides dedicated and exclusive (single tenant) access to CloudHSM instances, which are isolated from other AWS customers. Available in multiple Regions and Availability Zones, 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 Relational Database Service (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 CloudHSM when you write your own applications and continue to use the standard cryptographic libraries you're familiar with, 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 all of the CloudHSM API calls made from your account through 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[8] is a client-side encryption library, available in Java, Python, C, JavaScript, and a command line interface that supports Linux, macOS, and Windows. The AWS Encryption SDK 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[9] 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 & 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. 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 the *data protection by design* approach, 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

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."	SOC 1 / SSAE 16 / ISAE 3402 (formerly SAS 70) / SOC 2 / SOC 3 PCI DSS Level 1 ISO 9001 / ISO 27001 / ISO 27017 / ISO 27018 NIST FIPS 140-2 Common Cloud Computing Controls Catalog (C5)
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."	AWS Identity and Access Management (IAM) Amazon Cognito AWS WAF AWS CloudFormation AWS Systems Manager
Monitoring and Logging	"Each controller and, where applicable, the controller's representative, shall maintain a record of processing activities under its responsibility."	AWS CloudTrail AWS Config Amazon CloudWatch AWS Control Tower Amazon GuardDuty AWS Security Hub
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 Tools and SDKs 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

Document Revisions

Date	Description
October 2019	Updated to include the addition of new AWS services.
September 2018	Minor updates.
November 2017	First publication

Notes

[1] https://ec.europa.eu/info/law/law-topic/data-protection_en

[2] https://eur-lex.europa.eu/legal-content/EN/TXT/PDF/?uri=CELEX:32016R0679

[3] https://aws.amazon.com/compliance/programs/

[4] https://cispe.cloud/

[5] https://docs.aws.amazon.com/general/latest/gr/rande-manage.html

[6] https://aws.amazon.com/solutions/aws-landing-zone/

[7] https://aws.amazon.com/compliance/data-privacy/service-capabilities/

[8] https://docs.aws.amazon.com/crypto/latest/userguide/awscryp-service-encrypt.html

[9] https://docs.aws.amazon.com/crypto/latest/userguide/awscryp-service-ddb-client.html