Using AWS in the context of NCSC UK's Cloud Security Principles

October 2016

This paper has been archived

For the latest technical content, refer to the AWS Whitepapers & Guides page:
https://aws.amazon.com/whitepapers

Table of Contents

Abstract

This whitepaper is intended to assist organisations using Amazon Web Services (AWS) for United Kingdom (UK) OFFICIAL classified workloads in alignment with National Cyber Security Centre's (NCSC) Cloud Security Principles published under the Cloud Security Guidance. This document aims to help the reader understand:

- How AWS implements security processes and provides assurance over those processes for each of the Cloud Security Principles
- The role that the customer and AWS play in managing and securing content stored on AWS
- The way AWS services operate, including how customers can address security and risk management using AWS cloud services.

Scope

This whitepaper is based around typical questions asked by AWS customers when considering the implications of handling OFFICIAL information in relation to NCSC Cloud Security Principles. Our intention is to provide you with guidance that you can use to make an informed decision when performing risk assessments to help address common security requirements. This whitepaper is not legal advice for your specific use of AWS; we strongly encourage you to obtain appropriate compliance advice about your specific data protection and security requirements, as well as applicable laws relevant to your projects and datasets.

Considerations for public sector organisations

NCSC published the Cloud Security Guidance documents for public sector organisations that are considering the use of cloud services for handling OFFICIAL information on 23 April 2014. Under this guidance, HM Government information assets are currently classified into three categories: OFFICIAL, SECRET and TOP SECRET. Each information asset classification attracts a baseline set of security controls providing appropriate protection against typical threats.

NCSC Cloud Security Guidance includes a risk management approach to using cloud services, a summary of the Cloud Security Principles, and guidance on implementation of the Cloud Security Principles. Additionally, supporting guidance documents are included on recognised standards and definitions, separation requirements for cloud services and specific guidance on the measures that customers of Infrastructure as a Service (IaaS) offerings should consider taking.

This whitepaper provides guidance on how AWS aligns with Cloud Security Principles and the objectives of these principles as part of NCSC's Cloud Security Guidance. The legacy Impact Level accreditation scheme has been phased out and is no longer the mechanism used to describe the security properties of a system, including cloud services. Public sector

organisations are ultimately responsible for risk management decisions relating to the use of cloud services.

Gov.UK Digital Marketplace

Amazon Web Services currently provide the services listed on our UK G-Cloud page on the UK Government Digital Marketplace.

When using AWS services, customers maintain complete control over their content and are responsible for managing critical content security requirements, including:

- What content they choose to store on AWS
- Which AWS services are used with the content
- In what country that content is stored
- The format and structure of that content and whether it is masked, anonymised or encrypted
- Who has access to that content and how those access rights are granted, managed and revoked.

Because AWS customers retain control over their data, they also retain responsibilities relating to that content as part of the AWS shared responsibility model. This shared responsibility model is fundamental to understanding the respective roles of the customer and AWS in the context of the Cloud Security Principles.

Shared Responsibility Environment

Using AWS creates a shared responsibility model between customers and AWS. AWS operates, manages, and controls the components from the host operating system and virtualization layer down to the physical security of the facilities in which the services operate. In turn, customers assume responsibility for and management of the guest operating system

(including updates and security patches), other associated application software, as well as the configuration of the AWS-provided security group firewall. Customers should carefully consider the services they choose, as their responsibilities vary depending on the services they use, the integration of those services into their IT environments, and applicable laws and regulations. It is possible to enhance security and/or meet more stringent compliance requirements by leveraging technology such as host-based firewalls, host-based intrusion detection/ prevention, and encryption. AWS provides tools and information to assist customers in their efforts to account for and to validate that controls are operating effectively in their extended IT environment. More information can be found on the AWS Compliance center at http://aws.amazon.com/compliance.

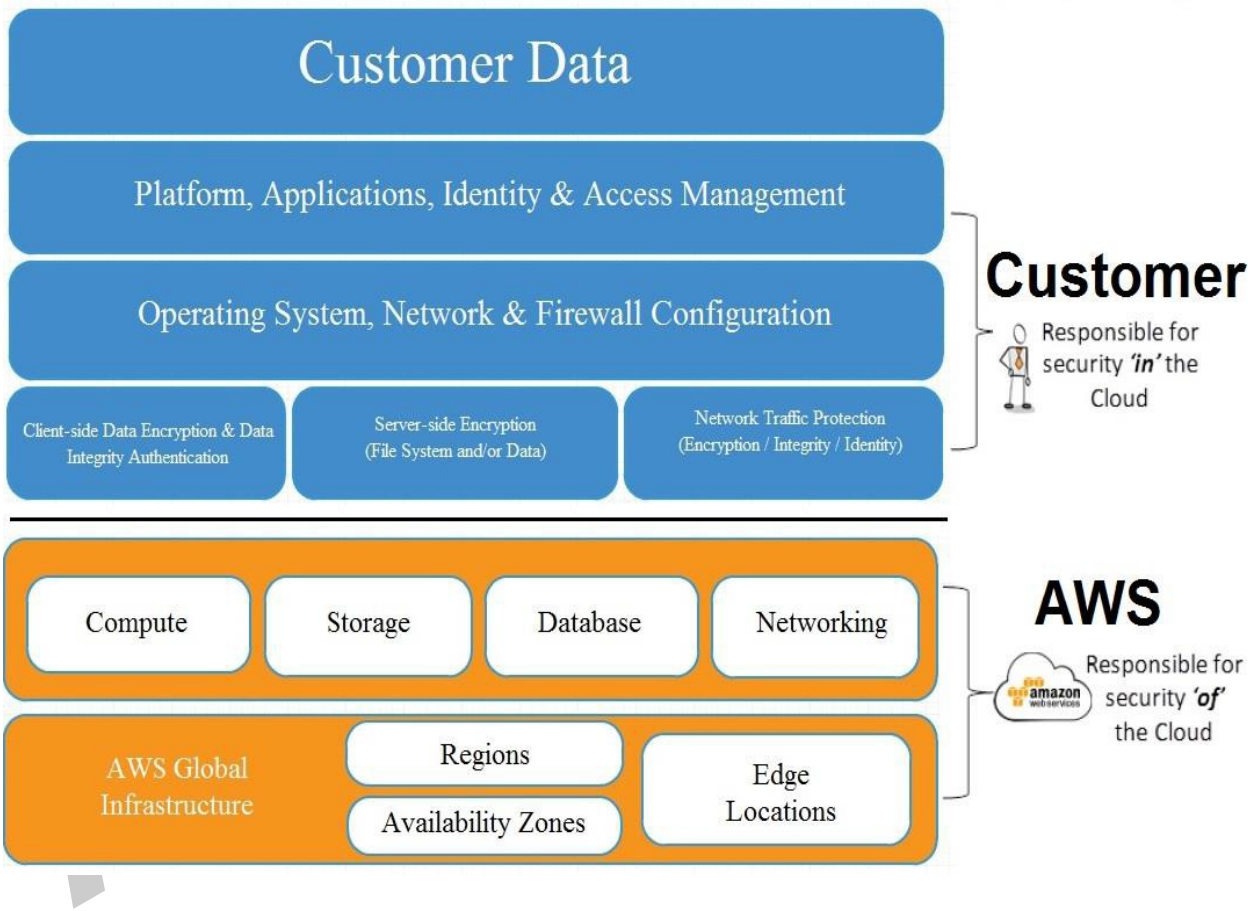

Implementing Cloud Security Principles in AWS

The Cloud Security Guidance published by NCSC lists 14 essential principles to consider when evaluating cloud services, and why these may be important to the public sector organisation. Cloud service users should decide which of the principles are important, and how much (if any) assurance the users require in the implementation of these principles.

The 14 Cloud Security Principles, their objectives and how AWS services can be used to implement these objectives are described with the related assurance approach.

Principle 1: Data in transit protection

Implementation approach

AWS uses various technologies to enable data in transit protection between the consumer and a service, within each service and between the services. Cloud infrastructure and applications often communicate over public links, such as the Internet, so it is important to protect data in transit when you run applications in the cloud. This involves protecting network traffic between clients and servers, and network traffic between servers. Further information on enabling network security for data protection is provided in the next section.

AWS Network Protection

The AWS network provides protection against network attacks. Procedures and mechanisms are in place to appropriately restrict unauthorized internal and external access to data, and access to customer data is appropriately segregated from other customers.

Examples include:

Distributed Denial of Service (DDoS) Attacks: AWS API endpoints are hosted on large, Internet-scale infrastructure and use proprietary

Data in Transit Protection

Consumer data transiting networks should be adequately protected against tampering (integrity) and eavesdropping (confidentiality). This should be achieved via a combination of:

- Network protection (denying your attacker access to intercept data)
- Encryption (denying your attacker the ability to read data).

Implementation objectives

Consumers should be sufficiently confident that:

- Data in transit is protected between the consumer's end user device and the service
- Data in transit is protected internally within the service
- Data in transit is protected between the service and other services (e.g. where Application Programming Interfaces (APIs) are exposed).

https://www.gov.uk/government/publications/implementing-the-cloud-security-principles/implementing-the-cloud-security-principles#principle-1-data-in-transit-protection

DDoS mitigation techniques. Additionally, AWS networks are multi-homed across a number of providers to achieve Internet access diversity.

Man in the Middle (MITM) Attacks: All of the AWS APIs are available via Secure Sockets Layer (SSL) protected endpoints, which provide server authentication. Amazon EC2 Amazon Machine Images (AMIs) automatically generate new Secure Shell (SSH) host keys on first boot and log them to the instance's console. Customers can then use the secure APIs to call the console and access the host keys before logging into the instance for the first time. Customers can use SSL for all of their interactions with AWS.

Internet Protocol (IP) Spoofing: The AWS-controlled, host-based firewall infrastructure will not permit an instance to send traffic with a source IP or Media Access Control (MAC) address other than its own.

Port Scanning: Unauthorized port scans by Amazon EC2 customers are a violation of the AWS Acceptable Use Policy. Violations of the AWS Acceptable Use Policy are taken seriously, and every reported violation is investigated. Customers can report suspected abuse via the contacts available on our website at: http://aws-portal.amazon.com/gp/aws/html-forms-controller/contactus/AWSAbuse. When unauthorized port scanning is detected by AWS, it is stopped and blocked. Port scans of Amazon EC2 instances are generally ineffective because, by default, all inbound ports on Amazon EC2 instances are closed and are only opened by the customer. Customers' strict management of security groups can further mitigate the threat of port scans. Customers can request permission to conduct scans of their cloud infrastructure as long as they are limited to the customer's instances and do not violate the AWS Acceptable Use Policy. Advance approval for these types of scans can be initiated by submitting a request via the AWS Vulnerability / Penetration Testing Request Form.

Customer Network Protection

Virtual Private Cloud (VPC): A VPC is an isolated portion of the AWS cloud within which customers can deploy Amazon EC2 instances into subnets that segment the VPC's IP address range (as designated by the customer) and isolate Amazon EC2 instances in one subnet from another. Amazon EC2 instances within a VPC are only accessible by a customer via an IPsec Virtual Private Network (VPN) connection that is established to the VPC.

IPsec VPN: an IPsec VPN connection connects a customer's VPC to another network designated by the customer. IPsec is a protocol suite for securing IP communications by authenticating and encrypting each IP packet of a data stream. Amazon VPC customers can create an IPsec VPN connection to their VPC by first establishing an Internet Key Exchange (IKE) security association between their Amazon VPC, VPN gateway, and another network gateway using a pre-shared key as the authenticator. Upon establishment, IKE negotiates an ephemeral key to secure future IKE messages. An IKE security association cannot be established unless there is complete agreement among the parameters, including SHA-1 authentication and AES 128-bit encryption. Next, using the IKE ephemeral key, keys are established between the VPN gateway and customer gateway to form an IPsec security

association. Traffic between gateways is encrypted and decrypted using this security association. IKE automatically rotates the ephemeral keys used to encrypt traffic within the IPsec security association on a regular basis to ensure confidentiality of communications.

API: Amazon VPC API calls are part of the Amazon EC2 WSDL. All API calls to create and delete VPCs, subnets, VPN gateways and IPsec VPN connections are all signed using an X.509 certificate and an associated private key or the customer's AWS Secret Access Key. Without access to the customer's Secret Access Key or X.509 certificate, Amazon EC2 API calls cannot be successfully made with that customer's key pair. In addition, API calls can be encrypted with SSL to maintain confidentiality.

AWS Encryption (Data in transit)

AWS supports both IPsec and SSL/TLS for protection of data in transit. IPsec is a protocol that extends the IP protocol stack, often in network infrastructure, and allows applications on upper layers to communicate securely without modification. SSL/TLS, on the other hand, operates at the session layer, and while there are third-party SSL/TLS wrappers, it often requires support at the application layer as well. For further details on AWS service specific data in transit security, please refer to the AWS Security Best Practices whitepaper.

Assurance approach

The data in transit protection principle and related processes within AWS services are subject to audit at least annually under ISO 27001:2013, ISO 9001, SOC 1, SOC 2, SOC 3 and PCI-DSS certification programs. These certifications, among others, are recognised by the European Union Agency for Network and Information Security (ENISA) under the Cloud Certification Schemes. The controls in relation to data in transit protection are validated independently at least annually under the certification programs. Based on the alternatives provided for selection within Cloud Security Principles guidance, AWS uses Service Provider Assertion in respect of region-specific requirements.

Principle 2: Asset protection and resilience

Asset protection and resilience

Consumer data, and the assets storing or processing it, should be protected against physical tampering, loss, damage or seizure.

https://www.gov.uk/government/publications/implementing-the-cloud-security-principles/implementing-the-cloud-security-principles#principle-2-asset-protection-and-resilience

Implementation approach

The AWS cloud is a globally available platform in which you can choose the geographic region in which your data is located. AWS data centers are built in clusters in various global regions. AWS calls

these data center clusters *Availability zones* (AZs). As of October 2016, AWS maintains 38 AZs organized into 14 regions globally. As an AWS customer, you are responsible for carefully selecting the Availability Zones where your systems will reside. You can choose to use one region, all regions, or any combination of regions using built-in features available within the AWS Management Console.

AWS regions and Availability Zones ensure that if you have location-specific requirements or regional data privacy policies, you can establish and maintain your private AWS environment in the appropriate location. ***You*** can choose to replicate and back up content in more than one region; AWS ***does not move*** customer data outside the region(s) you configure.

Availability Zones are designed for fault isolation. They are connected to multiple Internet Service Providers (ISPs) and different power grids. They are interconnected using high speed links, so applications can rely on Local Area Network (LAN) connectivity for communication between Availability Zones within the same region.

2.1 Physical location and legal jurisdiction

The locations, at which consumer data is stored, processed and managed from, must be identified so that organisations can understand the legal circumstances in which their data could be accessed without their consent.

Public sector organisations will also need to understand how data handling controls within the service are enforced, relative to UK legislation. Inappropriate protection of consumer data could result in legal and regulatory sanction or reputational damage.

Implementation objectives

Consumers should understand:

- What countries their data will be stored, processed and managed from and how this affects their compliance with relevant legislation
- Whether the legal jurisdiction(s) that the service provider operates within are acceptable to them

---

https://www.gov.uk/government/publications/implementing-the-cloud-security-principles/implementing-the-cloud-security-principles#principle-2-asset-protection-and-resilience

On March 6, 2015, the AWS data processing addendum, including the Model Clauses, was approved by the group of EU data protection authorities known as the Article 29 Working Party. This approval means that any AWS customer who requires the Model Clauses can now rely on the AWS data processing addendum as providing sufficient contractual commitments to enable international data flows in accordance with the Directive. For more detail on the approval from the Article 29 Working Party, please visit the Luxembourg Data Protection Authority webpage here:

http://www.cnpd.public.lu/en/actualites/international/2015/03/AWS/index.html.

AWS complies with Directive 95/46/EC of the European Parliament and of the Council of 24 October 1995 on the protection of individuals with regard to the processing of personal data and on the free movement of such data.

Most countries have data access laws, which purport to have extraterritorial application. An example of a US law with extra-territorial reach that is often mentioned in the context of cloud services is the US Patriot Act. The Patriot Act is not dissimilar to laws in many other developed nations that enable governments to obtain information with respect to investigations relating to international terrorism and other foreign intelligence issues. Any request for documents under the Patriot Act requires a court order demonstrating that the request complies with the law, including, for example, that the request is related to legitimate investigations.

Assurance approach

The legal jurisdiction sub-principle and related processes within AWS services are subject to audit at least annually under ISO 27001:2013 and PCI DSS Level 1 service-specific certification programs. These certifications are recognised by the European Union Agency for Network and Information Security (ENISA) under the Cloud Certification Schemes. The controls in relation to legal jurisdiction are validated independently at least annually under the certification programs. Based on the alternatives provided for selection within Cloud Security Principles guidance, AWS uses Service Provider Assertion in respect of region-specific requirements.

The physical location sub-principle and related processes are not validated independently within AWS compliance programs. Based on the alternatives provided for selection within Cloud Security Principles guidance, the controls in relation to physical location do not exist within the existing certification programs for them to be validated independently. Our ISO 27001:2013 and ISO 9001:2008 certifications list all the locations in scope of the independent annual audits. AWS uses Service Provider Assertion in respect of region-specific requirements.

2.2 Data centre security

The locations used to provide cloud services need physical protection against unauthorised access, tampering, theft or reconfiguration of systems. Inadequate protections may result in the disclosure, alteration or loss of data.

Implementation objectives

Consumers should be confident that the physical security measures employed by the provider are sufficient for their intended use of the service.

https://www.gov.uk/government/publications/implementing-the-cloud-security-principles/implementing-the-cloud-security-principles#principle-2-asset-protection-and-resilience

2.2 Data centre security

Implementation approach

Amazon has significant experience in securing, designing, constructing, and operating large-scale data centers. This experience has been applied to the AWS platform and infrastructure.

AWS provides data center physical access to approved employees and contractors who have a legitimate business need for such privileges. All individuals are required to present identification and are signed in. Visitors are escorted by authorised staff.

When an employee or contractor no longer requires these privileges, his or her access is promptly revoked, even if he or she continues to be an employee of Amazon or AWS. In addition, access is automatically revoked when an employee's record is terminated in Amazon's HR system.

Cardholder access to data center is reviewed quarterly. Cardholders marked for removal have their access revoked as part of the quarterly review.

Physical access is controlled both at the perimeter and at building ingress points by professional security staff utilizing video surveillance, intrusion detection systems and other electronic means. Authorized staff utilises multi-factor authentication mechanisms to access data center floors.

Assurance approach

The data centre security sub-principle and related processes within AWS services are subject to audit at least annually under ISO 27001:2013, AICPA SOC 1, SOC 2, SOC 3 and PCI-DSS certification programs. These certifications are recognised by ENISA under the Cloud Certification Schemes. The controls in relation to data centre security are validated independently at least annually under the certification programs.

2.3 Data at rest protection

Implementation approach

As AWS customers, you have access to various security and data protection features that allows sufficient confidence that data at rest is protected from unauthorised access. One of the widely used methods to protect data at rest in storage media is encryption. Within AWS there are several options for encrypting data, ranging from completely automated AWS encryption solutions (server-side) to manual, client-side options. Your decision to use a particular encryption model may be based on a variety of factors, including the AWS service(s) being used, your institutional policies, regulatory and business compliance requirements, your technical capability, specific requirements of the data, use certificate, and other factors. There are three different models for how you and/or AWS provide the encryption method and work with the key management infrastructure (KMI), as illustrated in the diagram below.

2.3 Data at rest protection

Consumer data should be protected when stored on any type of media or storage within a service to ensure that it is not accessible by local unauthorised parties. Without appropriate measures in place, data may be inadvertently disclosed on discarded, lost or stolen media.

Implementation objectives

Consumers should have sufficient confidence that storage media containing their data is protected from unauthorised access.

https://www.gov.uk/government/publications/implementing-the-cloud-security-principles/implementing-the-cloud-security-principles#principle-2-asset-protection-and-resilience

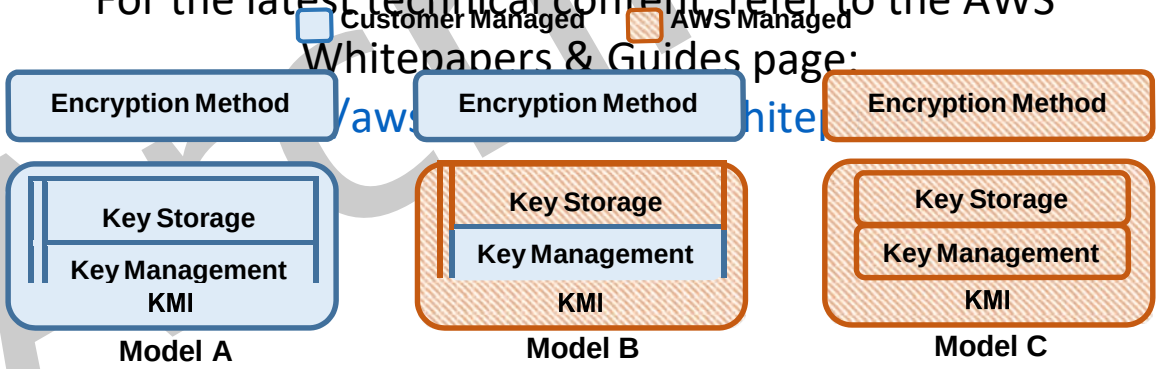

Customer manages the encryption method and entire KMI.

Customer manages the encryption method; AWS provides storage component of KMI while customer provides management layer of KMI.

AWS manages the encryption method and the entire KMI.

In addition to the client-side and server-side encryption features built into many AWS services, another common way to protect keys in a KMI is to use dedicated storage and data processing devices that perform cryptographic operations using keys on the devices. These devices, called hardware security modules (HSMs), typically provide tamper evidence or resistance to protect keys from unauthorized use. For customers who choose to use AWS encryption capabilities for controlled datasets, the AWS CloudHSM service is another encryption option within your AWS environment, giving you use of HSMs that are designed and validated to US government standards (NIST FIPS 140-2) for secure key management.

If you want to manage the keys that control encryption of data in AWS services, but don't want to manage the required KMI resources either within or external to AWS, you can leverage the AWS Key Management Service (KMS). 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 it uses HSMs to protect the security of your keys. AWS Key Management Service is integrated with other AWS services to help meet your regulatory and compliance needs. AWS KMS and other AWS services not listed on Digital Marketplace are available through our partner network. AWS KMS also allows you to implement key creation, rotation, and usage policies. AWS KMS is designed so that access to your master keys is restricted. 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 connect to the device. All access to update software on the service is controlled by multiple approved access capabilities and reviewed by an independent group within Amazon.

For more information about encryption options within the AWS environment, see Securing Data at Rest With Encryption, as well as the AWS CloudHSM product details page. To learn more about how AWS KMS works you can read the AWS Key Management Service Whitepaper.

To learn more about specific data at rest protection features in Amazon S3, Amazon EBS, Amazon RDS and Amazon Glacier, please refer to the AWS Security Best Practices Whitepaper.

For the implementation approach towards physical security controls to secure data at rest, please refer to the details described in Data Centre Security (Section 2.2) of this document.

Assurance approach

The data at rest protection sub-principle and related processes within AWS services are subject to audit at least annually under ISO 27001:2013, AICPA SOC 1, SOC 2, SOC 3 and PCI-DSS certification programs. These certifications are recognised by ENISA under the Cloud Certification Schemes. The controls in relation to data at rest protection are validated independently at least annually under the certification programs. Based on the alternatives provided for selection within Cloud Security Principles guidance, AWS uses Service Provider Assertion in respect of region-specific requirements.

2.4 Data sanitisation

Implementation approach

Helping to protect the confidentiality, integrity, and availability of our customers' systems and data is of the utmost importance to AWS, as is maintaining customer trust and confidence. AWS uses techniques described in industry-accepted standards to ensure that data is erased when resources are moved or re-provisioned, when they leave the service or when you request it to be erased.

AWS Data Erasure

Amazon EBS volumes are presented to you as raw unformatted block devices that have been wiped prior to being made available for use. Wiping occurs immediately before reuse as a mandatory process before re-provisioning. If you have procedures requiring that all data be wiped via a specific method, such as those detailed in DoD 5220.22-M ("National Industrial Security Program Operating Manual ") or NIST 800-88 ("Guidelines for Media Sanitization"), you have the ability to do so on Amazon EBS. You should conduct a specialized wipe procedure prior to deleting the volume for compliance with your established requirements.

Similarly, when deletion is requested for Amazon RDS database instance, the database instance is marked for deletion. An Amazon RDS automation sweeper deletes the instance from the Amazon RDS Storage System. At this point the instance is no longer accessible to the

2.4 Data sanitisation

The process of provisioning, migrating and de-provisioning resources should not result in unauthorised access to consumer data. Inadequate sanitisation of data could result in:

- Consumer data being retained by the service provider indefinitely
- Consumer data being accessible to other consumers of the service as resources are reused
- Consumer data being lost or disclosed on discarded, lost or stolen media.

Implementation objectives

Consumers should be sufficiently confident that:

- Their data is erased when resources are moved or re-provisioned, when they leave the service or when they request it to be erased
- Storage media which has held consumer data is sanitised or securely destroyed at the end of its life.

https://www.gov.uk/government/publications/implementing-the-cloud-security-principles/implementing-the-cloud-security-principles#principle-2-asset-protection-and-resilience

customer or AWS and unless the customer requested a 'delete with final snapshot copy', the instance cannot be restored and will not be listed by any of the tools or APIs.

AWS Secure Destruction

When a storage device has reached the end of its useful life, AWS procedures include a decommissioning process that is designed to prevent customer data from being exposed to unauthorized individuals. AWS uses the techniques detailed in DoD 5220.22-M ("National Industrial Security Program Operating Manual ") or NIST 800-88 ("Guidelines for Media Sanitization") to destroy data as part of the decommissioning process. All decommissioned magnetic storage devices are degaussed and physically destroyed in accordance with industry-standard practices.

Assurance approach

The data sanitisation sub-principle and related processes within AWS services are subject to audit at least annually under ISO 27001:2013 and PCI-DSS certification programs. These certifications are recognised by ENISA under the Cloud Certification Schemes. The controls in relation to data sanitisation are validated independently at least annually under the certification programs. Based on the alternatives provided for selection within Cloud Security Principles guidance, AWS uses Service Provider Assertion in respect of region-specific requirements.

2.5 Equipment disposal

Implementation approach

Helping to protect the confidentiality, integrity, and availability of our customers' systems and data is of the utmost importance to AWS, as is maintaining customer trust and confidence. AWS uses techniques described in industry-accepted standards to ensure that data is erased when resources are moved or re-provisioned, when they leave the service or when you request it to be erased.

When a storage device has reached the end of its useful life, AWS procedures include a decommissioning process that is designed to prevent customer data from being exposed to unauthorized individuals. AWS uses the techniques detailed in DoD 5220.22-M ("National Industrial Security Program Operating Manual ") or NIST 800-88 ("Guidelines for Media Sanitization") to destroy data as part of the decommissioning process. All decommissioned magnetic storage devices are degaussed and physically destroyed in accordance with industry-standard practices.

Assurance approach

The equipment protection sub-principle and related processes within AWS services are subject to audit at least annually under ISO 27001:2013 and PCI-DSS certification programs. These certifications are recognised by ENISA under the Cloud Certification Schemes. The controls in relation to equipment protection are validated independently at least annually under the certification programs. Based on the alternatives provided for selection within Cloud Security Principles guidance, AWS uses Service Provider Assertion in respect of region-specific requirements.

2.5 Equipment disposal

Once equipment used to deliver a service reaches the end of its useful life, it should be disposed of in a way that does not compromise the security of the service or consumer data stored in the service.

Implementation objectives

Consumers should be sufficiently confident that:

- All equipment potentially containing consumer data, credentials, or configuration information for the service is identified at the end of its life (or prior to being recycled)
- Any components containing sensitive data are sanitised, removed or destroyed as appropriate
- Accounts or credentials specific to redundant equipment are revoked to reduce their value to an attacker.

https://www.gov.uk/government/publications/implementing-the-cloud-security-principles/implementing-the-cloud-security-principles#principle-2-asset-protection-and-resilience

2.6 Physical resilience and availability

Implementation approach

The AWS Resiliency program encompasses the processes and procedures by which AWS identifies, responds to and recovers from a major event or incident within our environment. This program aims to provide you sufficient confidence that your business needs for availability commitment of the service including the ability to recover from outages are met. This program builds upon the traditional approach of addressing contingency management which incorporates elements of business continuity and disaster recovery plans and expands this to consider critical elements of proactive risk mitigation strategies such as engineering physically separate Availability Zones (AZs) and continuous infrastructure capacity planning.

AWS contingency plans and incident response playbooks are maintained and updated to reflect emerging continuity risks and lessons learned from past incidents. Plans are tested and updated through the due course of business (at least monthly) and the AWS Resiliency plan is reviewed and approved by senior leadership annually.

AWS has identified critical system components required to maintain the availability of the system and recover service in the event of outage. Critical system components (example: code bases) are backed up across multiple, isolated locations known as Availability Zones. Each Availability Zone runs on its own physically distinct, independent infrastructure, and is engineered to be highly reliable. Common points of failures like generators and cooling equipment are not shared across Availability Zones. Additionally, Availability Zones are physically separate, and designed such that even extremely uncommon disasters such as fires, tornados or flooding should only affect a single Availability Zone. AWS replicates critical system components across multiple Availability Zones and authoritative backups are maintained and monitored to ensure success replication.

AWS continuously monitors service usage to project infrastructure needs to support availability commitments and requirements. AWS maintains a capacity planning model to

2.6 Physical resilience and availability

Services have varying levels of resilience, which will affect their ability to operate normally in the event of failures, incidents or attacks. A service without guarantees of availability may become unavailable, potentially for prolonged periods, with attendant business impacts.

Implementation objectives

Consumers should be sufficiently confident that the availability commitment of the service, including their ability to recover from outages, meets their business needs.

https://www.gov.uk/government/publications/implementing-the-cloud-security-principles/implementing-the-cloud-security-principles#principle-2-asset-protection-and-resilience

assess infrastructure usage and demands at least monthly, and usually more frequently (e.g., weekly). In addition, the AWS capacity planning model supports the planning of future demands to acquire and implement additional resources based upon current resources and forecasted requirements.

Combined usage of Availability Zones and geographically distributed regions and numerous AWS services features provide customers with capabilities to design and architect resilient applications and platforms. AWS customers benefit from the aforementioned resiliency features when the architectures are designed towards multiple failure scenarios.

Assurance approach

The physical resilience and availability sub-principle and related processes are not validated independently within AWS compliance programs. Based on the alternatives provided for selection within Cloud Security Principles guidance, the controls in relation to physical resilience and availability do not exist within the existing certification programs for them to be validated independently. AWS publishes most up-to-the-minute information on service availability at status.aws.amazon.com. AWS uses Service Provider Assertion in respect of region-specific requirements.

Principle 3: Separation between consumers

Implementation approach

Helping to protect the confidentiality, integrity, and availability of our customers' systems and data is of the utmost importance to AWS, as is maintaining customer trust and confidence. Using multiple levels of security, AWS aims to provide you confidence that sufficient separation of data and management of the service exists from other consumers of the service.

Multiple Levels of Security

Security within Amazon EC2 is provided on multiple levels: the operating system (OS) of the host platform, the virtual instance OS or guest OS, firewalls, and signed API calls. Each of these items builds on the capabilities of the others. This helps prevent data contained within Amazon EC2 from being intercepted by unauthorized systems or users and to provide Amazon EC2 instances that are as secure as possible without sacrificing flexibility of configuration.

Packet sniffing by other tenants: Virtual instances are designed to prevent other instances running in promiscuous mode to receive or "sniff" traffic that is intended for a different virtual instance. While customers

Separation between consumers

Separation between different consumers of the service prevents one malicious or compromised consumer from affecting the service or data of another.

Some of the important characteristics which affect the strength and implementation of the separation controls are:

- The service model (e.g. IaaS, PaaS, SaaS) of the cloud service
- The deployment model (e.g. public, private or community cloud) of the cloud service
- The level of assurance available in the implementation of separation controls.

SaaS and PaaS services built upon IaaS offerings may inherit some of the separation properties of the underlying IaaS infrastructure.

Implementation objectives

Consumers should:

- Understand the types of consumers with which they share the service or platform
- Have confidence that the service provides sufficient separation of their data and service from other consumers of the service
- Have confidence that their management of the service is kept separate from other consumers (covered separately as part of Principle 9).

https://www.gov.uk/government/publications/implementing-the-cloud-security-principles/implementing-the-cloud-security-principles#principle-3-separation-between-consumers

can place interfaces into promiscuous mode, the hypervisor will not deliver any traffic to them that is not addressed to them. Even two virtual instances that are owned by the same customer located on the same physical host cannot listen to each other's traffic. While Amazon EC2 does provide protection against one customer inadvertently or maliciously attempting to view another's data, as a standard practice customers can encrypt sensitive traffic.

Customer instances have no access to raw disk devices, but instead are presented with virtualized disks. The AWS proprietary disk virtualization layer automatically erases every block of storage before making it available for use, which protects one customer's data from being unintentionally exposed to another. Customers can further protect their data using traditional filesystem encryption mechanisms, or, in the case of Elastic Block Store (EBS) volumes, enable AWS-managed disk encryption.

Firewall

Amazon EC2 provides a complete firewall solution, referred to as a Security Group; this mandatory inbound firewall is configured in a default deny-all mode and Amazon EC2 customers must explicitly open the ports needed to allow inbound traffic. The traffic may be restricted by any combination of protocol, port, and source (individual IP or Classless Inter-Domain Routing (CIDR) subnet, or another customer-defined security group). Customers launching instances in a Virtual Private Cloud (VPC) also have access to additional features such as restricting outbound traffic from an instance.

A VPC is an isolated portion of the AWS cloud within which customers can deploy Amazon EC2 instances into subnets that segment the VPC's IP address range (as designated by the customer) and isolate Amazon EC2 instances in one subnet from another. Amazon EC2 instances within a VPC are only accessible by a customer via an IPsec Virtual Private Network (VPN) connection that is established to the VPC.

Assurance approach

The separation between consumer's principle and related processes are not validated independently within AWS compliance programs. Based on the alternatives provided for selection within Cloud Security Principles guidance, the controls in relation to physical resilience and availability do not exist within the existing certification programs for them to be validated independently. AWS uses Service Provider Assertion in respect of region-specific requirements.

Principle 4: Governance framework

Implementation approach

AWS's Compliance and Security teams have established an information security framework and policies based on the Control Objectives for Information and related Technology (COBIT) framework and have effectively integrated the ISO 27001 certifiable framework based on ISO 27002 controls, American Institute of Certified Public Accountants (AICPA) Trust Services Principles, the PCI DSS v3.0, and the National Institute of Standards and Technology (NIST) Publication 800-53 Rev 4 (Recommended Security Controls for Federal Information Systems). AWS maintains the security policy, provides security training to employees, and performs application security reviews. These reviews assess the confidentiality, integrity, and availability of data, as well as conformance to the information security policy.

Governance framework

The service provider should have a security governance framework that coordinates and directs their overall approach to the management of the service and information within it.

Implementation objectives

The consumer has sufficient assurance that the governance framework and processes in place for the service are appropriate for their intended use of it.

https://www.gov.uk/government/publications/implementing-the-cloud-security-principles/implementing-the-cloud-security-principles#principle-4-governance-framework

As part of a globally accepted governance framework, AWS has achieved ISO 27001:2013 certification of our Information Security Management System (ISMS) covering AWS infrastructure, data centers, and many services. ISO 27001/27002 is a widely-adopted global security standard that sets out requirements and best practices for a systematic approach to managing company and customer information that's based on periodic risk assessments appropriate to ever-changing threat scenarios. In order to achieve the certification, a company must show it has a systematic and ongoing approach to managing information security risks that affect the confidentiality, integrity, and availability of company and customer information. This certification reinforces Amazon's commitment to providing significant information regarding our security controls and practices. AWS's ISO 27001:2013 certification includes all AWS data centers in all regions worldwide and AWS has established a formal program to maintain the certification.

AWS has an established information security organization managed by the AWS Security team and is led by the AWS Chief Information Security Officer (CISO). AWS Security establishes and maintains formal policies and procedures to delineate the minimum standards for logical access on the AWS platform and infrastructure hosts. The policies also identify functional responsibilities for the administration of logical access and security. Where applicable, AWS

Security leverages the information system framework and policies established and maintained by Amazon Corporate Information Security.

The aforementioned processes aim to provide you sufficient confidence that the governance framework and processes in place for the AWS services are appropriate for their intended use of it.

Assurance approach

The governance framework principle and related processes within AWS services are subject to audit at least annually under ISO 27001:2013, AICPA SOC 1, SOC 2, SOC 3 and PCI-DSS certification programs. These certifications are recognised by ENISA under the Cloud Certification Schemes. The controls in relation to governance framework are validated independently at least annually under the certification programs.

Principle 5: Operational security

5.1 Configuration and change management

Implementation approach

Operational security

The service provider should have processes and procedures in place to ensure the operational security of the service. The service will need to be operated and managed securely in order to impede, detect or prevent attacks against it.

https://www.gov.uk/government/publications/implementing-the-cloud-security-principles/implementing-the-cloud-security-principles#principle-5-operational-security

Software

AWS applies a systematic approach to managing change so that changes to customer impacting services are reviewed, tested, approved and well communicated.

Change management (CM) processes are based on Amazon change management guidelines and tailored to the specifics of each AWS service. These processes are documented and communicated to the necessary personnel by service team management.

The goal of AWS' change management processes is to prevent unintended service disruptions and maintain the integrity of service to the customer. Change details are documented in Amazon's CM workflow tool or another change management or deployment tool. Changes deployed into production environments are:

- Reviewed: peer reviews of the technical aspects of a change
- Tested: when applied will behave as expected and not adversely impact performance
- Approved: to provide appropriate oversight and understanding of business impact from service owners (management).

Changes are typically pushed into production in a phased deployment starting with lowest impact sites. Deployments are closely monitored so impact can be evaluated. Service owners have a number of configurable metrics that measure the health of the service's upstream dependencies. These metrics are closely monitored with thresholds and alarming in place (e.g., latency, availability, fatal errors, CPU utilization, etc.). Rollback procedures are documented in the Change Management (CM) ticket or other change management tool.

When possible, changes are scheduled during regular change windows. Emergency changes to production systems that require deviations from standard change management procedures are associated with an incident and are logged and approved as appropriate.

Infrastructure

AWS internally developed configuration management software is installed when new hardware is provisioned. These tools are run on all hosts to validate that they are configured and software is installed in a standard manner based on host classes and updated regularly. Only approved Systems Engineers and additional parties authorized through a permissions service may log in to the central configuration management servers.

Emergency, non-routine and other configuration changes to existing AWS infrastructure are authorized, logged, tested, approved and documented in accordance with industry norms for similar systems. Updates to AWS infrastructure are done in such a manner that in the vast majority of cases they will not impact the customer and their service use. AWS communicates with customers, either via email, or through the AWS Service Health Dashboard (http://status.aws.amazon.com) when service use may be adversely affected.

Assurance approach

The configuration and change management sub-principle and related processes within AWS services are subject to audit at least annually under ISO 27001:2013, AICPA SOC 1, SOC 2, SOC 3 and PCI-DSS certification programs. These certifications are recognised by ENISA under the Cloud Certification Schemes. The controls in relation to configuration and change management are validated independently at least annually under the certification programs.

5.2 Vulnerability management

Implementation approach

Amazon Web Services is responsible for protecting the global infrastructure that runs all of the services offered in the AWS cloud. Protecting this infrastructure is AWS's number one priority. AWS Security regularly scans all Internet facing service endpoint IP addresses for vulnerabilities (these scans do not include customer instances). AWS Security notifies the appropriate parties to remediate many identified vulnerabilities. In addition, external vulnerability threat assessments are performed regularly by independent security firms. Findings and recommendations resulting from these assessments are categorized and delivered to AWS leadership. These scans are done in a manner for the health and viability of the underlying AWS infrastructure and are not meant to replace the customer's own vulnerability scans required to meet their specific compliance requirements. Customers can request permission to conduct scans of their cloud infrastructure as long as they are limited to the customer's instances and do not violate the AWS Acceptable Use Policy. Advance approval for these types of scans can be

5.2 Vulnerability management

Occasionally, vulnerabilities will be discovered which, if left unmitigated, will pose an unacceptable risk to the service. Robust vulnerability management processes are required to identify, triage and mitigate vulnerabilities. Services which do not have effective vulnerability management processes will quickly become vulnerable to attack, leaving them at risk of exploitation using publicly known methods and tools.

Implementation objectives

Consumers should have confidence that:

- Potential new threats, vulnerabilities or exploitation techniques which could affect the service are assessed and corrective action is taken
- Relevant sources of information relating to threat, vulnerability and exploitation technique information are monitored by the service provider
- The severity of threats and vulnerabilities are considered within the context of the service and this information is used to prioritise implementation of mitigations
- Known vulnerabilities within the service are tracked until suitable mitigations have been deployed through a suitable change management process
- Service provider timescales for implementing mitigations to vulnerabilities found within the service are made available to them.

https://www.gov.uk/government/publications/implementing-the-cloud-security-principles/implementing-the-cloud-security-principles#principle-5-operational-security

initiated by submitting a request via the AWS Vulnerability / Penetration Testing Request Form.

In addition, the AWS control environment is subject to regular internal and external risk assessments. AWS engages with external certifying bodies and independent auditors to review and test the AWS overall control environment.

Assurance approach

The vulnerability management sub-principle and related processes within AWS services are subject to audit at least annually under ISO 27001:2013, AICPA SOC 1, SOC 2, SOC 3 and PCI-DSS certification programs. These certifications are recognised by ENISA under the Cloud Certification Schemes. The controls in relation to vulnerability management are validated independently at least annually under the certification programs.

Protective monitoring

Implementation approach

Systems within AWS are extensively instrumented to monitor key operational and security metrics. Alarms are configured to automatically notify operations and management personnel when early warning thresholds are crossed on key metrics. When a threshold is crossed, the AWS incident response process is initiated. The Amazon Incident Response team employs industry-standard diagnostic procedures to drive resolution during business-impacting events. Staff operates 24x7x365 coverage to detect incidents and manage the impact to resolution.

AWS security monitoring tools help identify several types of denial of service (DoS) attacks, including distributed, flooding, and software/logic attacks. When DoS attacks are identified, the AWS incident response process is initiated. In addition to the DoS prevention tools, redundant

5.3 Protective monitoring

Effective protective monitoring allows a service provider to detect and respond to attempted and successful attacks, misuse and malfunction. A service which does not effectively monitor for attacks and misuse will be unlikely to detect attacks (both successful and unsuccessful) and will be unable to quickly respond to potential compromises of consumer environments and data.

Implementation objectives

Consumers should have confidence that:

- Events generated in service components required to support effective identification of suspicious activity are collected and fed into an analysis system
- Effective analysis systems are in place to identify and prioritise indications of potential malicious activity.

https://www.gov.uk/government/publications/implementing-the-cloud-security-principles/implementing-the-cloud-security-principles#principle-5-operational-security

telecommunication providers at each region, as well as additional capacity protect against the possibility of DoS attacks.

Assurance approach

The protective monitoring sub-principle and related processes within AWS services are subject to audit at least annually under ISO 27001:2013, AICPA SOC 1, SOC 2, SOC 3 and PCI-DSS certification programs. These certifications are recognised by ENISA under the Cloud Certification Schemes. The controls in relation to protective monitoring are validated independently at least annually under the certification programs.

Incident management

Implementation approach

AWS has implemented a formal, documented incident response policy and program. The policy addresses purpose, scope, roles, responsibilities, and management commitment.

AWS utilizes a three-phased approach to manage incidents:

1. Activation and Notification Phase: Incidents for AWS begin with the detection of an event. This can come from several sources including:
   a. Metrics and alarms - AWS maintains an exceptional situational awareness capability, most issues are rapidly detected from 24x7x365 monitoring and alarming of real time metrics and service dashboards. The majority of incidents are detected in this manner. AWS

5.4 Incident management

An incident management process allows a service provider to respond to a wide range of unexpected events that affect the delivery of the service to consumers. Unless carefully pre-planned incident management processes are in place, poor decisions are likely to be made when incidents do occur.

Implementation objectives

Consumers should have confidence that:

- Incident management processes are in place for the service and are enacted in response to security incidents
- Pre-defined processes are in place for responding to common types of incident and attack
- A defined process and contact route exists for reporting of security incidents by consumers and external entities
- Security incidents of relevance to them will be reported to them in acceptable timescales and format.

https://www.gov.uk/government/publications/implementing-the-cloud-security-principles/implementing-the-cloud-security-principles#principle-5-operational-security

utilizes early indicator alarms to proactively identify issues that may ultimately impact Customers.

b. Trouble ticket entered by an AWS employee
c. Calls to the 24X7X365 technical support hotline. If the event meets incident criteria, then the relevant on-call support engineer will start an engagement utilizing AWS Event Management Tool system to start the engagement and page relevant program resolvers (e.g. Security team). The resolvers will perform an analysis of the incident to determine if additional resolvers should be engaged and to determine the approximate root cause.

2. Recovery Phase - the relevant resolvers will perform break fix to address the incident. Once troubleshooting, break fix and affected components are addressed, the call leader will assign next steps in terms of follow-up documentation and follow-up actions and end the call engagement.
3. Reconstitution Phase - Once the relevant fix activities are complete the call leader will declare that the recovery phase is complete. Post mortem and deep root cause analysis of the incident will be assigned to the relevant team. The results of the post mortem will be reviewed by relevant senior management, and relevant actions such as design changes etc. will be captured in a Correction of Errors (COE) document and tracked to completion.

In addition to the internal communication mechanisms detailed above, AWS has also implemented various methods of external communication to support its customer base and community. Mechanisms are in place to allow the customer support team to be notified of operational issues that impact the customer experience. A "Service Health Dashboard" is available and maintained by the customer support team to alert customers to any issues that may be of broad impact.

Assurance approach

The incident management sub-principle and related processes within AWS services are subject to audit at least annually under ISO 27001:2013, AICPA SOC 1, SOC 2, SOC 3 and PCI-DSS certification programs. These certifications are recognised by ENISA under the Cloud Certification Schemes. The controls in relation to incident management are validated independently at least annually under the certification programs.

Principle 6: Personnel security

Implementation approach

To ensure you are confident with the level of personnel checks, AWS conducts criminal background checks, as permitted by applicable law, as part of pre-employment screening practices for employees commensurate with the employee's position and level of access to AWS facilities.

As part of the on-boarding process, all personnel supporting AWS systems and devices sign a non-disclosure agreement prior to being granted access. Additionally, as part of orientation, personnel are required to read and accept the Acceptable Use Policy and the Amazon Code of Business Conduct and Ethics (Code of Conduct) Policy.

AWS maintains employee training programs to promote awareness of AWS information security requirements. Every employee is provided with the Company's Code of Business Conduct and Ethics and completes periodic Information Security training, which requires an acknowledgement to complete. Compliance audits are periodically performed to validate that employees understand and follow the established policies.

Personnel security

Consumers should be content with the level of security screening conducted on service provider staff with access to their information or with ability to affect their service.

Implementation objectives

Service provider staff should be subject to personnel security screening and security education for their role.

Personnel within a cloud service provider with access to consumer data and systems need to be trustworthy. Service providers need to make clear how they screen and manage personnel within any privileged roles. Personnel in those roles should understand their responsibilities and receive regular security training. More thorough screening, supported by adequate training, reduces the likelihood of accidental or malicious compromise of consumer data by service provider personnel.

https://www.gov.uk/government/publications/implementing-the-cloud-security-principles/implementing-the-cloud-security-principles#principle-6-personnel-security

Assurance approach

The personnel security principle and related processes within AWS services are subject to audit at least annually under ISO 27001:2013, AICPA SOC 1, SOC 2, SOC 3 and PCI-DSS certification programs. These certifications are recognised by ENISA under the Cloud Certification Schemes. The controls in relation to personnel security are validated independently at least annually under the certification programs. Based on the alternatives provided for

selection within Cloud Security Principles guidance, AWS uses Service Provider Assertion in respect of region-specific requirements.

Principle 7: Secure development

Implementation approach

AWS' development process follows secure software development best practices, which include formal design reviews by the AWS Security Team, threat modeling, and completion of a risk assessment. Static code analysis tools are run as a part of the standard build process, and all deployed software undergoes recurring penetration testing performed by carefully selected industry experts. Our security risk assessment reviews begin during the design phase and the engagement lasts through launch to ongoing operations.

In addition, refer to ISO 27001:2013 standard, Annex A, domain 12.5 for further details on AWS and our ISO 27001 certification standard.

Secure development

Services should be designed and developed to identify and mitigate threats to their security.

Services which are not designed securely may be vulnerable to security issues which could compromise consumer data, cause loss of service or enable other malicious activity.

Implementation objectives

Consumers should be content with the level of security screening conducted on service provider staff with access to their information or with ability to affect their service.

https://www.gov.uk/government/publications/implementing-the-cloud-security-principles/implementing-the-cloud-security-principles#principle-7-secure-development

Assurance approach

The secure development principle and related processes within AWS services are subject to audit at least annually under ISO 27001:2013, AICPA SOC 1, SOC 2, SOC 3 and PCI-DSS certification programs. These certifications are recognised by ENISA under the Cloud Certification Schemes. The controls in relation to secure development are validated independently at least annually under the certification programs.

Principle 8: Supply chain security

Implementation approach

In alignment with ISO 27001 standards, AWS hardware assets are assigned an owner and tracked and monitored by the AWS personnel with AWS proprietary inventory management tools. AWS procurement and supply chain teams maintain relationships with all AWS suppliers.

Personnel security requirements for third-party providers supporting AWS systems and devices are established in a Mutual Non-Disclosure Agreement between AWS' parent organization, Amazon.com, and the respective third-party provider. The Amazon Legal Counsel and the AWS Procurement team define AWS third party provider personnel security requirements in contract agreements with the third party provider. All persons working with AWS information must at a minimum meet the screening process for pre-employment background checks and sign a Non-Disclosure Agreement (NDA) prior to being granted access to AWS information.

Refer to ISO 27001 standards; Annex A, domain 7.1 for additional details. AWS has been validated and certified by an independent auditor to confirm alignment with the ISO 27001 certification standard.

Assurance approach

The supply chain security principle and related processes within AWS services are subject to audit at least annually

Supply chain security

The service provider should ensure that its supply chain satisfactorily supports all of the security principles that the service claims to implement.

Cloud services often rely upon third party products and services. Those third parties can have an impact on the overall security of the services. If this principle is not implemented then it is possible that supply chain compromise can undermine the security of the service and affect the implementation of other security principles.

Implementation objectives

The consumer understands and accepts:

- How their information is shared with, or accessible by, third party suppliers and their supply chains
- How the service provider's procurement processes place security requirements on third party suppliers and delivery partners
- How the service provider manages security risks from third party suppliers and delivery partners
- How the service provider manages the conformance of their suppliers with security requirements
- How the service provider verifies that hardware and software used in the service are genuine and have not been tampered with.

https://www.gov.uk/government/publications/implementing-the-cloud-security-principles/implementing-the-cloud-security-principles#principle-8-supply-chain-security

under ISO 27001:2013, AICPA SOC 1, SOC 2, SOC 3 and PCI-DSS certification programs. These certifications are recognized by ENISA under the Cloud Certification Schemes. The controls in relation to supply chain security are validated independently at least annually under the certification programs.

Principle 9: Secure consumer management

Implementation approach

AWS Identity and Access Management (IAM) provides you with controls and features to provide confidence that authenticated and authorised users have access to specified services and interfaces. AWS IAM allows you to create multiple users and manage the permissions for each of these users within your AWS Account. A user is an identity (within an AWS Account) with unique security credentials that can be used to access AWS Services. AWS IAM eliminates the need to share passwords or keys, and makes it easy to enable or disable a user's access as appropriate.

AWS IAM enables you to implement security best practices, such as least privileged, by granting unique

Secure consumer management

Consumers should be provided with the tools required to help them securely manage their services. Management interfaces and procedures are a vital security barrier in preventing unauthorised people accessing and altering consumers' resources, applications and data.

9.1 Authentication of consumers to management interfaces and within support channels

In order to maintain a secure service, consumers need to be securely authenticated before being allowed to perform management activities, report faults or request changes to the service. These activities may be conducted through a service management web portal, or through other support channels (such as telephone or email) and are likely to facilitate functions such as provisioning new service elements, managing user accounts and managing consumer data. It is important that service providers ensure any management requests which could have a security impact are performed over secure and authenticated channels. If consumers are not strongly authenticated then an attacker posing as them could perform privileged actions undermining the security of their service or data.

Implementation objectives

The consumer:

- Has sufficient confidence that only authorised individuals from the consumer organisation are able to authenticate to and access management interfaces for the service (Principle 10 should be used to assess the risks of different approaches to meet this objective)
- Has sufficient confidence that only authorised individuals from the consumer organisation are able to perform actions affecting the consumer's service through support channels.

https://www.gov.uk/government/publications/implementing-the-cloud-security-principles/implementing-the-cloud-security-principles#principle-9-secure-consumer-management

credentials to every user within your AWS Account and only granting permission to access the AWS services and resources required for the users to perform their jobs. AWS IAM is secure by default; new users have no access to AWS until permissions are explicitly granted.

AWS IAM is also integrated with the AWS Marketplace, so that you can control who in your organization can subscribe to the software and services offered in the Marketplace. Since subscribing to certain software in the Marketplace launches an EC2 instance to run the software, this is an important access control feature. Using AWS IAM to control access to the AWS Marketplace also enables AWS Account owners to have fine-grained control over usage and software costs.

AWS IAM enables you to minimize the use of your AWS Account credentials. Once you create AWS IAM user accounts, all interactions with AWS Services and resources should occur with AWS IAM user security credentials. More information about AWS IAM is available on the AWS website: http://aws.amazon.com/iam/.

Delegate API Access to AWS Services Using IAM Roles

AWS supports a very important and powerful use case with AWS Identity and Access Management (IAM) roles in combination with IAM users to enable cross-account API access or delegate API access within an account. This functionality gives better control and simplifies access management when managing services and resources across multiple AWS accounts. You can enable cross-account API access or delegate API access within an account or across multiple accounts without having to share long-term security credentials.

When you assume an IAM role, you get a set of temporary security credentials that have the permissions associated with the role. You can use these temporary security credentials instead of your long-term security credentials in calls to AWS services. Users interact with the service with the permissions granted to the IAM role assumed. This reduces the potential attack surface area by having to create and manage fewer user credentials, and users don't have to remember multiple passwords.

Assurance approach

The secure consumer management sub-principle and related processes within AWS services are subject to audit at least annually under ISO 27001:2013, AICPA SOC 1, SOC 2, SOC 3 and PCI-DSS certification programs. These certifications are recognised by ENISA under the Cloud Certification Schemes. The controls in relation to secure consumer management are validated independently at least annually under the certification programs. Based on the alternatives provided for selection within Cloud Security Principles guidance, AWS uses Service Provider Assertion in respect of region-specific requirements.

Separation and access control within management interfaces

Implementation approach

API calls to launch and terminate instances, change firewall parameters, and perform other functions are all signed by your Amazon Secret Access Key, which could be either the AWS Accounts Secret Access Key or the Secret Access key of a user created with AWS IAM. Without access to your Secret Access Key, Amazon EC2 API calls cannot be made on your behalf. In addition, API calls can be encrypted with SSL to maintain confidentiality. Amazon recommends always using SSL-protected API endpoints.

AWS IAM also enables you to further control what APIs a user has permissions to call to manage a specific resource.

Cross-Account Access for better identity management

In AWS, assuming a role is a security principle that enables the user to assign policies that grant permissions to perform actions on AWS resources. Unlike with a user account, you don't sign in to a role. Instead, you are already signed in as a user, and then you switch to the role, temporarily giving up your original user permissions and assuming the permissions of the role. When you are done using the role, you revert to your user's permissions again.

As documented in the IAM User Guide, an administrator creates a role in an account with resources to be managed, and then specifies the AWS account IDs that are

9.2 Separation and access control within management interfaces

Many cloud services are managed via web applications or APIs. These interfaces are a key part of the service's security. If consumers are not adequately separated within management interfaces then one consumer may be able to affect the service, or modify data belonging to another.

Implementation objectives

The consumer:

- Has sufficient confidence that other consumers cannot access, modify or otherwise affect their service management
- Can manage the risks of their own privileged access, e.g. through 'principle of least privilege', providing the ability to constrain permissions given to consumer administrators
- Understands how management interfaces are protected (see Principle 11) and what functionality is available via those interfaces.

https://www.gov.uk/government/publications/implementing-the-cloud-security-principles/implementing-the-cloud-security-principles#principle-9-secure-consumer-management

trusted to use the role. The administrators of the trusted accounts then grant permissions to specific users who can switch to the role.

Delegating access through roles this way can help you improve your security posture by simplifying the management of credentials. Instead of having to provide your users with sign-in credentials for every account that they need to access, users only need one set of sign-in credentials. This leads to a reduction in the potential attack surface area by having fewer user credentials that you have to create and manage, and your users don't have to remember multiple passwords.

This feature can be used to help improve security within a single account. When you create a typical user, you give that user permissions to access all of the resources needed to do the job - even the most sensitive and rarely accessed resources. Ideally, a user shouldn't have any access to the sensitive and critical resources until actually needed to keep to the security principle of "least access". The ability to delegate permissions to a role and allow a user to switch to the role solves this dilemma. Grant the user only those permissions that allow access to the normal day-to-day managed resources and *not* to the sensitive resources. Instead, grant to a role the permissions to access sensitive resources. The user can switch to the role when needing to use those resources and then switch right back to their user account. This feature helps reduce the attack surface area.

Assurance approach

The separation and access control within management interfaces sub-principle and related processes within AWS services are subject to audit at least annually under ISO 27001:2013, AICPA SOC 1, SOC 2, SOC 3 and PCI-DSS certification programs. These certifications are recognised by ENISA under the Cloud Certification Schemes. The controls in relation to separation and access control within management interfaces are validated independently at least annually under the certification programs. Based on the alternatives provided for selection within Cloud Security Principles guidance, AWS uses Service Provider Assertion in respect of region-specific requirements.

Principle 10: Identity and authentication

Implementation approach

AWS provides a number of ways for you to identify users and securely access your AWS Account. A complete list of credentials supported by AWS can be found on the Security Credentials page under 'Your Account'. AWS also provides additional security options that enable you to further protect your AWS Account and control access: AWS Identity and Access Management (AWS IAM), key management and rotation, temporary security credentials, and multi-factor authentication (MFA). AWS IAM enables you to minimize the use of your AWS Account credentials. Once you create AWS IAM user accounts, all interactions with AWS Services and resources should occur with AWS IAM user security credentials. More information about AWS IAM is available on the AWS website: http://aws.amazon.com/iam/

Host Operating System: Administrators with a business need to access the management plane are required to use multi-factor authentication to gain access to purpose-built administration hosts. These administrative hosts are systems that are specifically designed, built, configured, and hardened to protect the management plane of the cloud. All such access is logged and audited. When an employee no longer has a business need to access the management plane, the privileges and access to these hosts and relevant systems can be revoked.

Guest Operating System: Virtual instances are completely controlled by you, the customer. You have full root access or administrative control over accounts, services, and applications. AWS does not have any access rights to your instances or the guest OS. AWS recommends a base set of security best practices to include disabling password-only access to your guests, and utilizing some form of multi-factor authentication to gain access to your instances (or at a minimum certificate-based SSH Version 2 access). Additionally, you should employ a privilege escalation mechanism with logging on a per-user basis. For example, if the guest OS is Linux,

Identity and authentication

Consumer and service provider access to all service interfaces should be constrained to authenticated and authorised individuals.

All cloud services will have some requirement to identify and authenticate users wishing to access service interfaces. Weak authentication or access control may allow unauthorised changes to a consumer's service, theft or modification of data, or denial of service.

Implementation objectives

Consumers should have sufficient confidence that identity and authentication controls ensure users are authorised to access specific interfaces.

https://www.gov.uk/government/publications/implementing-the-cloud-security-principles/implementing-the-cloud-security-principles#principle-10-identity-and-authentication

after hardening your instance you should utilize certificate-based SSHv2 to access the virtual instance, disable remote root login, use command-line logging, and use 'sudo' for privilege escalation. You should generate your own key pairs in order to guarantee that they are unique, and not shared with other customers or with AWS.

AWS also supports the use of the Secure Shell (SSH) network protocol to enable you to log in securely to the EC2 instances. Authentication for SSH used with AWS is via a public/private key pair to reduce the risk of unauthorized access to your instance. You can also connect remotely to your Windows instances using Remote Desktop Protocol (RDP) by utilizing an RDP certificate generated for your instance.

AWS IAM enables you to implement security best practices, such as least privilege, by granting unique credentials to every user within your AWS Account and only granting permission to access the AWS services and resources required for the users to perform their jobs. AWS IAM is secure by default; new users have no access to AWS until permissions are explicitly granted.

AWS IAM is also integrated with the AWS Marketplace, so that you can control who in your organization can subscribe to the software and services offered in the Marketplace. Since subscribing to certain software in the Marketplace launches an EC2 instance to run the software, this is an important access control feature. Using AWS IAM to control access to the AWS Marketplace also enables AWS Account owners to have fine-grained control over usage and software costs.

Assurance approach

The identity and authentication principle and related processes within AWS services are subject to audit at least annually under ISO 27001, SOC 1, SOC 2, SOC 3 and PCI-DSS certification programs. These certifications are recognised by ENISA under the Cloud Certification Schemes. The controls in relation to identity and authentication are validated independently at least annually under the certification programs. Based on the alternatives provided for selection within Cloud Security Principles guidance, AWS uses Service Provider Assertion in respect of region-specific requirements.

Principle 11: External interface protection

Implementation approach

Helping to protect the confidentiality, integrity, and availability of our customers' systems and data is of the utmost importance to AWS, as is maintaining customer trust and confidence.

The AWS network has been architected to permit you to select the level of security and resiliency appropriate for your workload. To enable you to build geographically dispersed, fault-tolerant web architectures with cloud resources, AWS has implemented a world-class network infrastructure that is carefully monitored and managed.

Secure Network Architecture

Network devices, including firewall and other boundary devices, are in place to monitor and control communications at the external boundary of the network and at key internal boundaries within the network. These boundary devices employ rule sets, access control lists (ACL), and configurations to enforce the flow of information to specific information system services.
ACLs, or traffic flow policies, are established on each managed interface, which manage and enforce the flow of traffic. ACL policies are approved by Amazon Information Security. These policies are automatically pushed using AWS's ACL-Manage tool, to help ensure these managed interfaces enforce the most up-to-date ACLs.

Secure Access Points

AWS has strategically placed a limited number of access points to the cloud to allow for a more comprehensive monitoring of inbound and outbound communications and network traffic. These customer access points are called API endpoints, and they allow secure HTTP access (HTTPS), which allows you to establish a secure communication session with your storage or compute instances within AWS. In addition, AWS has implemented network devices

External interface protection
All external or less trusted interfaces of the service should be identified and have appropriate protections to defend against attacks through them.

If an interface is exposed to consumers or outsiders and it is not sufficiently robust, then it could be subverted by attackers in order to gain access to the service or data within it. If the interfaces exposed include private interfaces (such as management interfaces) then the impact may be more significant.

Consumers can use different models to connect to cloud services which expose their enterprise systems to varying levels of risk.

Implementation objectives

- The consumer understands how to safely connect to the service whilst minimising risk to the consumer's systems
- The consumer understands what physical and logical interfaces their information is available from
- The consumer has sufficient confidence that protections are in place to control access to their data
- The consumer has sufficient confidence that the service can determine the identity of connecting users and services to an appropriate level for the data or function being accessed.

https://www.gov.uk/government/publications/implementing-the-cloud-security-principles/implementing-the-cloud-security-principles#principle-11-external-interface-protection

that are dedicated to managing interfacing communications with Internet service providers (ISPs). AWS employs a redundant connection to more than one communication service at each Internet-facing edge of the AWS network. These connections each have dedicated network devices.

Transmission Protection

You can connect to an AWS access point via HTTP or HTTPS using Secure Sockets Layer (SSL), a cryptographic protocol that is designed to protect against eavesdropping, tampering, and message forgery. For customers who require additional layers of network security, AWS offers the Amazon Virtual Private Cloud (VPC), which provides a private subnet within the AWS cloud, and the ability to use an IPsec Virtual Private Network (VPN) device to provide an encrypted tunnel between the Amazon VPC and your data center.

Network Monitoring and Protection

AWS utilizes a wide variety of automated monitoring systems to provide a high level of service performance and availability. AWS monitoring tools are designed to detect unusual or unauthorized activities and conditions at ingress and egress communication points. These tools monitor server and network usage, port scanning activities, application usage, and unauthorized intrusion attempts. The tools have the ability to set custom performance metrics thresholds for unusual activity.

Systems within AWS are extensively instrumented to monitor key operational metrics. Alarms are configured to automatically notify operations and management personnel when early warning thresholds are crossed on key operational metrics. An on-call schedule is used so personnel are always available to respond to operational issues. This includes a pager system so alarms are quickly and reliably communicated to operations personnel.

Documentation is maintained to aid and inform operations personnel in handling incidents or issues. If the resolution of an issue requires collaboration, a conferencing system is used which supports communication and logging capabilities. Trained call leaders facilitate communication and progress during the handling of operational issues that require collaboration. Post-mortems are convened after any significant operational issue, regardless of external impact, and Cause of Error (COE) documents are drafted so the root cause is captured and preventative actions are taken in the future. Implementation of the preventative measures is tracked during weekly operations meetings.

AWS security monitoring tools help identify several types of denial of service (DoS) attacks, including distributed, flooding, and software/logic attacks. When DoS attacks are identified, the AWS incident response process is initiated. In addition to the DoS prevention tools, redundant

telecommunication providers at each region as well as additional capacity protect against the possibility of DoS attacks.

Assurance approach

The external interface protection principle and related processes within AWS services are subject to audit at least annually under ISO 27001:2013, AICPA SOC 1, SOC 2, SOC 3 and PCI-DSS certification programs. These certifications are recognised by ENISA under the Cloud Certification Schemes. The controls in relation to external interface protection are validated independently at least annually under the certification programs. Based on the alternatives provided for selection within Cloud Security Principles guidance, AWS uses Service Provider Assertion in respect of region-specific requirements.

Principle 12: Secure service administration

Implementation approach

User Access

Procedures exist so that Amazon employee and contractor user accounts are added, modified, or disabled in a timely manner and are reviewed on a periodic basis. In addition, password complexity settings for user authentication to AWS systems are managed in compliance with Amazon's Corporate Password Policy.

Account Provisioning

The responsibility for provisioning employee and contractor access is shared across Human Resources (HR), Corporate Operations and Service Owners.

A standard employee or contractor account with minimum privileges is provisioned in a disabled state when a hiring manager submits his or her new employee or contractor onboarding request in Amazon's HR system. The account is automatically enabled when the employee's record is activated in

Secure service administration
The methods used by the service provider's administrators to manage the operational service should be designed to mitigate any risk of exploitation that could undermine the security of the service.

The security of a cloud service is closely tied to the security of the service provider's administration systems. Access to service administration systems gives an attacker high levels of privilege and the ability to affect the security of the service. Therefore the design, implementation and management of administration systems should reflect their higher value to an attacker.

Implementation objectives
Consumers have sufficient confidence that the technical approach the service provider uses to manage the service does not put their data or service at risk.

https://www.gov.uk/government/publications/implementing-the-cloud-security-principles/implementing-the-cloud-security-principles#principle-12-secure-service-administration

Amazon's HR system. First time passwords are set to a unique value and are required to be changed on first use.

Access to other resources including Services, Host, Network devices, and Windows and UNIX groups is explicitly approved in Amazon's proprietary permission management system by the appropriate owner or manager. Requests for changes in access are captured in the Amazon permissions management tool audit log. When changes in an employee's job function occur, continued access must be explicitly approved to the resource or it will be automatically revoked.

Periodic Account Review

Accounts are reviewed every 90 days; explicit re-approval is required or access to the resource is automatically revoked.

Access Removal

Access is automatically revoked when an employee's record is terminated in Amazon's HR system. Windows and UNIX accounts are disabled and Amazon's permission management system removes the user from all systems.

Password Policy

Access and administration to logical security has access mechanisms, users IDs, passwords and Kerberos to authenticate users to services, resources and devices as well as to authorize the appropriate level of access for the user. AWS Security has established a password policy with required configurations and expiration intervals.

Administrators with a business need to access the management plane are required to use multifactor authentication to gain access to purpose-built administration hosts. These administrative hosts are systems that are specifically designed, built, configured, and hardened to protect the management plane of the cloud. All such access is logged and audited. When an employee no longer has a business need to access the management plane, the privileges and access to these hosts and relevant systems are revoked.

Assurance approach

The secure service administration principle and related processes within AWS services are subject to audit at least annually under ISO 27001:2013, AICPA SOC 1, SOC 2, SOC 3 and PCI-DSS certification programs. These certifications are recognised by ENISA under the Cloud Certification Schemes. The controls in relation to secure service administration are validated independently at least annually under the certification programs. Based on the alternatives provided for selection within Cloud Security Principles guidance, AWS uses Service Provider Assertion in respect of region-specific requirements.

Principle 13: Audit information provision to consumers

Implementation approach

AWS CloudTrail is a service that provides audit records for AWS customers and delivers audit information in the form of log files to a specified storage bucket. The recorded information includes the identity of the API caller, the time of the API call, the source IP address of the API caller, the request parameters, and the response elements returned by the AWS service.

CloudTrail provides a history of AWS API calls for customer accounts, including API calls made via the AWS Management Console, AWS SDKs, command line tools, and higher-level AWS services (such as AWS CloudFormation). The AWS API call history produced by CloudTrail enables security analysis, resource change tracking, and compliance auditing.

The logfile objects written to S3 are granted full control to the bucket owner. The bucket owner thus has full control over whether to share the logs with anyone else.

This feature enables AWS customers and provides confidence to meet their needs for investigating service misuse or incidents.

Audit information provision to consumers

Consumers should be provided with the audit records they need to monitor access to their service and the data held within it.

The type of audit information available to consumers will have a direct impact on their ability to detect and respond to inappropriate or malicious usage of their service or data within reasonable timescales.

Implementation objectives

Consumers are:

- Aware of the audit information that will be provided to them, how and when it will be made available to them, the format of the data, and the retention period associated with it.
- Confident that the audit information available will allow them to meet their needs for investigating misuse or incidents.

https://www.gov.uk/government/publications/implementing-the-cloud-security-principles/implementing-the-cloud-security-principles#principle-13-audit-information-provision-to-consumers

More details on AWS CloudTrail and further information on audit records can be requested at http://aws.amazon.com/cloudtrail. A latest version of CloudTrail User Guide is available at http://awsdocs.s3.amazonaws.com/awscloudtrail/latest/awscloudtrail-ug.pdf.

Assurance approach

The audit information provision to the consumer's principle and related processes within AWS services are subject to audit at least annually under ISO 27001:2013, AICPA SOC 1, SOC 2,

SOC 3 and PCI-DSS certification programs. These certifications are recognised by ENISA under the Cloud Certification Schemes. The controls in relation to audit information provision to consumers are validated independently at least annually under the certification programs. Based on the alternatives provided for selection within Cloud Security Principles guidance, AWS uses Service Provider
Assertion in respect of region-specific requirements.

Principle 14: Secure use of the service by the consumer

Implementation approach

AWS has implemented various methods of external communication to support you and the wider customer base and the community. AWS has published a public Acceptable Use Policy that provides guidance and informs consumers on acceptable use of AWS services. This policy includes guidance on illegal, harmful, or offensive content, security violations, network abuse and e-mail or message abuse with information on monitoring and enforcement of the policy. Additionally, guidance is provided on reporting violations of the Acceptable Use Policy.

Mechanisms are in place to allow the customer support team to be notified of operational issues that impact the customer experience. A "Service Health Dashboard" is

Secure use of the service by the consumer

Consumers have certain responsibilities when using a cloud service in order for their use of it to remain secure, and for their data to be adequately protected.

The security of cloud services and the data held within them can be undermined by poor use of the service by consumers. The extent of the responsibility on the consumer for secure use of the service will vary depending on the deployment models of the cloud service, specific features of an individual service and the scenario in which the consumers intend to the use the service. IaaS and PaaS offerings are likely to require the consumer to be responsible for significant aspects of the security of their service.

Implementation objectives

- The consumer understands any service configuration options available to them and the security implications of choices they make
- The consumer understands the security requirements on their processes, uses, and infrastructure related to the use of the service

The consumer can educate those administrating and using the service in how to use it safely and securely.

https://www.gov.uk/government/publications/implementing-the-cloud-security-principles/implementing-the-cloud-security-principles#principle-14-secure-use-of-the-service-by-the-consumer

available and maintained by the customer support team to alert customers to any issues that may be of broad impact. The AWS Security Center is available to provide you with security and compliance details about AWS.

Customers can also subscribe to AWS Support offerings that include direct communication with the customer support team and proactive alerts to any customer impacting issues.

Using the Trusted Advisor Tool

Some AWS Support plans include access to the Trusted Advisor tool, which offers a one-view snapshot of your service and helps identify common security misconfigurations, suggestions for improving system performance, and underutilized resources.

Trusted Advisor checks for compliance with the following security recommendations:

- Limited access to common administrative ports to only a small subset of addresses. This includes ports 22 (SSH), 23 (Telnet) 3389 (RDP), and 5500 (VNC)
- Limited access to common database ports. This includes ports 1433 (MSSQL Server), 1434 (MSSQL Monitor), 3306 (MySQL), Oracle (1521) and 5432 (PostgreSQL)
- IAM is configured to help ensure secure access control of AWS resources
- Multi-factor authentication (MFA) token is enabled to provide two-factor authentication for the root AWS account

Assurance approach

The secure use of the service by the consumer principle and related processes are not validated independently within AWS compliance programs. Based on the alternatives provided for selection within Cloud Security Principles guidance, the controls in relation to secure use of the service by the consumer do not exist within the existing certification programs for them to be validated independently. AWS publishes guidance on configuration options and the relative impacts on security regularly through various communication channels like local summit sessions, webinars, blogs, and training and guidance documents. AWS uses Service Provider Assertion in respect of region-specific requirements.

Conclusion

The AWS cloud platform provides a number of important benefits to UK public sector organisations and enables you to meet the objectives of the fourteen Cloud Security Principles. While AWS delivers these benefits and advantages through our services and features, the individual public sector organisations are ultimately responsible for risk management decisions relating to the use of secure cloud services for OFFICIAL information. Using the information presented in this whitepaper, we encourage you to use AWS services for your organisations to manage security and the related risks appropriately.

For AWS, security is always our top priority. We deliver services to hundreds of thousands of businesses including enterprises, educational institutions, and government agencies in over 190 countries. Our customers include government agencies, financial services and healthcare providers who leverage the benefits of AWS while retaining control and responsibility for their data including some of their most sensitive information.

AWS services are designed to give customers flexibility over how they configure and deploy their solutions as well as control over their content, including where it is stored, how it is stored and who has access to it and the security configuration environment. AWS customers can build their own secure applications and store content securely on AWS.

Additional Resources

To help customers further understand how they can address their privacy and data protection requirements, customers are encouraged to read the risk, compliance and security whitepapers, best practices, checklists and guidance published on the AWS website. This material can be found at:

AWS Compliance: http://aws.amazon.com/compliance
AWS Security Center: http://aws.amazon.com/security

AWS also offers training to help customers learn how to design, develop, and operate available, efficient, and secure applications on the AWS cloud and gain proficiency with AWS services and solutions. We offer free instructional videos, self-paced labs, and instructor-led classes. Further information on AWS training is available at http://aws.amazon.com/training/.

AWS certifications certify the technical skills and knowledge associated with best practices for building secure and reliable cloud-based applications using AWS technology. Further information on AWS certifications is available at http://aws.amazon.com/certification/.

If further information is required, please contact AWS at: https://aws.amazon.com/contact-us/ or contact the local AWS account representative.

Appendix – AWS Platform Benefits

When designing and implementing large, cloud-based applications, it's important to consider how infrastructure will be managed to ensure the cost and complexity of running such systems is minimized. When organisations first begin using AWS platform, it is easy to manage EC2 instances just like regular virtualised servers running in a data center.

However, as the architecture evolves and changes are made over time, the instances will inevitably begin to diverge from their original specification, which can lead to inconsistencies with other instances in the same environment. This divergence from a known baseline can become a huge challenge when managing large fleets of instances across multiple environments. Ultimately, it will lead to service issues because these environments will become less predictable and more difficult to maintain.

The AWS platform provides a rich and diverse set of tools to address this challenge with a different approach. By using the AWS platform and features, public sector organisations can specify and manage the desired end state of the infrastructure independently of the instances and other running components.

When technology teams start to think of infrastructure as being defined independently of the running instances and other components in the environments, they can take greater advantage of the benefits of dynamic cloud environments:

Software-defined infrastructure – By defining infrastructure using a set of software artifacts, many of the tools and techniques that are used in the development of software components can be leveraged. This includes managing the evolution of infrastructure in a version control system, as well as using continuous integration (CI) processes to continually test and validate infrastructure changes before deploying them to production.

Auto Scaling and self-healing – If new instances are provisioned automatically from a consistent specification, Auto Scaling groups can be used to manage the number of instances in an EC2 fleet. For example, a condition to add new EC2 instances in increments can be set to the Auto Scaling group when the average utilization of EC2 fleet is high. Auto Scaling can also be used to detect impaired EC2 instances and unhealthy applications, and replace the instances without intervention.

Fast environment provisioning – Consistent environments can be provisioned quickly and easily, which opens up new ways of working within teams. For example, a new environment can be provisioned to allow testers to validate a new version of an application in their own, personal test environments that are isolated from other changes.

Reduce costs – Now that environments can be provisioned quickly, the option is always there to remove them when they are no longer needed. This reduces costs because customers are charged only for the resources that are used.

Blue-green deployments – Application teams can deploy new versions of application by provisioning new instances (containing a new version of the code) beside the existing infrastructure. Traffic can be switched between environments in an approach known as blue-green deployments. This has many benefits over traditional deployment strategies, including the ability to quickly and easily roll back a deployment in the event of an issue.

In addition to the implementation and assurance approaches detailed in this whitepaper for each Cloud Security Principle, public sector organisations adopting cloud technologies should take into consideration the additional benefits of AWS platform within the risk assessment and management frameworks. Whilst a secure and compliant public cloud environment is necessary for handling government OFFICIAL information, the AWS platform and security features that scale and enable resilience to change are equally important to consider.