G-Cloud 15 Amazon Web Services EMEA SARL, UK Branch (AWS)

AWS Services

Service Definition Document

Lot 1a Infrastructure as a Service (IaaS) and Platform as a Service (PaaS)

Lot 2a Infrastructure as a Service (I-SaaS)

Lot 2b Software as a Service (SaaS)

30 January 2026

G-Cloud 15 Amazon Web Services EMEA SARL, UK Branch (AWS)

AWS Services

Service Definition Document

Lot 1a Infrastructure as a Service (IaaS) and Platform as a Service (PaaS)

Lot 2a Infrastructure as a Service (I-SaaS)

Lot 2b Software as a Service (SaaS)

Table of Contents

INTRODUCTION

Please note that we have consolidated common elements of each Service Offering and have provided descriptions for these common elements that apply equally to each Service Offering—see "CROSS-SERVICE DEFINITIONS". To find out more about AWS on G-Cloud and AWS Cloud services, visit us at AWS on G-Cloud UK.

The AWS Free Tier enables you to gain free, hands-on experience with AWS products and services. The AWS Free Tier provides customers the ability to explore and try out AWS services free of charge up to specified limits for each service, with offerings available for up to 12 months for traditional accounts or through a $200 credit system (valid for 6 months) for accounts created after July 15, 2025.. See http://aws.amazon.com/free/ for service-specific details.

CROSS-SERVICE DEFINITIONS

1 Shared Responsibility Model

Customers and AWS share security and compliance responsibilities. As shown in Figure 1, AWS operates the services and underlying infrastructure that make up the cloud. We secure both the facilities that house the hardware and the foundational cloud technologies, like virtualization tools. You are responsible for managing your data, guest operating systems, and applications, as well as configuring the security of your AWS environment.

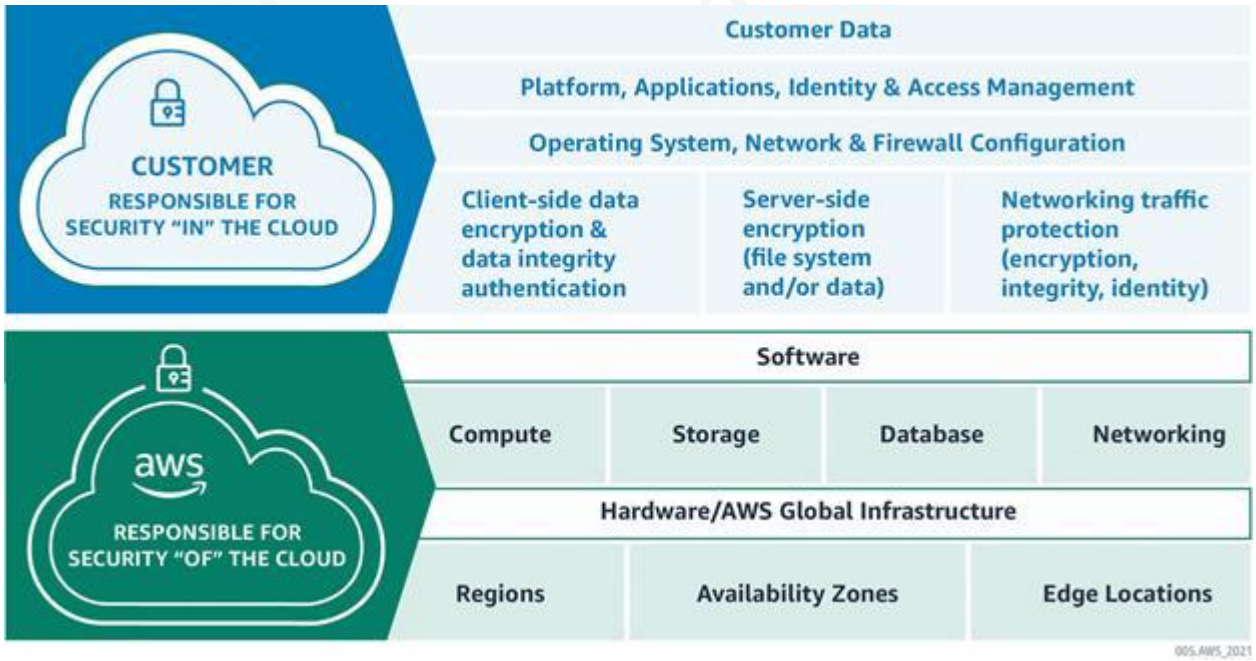

Figure 1. The Shared Responsibility Model.

This model establishes that AWS is responsible for the underlying cloud infrastructure, while customers are responsible for their data and applications. The customer's specific responsibilities depend on the AWS services deployed.

AWS Security Responsibilities: Security of the Cloud

AWS is responsible for *security of the cloud*. This means protecting the infrastructure that runs all of the services we offer. This infrastructure is composed of hardware, software, networking, and facilities. To protect that infrastructure, we designed and built it with the unique needs of the cloud in mind.

We use redundant and layered controls, continuous validation and testing, and automation to monitor and protect our underlying infrastructure 24/7. We replicate these controls in every new data center or service. We built our data centers and network architecture to meet the requirements of our most security-sensitive customers. This includes the Intelligence Community, DoD, and civilian agencies, as well as organizations in regulated industries like healthcare, aerospace and defense, energy, and financial services. As a result, all our customers get a resilient infrastructure designed for high security.

AWS has implemented sophisticated technical and physical measures against unauthorized access to protect our customers' security. Customers can validate the security controls in place within the AWS environment through our certifications and reports, including the AWS System and Organization Control (SOC) 1, 2, and 3 reports; International Organization for Standardization (ISO) 27001, 27017, 27018, and 9001 certifications; and Payment Card Industry (PCI) Data Security Standard (DSS) compliance reports. The ISO 27018 certification demonstrates that we have a system of controls in place for the privacy protection of customer content. Independent third-party auditors produce these certifications and reports, which attest to the design and operating effectiveness of AWS security controls. You can read more about AWS compliance certifications, reports, and alignment with industry practices and standards at http://aws.amazon.com/compliance/.

Customer Responsibilities: Security in the Cloud

AWS customers are responsible for *security in the cloud*. This refers to your applications, content, and systems. You have complete control over which services you use and whom you empower to access your content and services, including what credentials are required to gain access. When using AWS services, you retain ownership and control of your content.

Because you, rather than AWS, control these important factors, you are responsible for securing any content that you store or process using our services and any content you connect to AWS infrastructure. This means you are responsible for the guest operating system, applications on your compute instances, and content stored and processed in AWS storage, databases, or other services. You will need to set up access control policies so that only authorized individuals can access resources.

Security Responsibilities by Service Layer

The division of responsibilities varies significantly based on the AWS service layer selected.

For Infrastructure as a Service (IaaS) offering like Amazon EC2, customers have responsibilities including guest operating system management, application software and utilities, security group configuration, network access control lists, identity and access management, and data encryption.

For Platform as a Service (PaaS) offerings such as managed services, AWS assumes additional responsibilities including operating system maintenance and patching, database or platform software updates, network configuration and security, backup and recovery infrastructure, and platform-level security controls. Customers focus on data management, encryption options, IAM policies, and application-level access controls.

For serverless services like AWS Lambda, AWS manages the most comprehensive set of responsibilities including infrastructure and foundation services, operating system and runtime environment, application platform and execution environment, server provisioning and capacity management, and network configuration and security. This frees customers to focus on securing application code, storing and accessing sensitive data, implementing identity and access management, and maintaining monitoring and logging.

Shared Responsibility and Operational Controls

Just as AWS and our customers share responsibilities to operate the IT environment, we also share responsibilities to manage, operate, and verify IT controls. As every customer's deployment in AWS is different, customers can shift management of certain controls to AWS, resulting in a distributed control environment. The following are examples of controls managed by AWS, our customers, or both AWS and customers.

Inherited Controls

We reduce your security burden by managing the controls associated with AWS's physical infrastructure (inherited controls).

Shared Controls

Some controls apply to both the infrastructure layer (AWS responsibility) and customer layers (customer responsibility), but in completely separate contexts or perspectives (shared controls). With shared controls, AWS provides the requirements for the infrastructure, and the customer provides their own control implementation within their AWS services. Examples of these shared controls include the following:

- **Patch Management:** AWS is responsible for patching and fixing flaws within the infrastructure, but customers are responsible for patching their guest operating system and applications.
- **Configuration Management:** AWS maintains the configuration of our infrastructure devices, but customers are responsible for configuring their own guest operating systems, databases, and applications.
- **Awareness and Training:** AWS trains AWS employees, but customers must train their own employees.

Customer-Specific Controls

Some controls are the sole responsibility of the customer based on the applications they deploy within the AWS Cloud (customer-specific controls). Examples include service and communications protection or zone security, which may require a customer to route or zone data within specific security environments. Customers also control how they use their account and what content moves into and out of the account. They also need visibility into vulnerabilities that can threaten their applications, content, and systems.

Operational Responsibilities Definition

Beyond security, the shared responsibility model encompasses operational responsibilities. AWS manages the global infrastructure, service availability, and platform maintenance, while customers manage their workload configuration, data governance, access management, and application performance optimization. This specialization allows AWS and customers to focus on their respective areas of expertise, creating measurable outcomes across security, operations, and cost management while maintaining customer control over critical business decisions.

Shared Responsibility and Customer Data

AWS classifies customer data into two categories: customer content and account information.

Customer Content

With AWS, customers get the ability to store their content on cloud infrastructure and access it from almost anywhere over the internet. As a customer, you maintain full control of your content that you upload to the AWS services under your AWS account. This means that you

maintain responsibility for configuring access to AWS services and resources. We provide an advanced set of access, encryption, and logging features to help you do this effectively. This includes APIs through which you can configure access control permissions for any of the services you develop or deploy in an AWS environment. We do not access or use your content for any purpose without your agreement. We never use your content or derive information from it for marketing or advertising purposes.

We define ***customer content*** as software (including machine images), data, text, audio, video, or images that a customer or any end user transfers to us for processing, storage, or hosting by AWS services. For example, customer content includes content that a customer or any end user stores in Amazon Simple Storage Service (S3). Customer content does not include account information which we describe below. Customer content also does not include information included in resource identifiers, metadata tags, usage policies, permissions, and similar items related to the management of AWS resources. The terms of the AWS Customer Agreement and the AWS Service Terms apply to customer content.

As an AWS customer—and in keeping with the shared responsibility model for cloud—you maintain ownership and control over your content. AWS provides tools and guidance to support you in owning and managing your data. With these tools, you can:

- Determine where your content will be stored, including the type of storage and geographic region of that storage.
- Choose the secured state of your content—we offer customers strong encryption for your content in transit and at rest, and we provide you with the option to manage your own encryption keys.
- Manage access to your content and access to AWS services and resources through users, groups, permissions, and credentials that you control.

Legal Requests for Customer Content

AWS is vigilant about customer privacy. We will not disclose customer content unless we are required to do so to comply with the law or binding order of a governmental body. If a governmental body sends AWS a demand for customer content, we will attempt to redirect the governmental body to request that data directly from the customer. If compelled to disclose customer content to a governmental body, we will give customers reasonable notice of the demand to allow the customer to seek a protective order or other appropriate remedy unless AWS is legally prohibited from doing so. Governmental and regulatory bodies need to follow the applicable legal process to obtain valid and binding orders. We review all orders and object to overbroad or otherwise inappropriate ones. Unless prohibited from doing so, AWS notifies customers before disclosing customer content so they can seek protection from disclosure. AWS customers can encrypt their customer content, and we provide customers with the option to manage their own encryption keys.

Amazon, our parent company, knows that customers care deeply about privacy and data security and regularly publishes Amazon Information Request Reports about the types and volume of information requests that it receives, including the following:

- Amazon does not disclose customer information in response to government demands unless it is required to do so to comply with a legally valid and binding order. Unless prohibited from doing so or there is clear indication of illegal conduct in connection with the use of Amazon products or services, Amazon notifies customers before disclosing content information.
- Where Amazon needs to act to protect customers, it does. Amazon has repeatedly challenged government demands for customer information that were considered overbroad, winning decisions that have helped to set the legal standards for protecting customer speech and privacy interests. Amazon also advocates in Congress to modernize outdated privacy laws to require law enforcement to obtain a

search warrant from a court to get the content of customer communications. That is the appropriate standard, and it is the standard Amazon and its subsidiaries follow.

- While Amazon recognizes the legitimate needs of law enforcement agencies to investigate criminal and terrorist activity and cooperate with them when they observe legal safeguards for conducting such investigations, it opposes legislation mandating or prohibiting security or encryption technologies that would weaken the security of products, systems, or services its customers use, whether they be individual consumers or business customers. For AWS customers, we offer strong encryption as one of many standard security features, and we provide them the option to manage their own encryption keys. We publish security best practices documents on our website and encourage our clients to use these measures to protect sensitive content.
- Amazon is a member of numerous associations focused on protecting privacy and security, and AWS in particular has achieved a number of internationally recognized certifications and accreditations demonstrating compliance with third-party assurance frameworks. AWS clients have control over their content and where it resides.

Amazon Information Request Reports dating back to 2015 can be found on Amazon's Law Enforcement Information Requests page.

Account Information

We define account information as information about a customer that a customer provides to us in connection with the creation or administration of a customer account. For example, account information includes names, usernames, phone numbers, email addresses, and billing information associated with a customer account. The information practices described in the AWS Privacy Notice apply to account information.

Customer Virtual Instances

Customer virtual instances are solely controlled by the customer. AWS personnel do not have the ability to log into customer instances. AWS customers manage the creation and deletion of their data on AWS, maintain control of access permissions, and manage appropriate data retention policies and procedures.

Documentation To Support Your Security Objectives

AWS can help you handle your share of security of the infrastructure. We equip customers with documentation for all our services to guide you through proper configuration for security. We also offer hundreds of tools and features to help you meet your security objectives, including services for network security, configuration management, access control, and data encryption. Additionally, we offer integrated Partner solutions and support in the form of AWS customer enablement services to help you design, implement, and operate your security environment. For further information, refer to Shared Responsibilities at https://docs.aws.amazon.com/wellarchitected/latest/security-pillar/shared-responsibility.html

2 Technical Architecture & Service Delivery

2.1 Infrastructure Details

Cloud deployment model

Based on the nature of your application and the underlying services (compute, storage, database, etc.) it requires, you can use AWS services to create a flexible deployment model—a cloud based application, a hybrid deployment, or a private cloud (on premises)—that you can tailor to fit the needs of both your application and your organization. AWS

supports all four NIST-defined cloud deployment models, providing organizations with flexible infrastructure options aligned with their specific requirements and compliance needs.

Public Cloud. AWS cloud infrastructure is provisioned for open use by the public and exists on AWS premises. AWS cloud delivers massive scale, continuous innovation, and cost efficiency that traditional private cloud implementations cannot match. Through services like Amazon VPC, organizations create logically isolated environments with complete control over their virtual networking, achieving the privacy benefits of dedicated infrastructure while accessing the full breadth of AWS capabilities.

Private Cloud. AWS enables private cloud deployments through AWS Outposts, which extends AWS infrastructure, services, APIs, and tools to virtually any data center, colocation space, or on-premises facility for a truly consistent hybrid experience. Private cloud infrastructure is provisioned for exclusive use by a single organization providing enhanced control over security, compliance, and configuration while maintaining consistent AWS APIs and tools.

Hybrid Cloud. AWS provides hybrid cloud solutions that compose two or more distinct cloud infrastructures bound by standardized technology enabling data and application portability. Key hybrid services include AWS Direct Connect for dedicated network connections, AWS Outposts for on-premises AWS infrastructure, and Amazon ECS Anywhere for container orchestration across environments. These solutions bridge existing investments with cloud capabilities, supporting data center extension, application migration, and local data requirements.

Community Cloud. AWS supports community cloud models through AWS Organizations for centralized management across multiple accounts with shared governance requirements. The AWS European Sovereign Cloud represents a significant advancement in community cloud capabilities, designed specifically for European government organizations and highly regulated industries. This independent cloud infrastructure will be physically and logically separated from other AWS Regions, operating within the EU with its own separate authentication, operations, and control systems. The European Sovereign Cloud will offer the same AWS services, features, and security standards while meeting the highest levels of sovereignty requirements for regulated industries and public sector organizations in Europe.

Infrastructure location

AWS offers 240+ fully featured services from data centers globally, refer to https://aws.amazon.com/about-aws/global-infrastructure/. The AWS Cloud spans 120 Availability Zones within 38 Geographic Regions, including the Europe (London) Region launched in 2016 with 3 Availability Zones. Whether you need to deploy your application workloads across the globe in a single click, or you want to build and deploy specific applications closer to your end users with single-digit millisecond latency, AWS provides cloud infrastructure where and when you need it.

AWS Implementation of NIST 5 Essential Characteristics

AWS implements all NIST 5 Essential Characteristics through specific technical capabilities that demonstrate comprehensive cloud computing functionality.

On-Demand Self-Service. AWS enables consumers to unilaterally provision computing capabilities automatically without requiring human interaction. AWS allocates resources using API calls, reducing provisioning time to minutes. Access methods include the AWS Management Console for web-based interface capabilities, AWS APIs for direct programmatic access, and the AWS Command Line Interface for unified command-line management.

Broad Network Access. AWS implements broad network access through global infrastructure comprising over 400 Points of Presence across more than 300 cities worldwide. The global infrastructure provides low latency and high network quality through a fully redundant 100 GbE fiber network backbone. Services like Amazon CloudFront, AWS Wavelength, and AWS Local Zones extend network access capabilities to edge locations and 5G networks.

Resource Pooling. AWS resource pooling uses multi-tenancy where physical and virtual resources are dynamically assigned according to consumer demand. AWS has massive economies of scale, enabling virtually unlimited resources for less cost than maintaining independent data centers. Multi-tenant architecture patterns include Silo, Pool, and Bridge models, with services like Amazon DynamoDB and AWS Lambda optimized for elastic resource sharing.

Rapid Elasticity. AWS rapid elasticity operates through Amazon EC2 Auto Scaling, which automatically launches or terminates instances according to predefined conditions. Organizations can deploy hundreds to thousands of servers in minutes to meet workload demands. Elasticity helps drastically reduce spending by matching resource utilization to demand at the lowest possible cost. AWS Lambda provides built-in elastic scalability, scaling automatically from few requests to thousands per second.

Measured Service. AWS implements a measured service through pay-as-you-go pricing where customers pay only for computing resources consumed. Amazon CloudWatch serves as the core monitoring component, collecting data from over 70 AWS services. AWS generates detailed billing reports and provides cost management tools including AWS Cost Explorer and AWS Budgets for tracking usage and costs.

The NIST Cyber Security Framework (CSF). The whitepaper, 'Aligning to the NIST CSF in the AWS Cloud', helps you understand how AWS cloud offerings align to the NIST CSF, and the role of third-party validations confirming AWS services' alignment with the NIST CSF risk management practices. You can access the whitepaper at https://d0.awsstatic.com/whitepapers/compliance/NIST_Cybersecurity_Framework_CSF_Jan_2025.pdf

2.2 Service Delivery & Access

Service Delivery

AWS delivers its cloud services primarily through secure connections over the public internet, using HTTPS/TLS encryption to protect data in transit and approved networks.

The AWS Global Network offers multiple layers of encryption for data in transit, digitally signed routing information, and unique threat intelligence. UK government organizations can connect through networks including PSN, HSCN, and Janet.

For UK government approved networks, AWS supports connectivity options through AWS Direct Connect, which establishes dedicated network connections from premises to AWS.

AWS VPN solutions meet NCSC Foundation profile requirements and can be used with systems running at OFFICIAL classification level. AWS PrivateLink provides private connectivity between VPCs, AWS services, and on-premises applications.

Supplier Portal and Interface Accessibility

AWS demonstrates strong commitment to accessibility compliance through systematic evaluation against WCAG 2.1 and 2.2 Level AA standards, Section 508, and EN 301 549. AWS has received Accessibility Conformance Reports (ACRs) for over 160 AWS services as

of 2024. These ACRs provide assessments against WCAG Level A and AA requirements and are available through AWS Artifact.

AWS demonstrates a broader commitment to accessibility across its products—focusing on keyboard operability, color contrast, accessible labels, and screen reader compatibility. AWS generally integrates accessibility from the design phase and works with experts to help ensure accessible code. AWS customers can use almost all services directly without retrofitting or adjusting them for accessibility. The AWS Command Line Interface (CLI) can be used by customers who use assistive technology such as Job Access with Speech (JAWS), NonVisual Desktop Access (NVDA), and alternative input devices. Throughout the service development cycle, AWS developers use automated accessibility testing tools and manual testing with assistive technologies.

Multi-cloud Support

At AWS, we give customers the freedom to choose technology that best suits their needs.

For example, you can use the same service (AWS Systems Manager) to patch and update Amazon Elastic Compute Cloud (Amazon EC2) instances, servers running on-premises, and servers provided by other cloud providers. Similarly, you can use Amazon CloudWatch to monitor applications, compute resources, and other cloud resources in all those environments.

Additionally, Oracle Database@AWS allows organizations to run Oracle workloads on AWS while maintaining compatibility with Oracle Cloud Infrastructure, creating a bridge between cloud platforms. For data movement across clouds, AWS DataSync supports transfers between AWS and other providers including Google Cloud, Azure, and Oracle Cloud, allowing consistent data access regardless of location.

Security remains paramount in multi-cloud environments, with Amazon Security Lake centralizing security data from AWS, Azure, GCP, and on-premises sources into a unified data lake using the Open Cybersecurity Schema Framework. For comprehensive visibility, Amazon CloudWatch extends monitoring capabilities across clouds, allowing organizations to query metrics from AWS, Azure Monitor, and other environments through a single interface.

AWS's approach focuses on practical solutions that address real multi-cloud challenges while applying AWS's strengths in security, scalability, and operational excellence. The AWS Solutions for Hybrid and Multicloud page contains additional examples of our extension-based approach to adding new capabilities as well as documentation on effective multi-cloud strategies.

We allow free data transfer out to the internet (DTO) when you want to move outside of AWS. We are committed to helping you be successful regardless of your approach.

We work closely with customers to develop a deep understanding of your multi-cloud goals. We can partner with you to help build the business case for your multi-cloud strategy and guidance on best practices. We also offer hands-on multicloud training for builders. We provide prescriptive guidance and professional services to set up and operate a Cloud Center of Excellence (CCOE) for a multi-cloud environment.

2.3 Service Granularity & Measurement

Metrics and Reporting

AWS provides flexible metric reporting through Amazon CloudWatch with granularity options tailored to different monitoring needs. The service offers two primary monitoring tiers: Basic Monitoring, which collects metrics at 5-minute intervals at no additional charge, and Detailed

Monitoring, which provides enhanced 1-minute interval metrics for an additional fee. Service-specific granularity varies across the AWS environments—Amazon S3 offers daily storage metrics for free while providing optional 1-minute request metrics at additional cost. For details, refer to ttps://docs.aws.amazon.com/AmazonS3/latest/userguide/cloudwatch-monitoring.html

CloudWatch's tiered data retention policy keeps 1-minute data points for 15 days, 5-minute data for 63 days, and hourly metrics for 455 days (15 months), allowing both real-time operational monitoring and long-term trend analysis. For details, refer to https://docs.aws.amazon.com/ec2/latest/devguide/monitor.html This multi-tiered approach allows organizations to balance monitoring precision with cost considerations across their AWS infrastructure.

Additionally, monitoring services include Amazon EventBridge for viewing AWS Backup events, AWS CloudTrail for monitoring API calls, and Amazon SNS for event notifications

Billing and Usage

AWS offers comprehensive billing transparency through multiple exportable formats and proactive notification systems. The AWS Cost and Usage Report delivers detailed billing data in CSV format, downloadable directly from the Billing Console or automatically delivered to an Amazon S3 bucket for integration with analytics tools. AWS Cost Explorer enables viewing and analyzing costs and usage with up to 12 months of historical data and 12-month forecasting capabilities.

For usage monitoring, AWS provides multi-layered notification capabilities through CloudWatch Billing Alarms that trigger when spending approaches or exceeds predefined thresholds. These alerts can be delivered via email, SMS, or through other AWS services. AWS Budgets further enhances cost control by setting custom budgets with email or SNS notifications when thresholds are exceeded, while AWS Cost Explorer generates downloadable reports with customizable filters for detailed spending analysis. Together, these tools provide financial visibility and proactive cost management across all AWS services.

Service Level Agreements

AWS maintains Service Level Agreements (SLAs) for all its services, documenting specific up-time commitments and service credit policies. The central repository for all AWS SLAs is accessible at https://aws.amazon.com/legal/service-level-agreements/, where customers can review detailed terms for each service.

3 Pricing & Commercial Model

AWS offers comprehensive pricing structures for diverse organizational needs while supporting UK government procurement requirements through the G-Cloud 15 framework. The pricing model encompasses multiple options from pay-as-you-go flexibility to long-term commitment savings, all available in GBP currency with transparent cost optimization tools.

Pricing Structure Overview

AWS's foundational pricing model operates on a pay-as-you-go basis, allowing customers to pay only for resources consumed without long-term commitments. On-Demand Instances enable payment for compute capacity by the hour or second with no long-term commitments, transforming large, fixed costs into smaller variable costs. This approach allows organizations to adapt to changing business needs without overcommitting budgets and improves responsiveness to changes.

Most workloads start as on-demand, but alternative pricing constructs can provide additional discounts based on utilization and capacity needs. AWS provides multiple pricing models to optimize costs based on workload characteristics and organizational requirements.

Savings Plans Framework. AWS Savings Plans serve as a primary cost optimization tool, offering flexible pricing models that provide significant savings compared to On-Demand pricing. AWS offers multiple Savings Plans types: Compute Savings Plans provide up to 66% savings and automatically apply to EC2, Lambda, and Fargate usage regardless of instance family, size, region, or operating system. EC2 Instance Savings Plans offer the lowest prices with savings up to 72% for commitment to specific instance families in a region. Database Savings Plans provide up to 35% savings on AWS database services with 1-year commitments. SageMaker Savings Plans offer up to 64% savings on Amazon SageMaker usage.

Volume Discounts and Enterprise Benefits

AWS provides volume discounts as usage grows, with customers able to take advantage of up to 72% discounts when committing to Savings Plans or Reserved Instances for predictable usage. Purchasing large numbers of Amazon EC2 Reserved Instances in an AWS Region provides discounts on upfront fees and hourly fees for future RI purchases. Enterprise customers benefit from consolidated billing through AWS Organizations, allowing multiple accounts to benefit from volume discounts and shared Reserved Instance benefits.

Pay-as-You-Go Availability

AWS offers you a pay-as-you-go approach for pricing for the majority of our cloud services. With AWS you pay only for the individual services you need, for as long as you use them, and without requiring long-term contracts or complex licensing (though Reserved Instances offer optional minimum commitments for cost savings). The pay-as-you-go model helps organizations to scale resources dynamically without upfront investments, supporting the agility requirements of modern government operations.

GBP Currency Support

AWS supports billing in British Pounds (GBP) and other local currencies, allowing UK-based organizations to manage AWS costs in their preferred currency for better financial planning and reporting.

Cost Optimization Tools and Mechanisms

AWS provides comprehensive cost optimization tools including AWS Cost Explorer, which allows viewing and analyzing costs and usage with up to 12 months of historical data and 12-month forecasting. AWS Cost Optimization Hub quantifies estimated cost savings of recommendations, incorporating specific AWS pricing and discounts such as Reserved Instances and Savings Plans. AWS Compute Optimizer uses machine learning to analyze historical utilization metrics and recommend optimal resource configurations.

AWS follows established cost optimization principles that apply across nearly all environments, focusing on rightsizing resources to match actual needs, increasing elasticity by turning off resources when not needed, using optimal pricing models based on workload characteristics, optimizing storage by selecting appropriate tiers, and measuring, monitoring, and improving through cost allocation tagging and continuous monitoring.

FOCUS 1.2 Standards Commitment

AWS achieved general availability of FOCUS 1.2 on November 19, 2025, allowing standardized cost and usage data across multi-cloud environments with enhanced capabilities for invoice reconciliation, capacity reservation tracking, and improved data

normalization.[1] AWS serves as a steering committee member and contributor to the FOCUS project, having been actively involved in FOCUS specification development since joining the FinOps Foundation as a premier member in October 2023.

FOCUS 1.2 includes significant enhancements over previous versions, with 14 additional columns compared to FOCUS 1.0, supporting hourly, daily, and monthly time granularity, and new features including InvoiceId for direct correlation with AWS invoices, CapacityReservationId and CapacityReservationStatus columns for tracking reserved capacity, and enhanced integration capabilities with SaaS provider FOCUS 1.2 datasets. Organizations can access AWS Data Exports for FOCUS 1.2 through the AWS Billing and Cost Management console.

Spot Instances for Cost-Sensitive Workloads

For fault-tolerant workloads, AWS Spot Instances provide access to unused Amazon EC2 capacity at up to 90% discount compared to on-demand prices. This option enables significant cost reductions for appropriate workload types while maintaining performance requirements.

Public Visibility and Transparency

AWS maintains publicly visible pricing at https://aws.amazon.com/pricing for all services, providing transparency for accurate cost planning. The AWS Pricing Calculator allows estimation of costs for running AWS services, supporting informed procurement decisions. This transparency aligns with G-Cloud 15 requirements for clear pricing visibility and supports competitive procurement processes.

4 Operational Capabilities

4.1 Service Management

AWS provides service management capabilities through various integrated tools and services that help with backups, scaling, monitoring, and cost management.

Backups and Recovery

AWS Backup serves as the centralized, fully managed service that automates data protection across AWS services and hybrid workloads. Using the AWS global infrastructure, AWS Backup is designed to achieve durability of 99.999999999% (11 nines). The service provides centralized console management, automated backup scheduling, backup retention management, and comprehensive monitoring and alerting capabilities.

AWS Backup integrates with Amazon CloudWatch to provide monitoring capabilities, with dashboard displays showing job metrics by count and customizable time frames ranging from 1 hour to 1 week. The monitoring system includes Amazon EventBridge for viewing AWS Backup events, AWS CloudTrail for monitoring API calls with source IP and user identification, and Amazon SNS for subscribing to backup, restore, and copy events.

AWS Backup Audit Manager provides continuous evaluation of backup activity and generates audit reports that help demonstrate compliance with regulatory requirements. Third-party auditors assess the security and compliance of AWS Backup as part of multiple AWS compliance programs, including SOC, PCI, FedRAMP, HIPAA, and others.

[1] FOCUS (FinOps Open Cost and Usage Specification) is an open specification supported by the FinOps Foundation that standardizes cost and usage data.

Amazon Data Lifecycle Manager automates the creation, retention, and deletion of Amazon EBS snapshots, helping you protect valuable data through regular backup schedules, maintain compliance with retention requirements, and reduce storage costs by removing outdated backups. This service, combined with Amazon CloudWatch and Amazon CloudTrail, provides a complete backup solution for EBS volumes at no additional cost.

For Windows environments, you can create VSS application-consistent snapshots using AWS CLI, PowerShell, or the AWSEC2-ManageVssIO SSM Document, which helps ensure database integrity during backups. AWS also supports manual snapshot management, though automated solutions are recommended. For details, refer to https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/application-consistent-snapshots.html

Scaling Capabilities

AWS provides multiple scaling options. Amazon EC2 Auto Scaling maintains application availability by automatically adding or removing EC2 instances according to the conditions you define. It can detect when instances or applications are unhealthy and replace them automatically to maintain availability.

For Spot Fleet instances, AWS offers automatic scaling through integration with Amazon CloudWatch and Application Auto Scaling. You can implement step scaling policies that respond to CloudWatch alarms or target tracking policies that maintain specific metric values. AWS Lambda also supports automatic scaling with provisioned concurrency. Using Application Auto Scaling, you can create target tracking scaling policies that adjust provisioned concurrency levels based on utilization metrics.

Usage Monitoring

Amazon CloudWatch is the primary service for monitoring AWS resources and applications. It collects and tracks metrics, monitors log files, sets alarms, and automatically reacts to changes in your AWS environment.

The CloudWatch dashboard shows current alarms and status, graphs of alarms and resources, and service health status. It allows you to graph monitoring data to troubleshoot issues, search and browse metrics, create and edit alarms, and view at-a-glance overviews of your resources. For EC2 instances, CloudWatch metrics can be aggregated by Auto Scaling group, allowing you to view statistics across all instances in a group.

Overspend Alerts

AWS Budgets allows organizations to set custom budgets to track costs and usage, with alerts triggered when actual or forecasted costs exceed defined thresholds. AWS Cost Anomaly Detection uses machine learning to continuously monitor cost and usage patterns, automatically learning normal spending behavior and alerting only when genuine anomalies occur. Amazon CloudWatch creates billing alerts that notify users when AWS service usage exceeds defined thresholds, with email notifications sent when usage exceeds specified amounts. The system analyzes spending patterns within an organisational context, calibrating anomaly detection according to how AWS accounts are organized by department, project, or functional area.

Organisations can implement cost control systems focusing on four main components: granting appropriate permissions to users, budgeting spend with custom thresholds, monitoring and analyzing cost progression toward limits, and taking actions to reduce unintentional costs.

4.2 Data Protection

Data Protection Between Buyer and Supplier

AWS implements comprehensive data protection measures across its services to secure data in transit, at rest, and during processing. These measures are part of AWS's shared responsibility model, where AWS secures the infrastructure while customers maintain responsibility for their data security configurations.

AWS data encryption capabilities span three critical protection states: encryption at rest, encryption in transit, and encryption in use. AWS Key Management Service integrates with over 120 AWS services, providing encryption capabilities across compute, storage, database, and networking services. All traffic between AWS data centers is transparently encrypted at the physical layer before it leaves AWS-secured facilities. All traffic within the AWS network is encrypted by default, and network traffic within virtual private clouds cannot be spoofed or sniffed. AWS CloudHSM provides dedicated HSM instances validated at FIPS 140-2 Level 3, critical for demonstrating compliance with security and privacy regulations including PCI DSS, GDPR, HIPAA, and FedRAMP.

Data Protection Within Supplier Network

AWS implements defense-in-depth encryption strategies across multiple layers including physical and perimeter security with transparent encryption between data centers, internal network encryption by default, host-level security through the AWS Nitro System, and data protection. The AWS Nitro System provides enhanced security through hardware-based isolation and dedicated security chips.

AWS supports 143 security standards and compliance certifications, including PCI-DSS, HIPAA/HITECH, FedRAMP, GDPR, FIPS 140-3, and NIST 800-171, helping customers satisfy compliance requirements around the globe. For UK public sector organizations, AWS is compliant with NHS Digital Data Security and Protection Toolkit, ISO 27001/27017/27018, Cyber Essentials Plus, GDPR/Data Protection Act 2018, and NCSC Cloud Security Principles.

Connected Public Sector Networks

AWS supports UK government networks through dedicated connectivity options. AWS Direct Connect delivers data through a private network connection between facilities and AWS, bypassing the public internet for secure access to workloads. The AWS VPN configuration meets NCSC Foundation profile requirements and can be used with systems running at OFFICIAL classification level. For NHS organizations specifically, AWS is connected to HSCN through several partner organizations, allowing NHS Trusts to access AWS Cloud services through HSCN if desired. Additionally, services such as AWS PrivateLink provide private connectivity between VPCs and AWS services without exposing traffic to the public internet, reducing exposure to brute force and DDoS attacks.

The AWS EU (London) Region has achieved Public Services Network (PSN) assurance, allowing UK public sector organizations to connect through PSN-certified AWS Direct Connect partners. This certification demonstrates that AWS meets PSN's stringent requirements for handling UK OFFICIAL classified data. AWS infrastructure in the London Region is PSN-assured, so customers don't need to include AWS infrastructure in their own PSN certification process.

5 Support & User Enablement

User Support

The AWS Support Center serves as the cornerstone of our support infrastructure, where we provide a unified portal for users to access documentation, submit support tickets, and monitor the health of AWS services in real-time. AWS offers a tiered support model that scales with business needs. This includes Business Support+, Enterprise Support, and Unified Operations premier support plans. The three primary support plans are designed to accelerate cloud operations. This is complemented by AWS community forums at https://repost.aws/ where developers and architects freely exchange knowledge, troubleshooting tips, and implementation strategies.

Business Support+ serves as the minimum recommended plan for production workloads, providing response times of less than 30 minutes for business/mission-critical system issues, less than 1 hour for production system issues, and less than 4 hours for production system impairment. Pricing follows a tiered structure with a minimum of £29 per month per account or percentage-based fees: 9% of monthly AWS charges up to £10K, 7% from £10K-£80K, 5% from £80K-£250K, and 3% over £250K.

Enterprise Support is recommended for business-critical workloads across organizations seeking expert guidance offering response times as fast as 15 minutes for business/mission-critical issues. This plan includes designated Technical Account Managers who combine deep AWS knowledge with understanding of business objectives, plus proactive services including AWS Well-Architected Reviews, architectural guidance, and planned events support. Pricing requires a minimum of £5K per month or tiered percentages: 10% up to £150K, 7% from £150K-£500K, 5% from £500K-£1M, and 3% over £1M.

Unified Operations targets mission-critical workloads requiring enhanced resilience and application-specific expertise, providing response times as fast as 5 minutes from Incident Management Engineers for business-critical issues. This premium tier includes 24/7 workload monitoring, designated Domain Specialist Engineers, and Incident Management Engineers, with AWS Countdown Premium and AWS Incident Detection and Response included at no additional cost. Pricing requires a minimum of £50K per month or tiered percentages: 10% up to £1M, 6% from £1M-£5M, and 5% over £5M.

All Premium Support plans include unlimited 24/7 contextual recommendations through AI-powered troubleshooting, providing instant, context-aware assistance that understands specific customer environments. Support channels encompass 24/7 phone, web, chat, and email access, AWS Support App in Slack integration, unlimited cases and contacts, third-party software support, and AWS Support API access. All plans are billed monthly with no long-term contracts, though customers must wait a minimum of 30 days to cancel subscriptions.

User Support Accessibility

AWS provides multiple accessible communication channels to accommodate different user needs and preferences. The AWS Command Line Interface serves as a critical accessibility accommodation, enabling users with disabilities to access services through assistive technology such as Job Access with Speech (JAWS), Nonvisual Desktop Access (NVDA), and alternative input devices. The AWS Support App in Slack integration allows users to manage support cases directly within Slack without requiring separate console access. Amazon Q Developer provides AI-powered conversational assistance accessible from anywhere in the AWS Management Console, enabling users to create support cases and access help through natural language interactions.

AWS demonstrates strong commitment to accessibility through systematic evaluation against WCAG 2.1 and 2.2 Level AA standards, Section 508, and EN 301 549. AWS has received Accessibility Conformance Reports for over 160 AWS services, with comprehensive coverage across the service portfolio. These reports are available through AWS Artifact, providing customers with comprehensive accessibility compliance documentation.

AWS uses accessibility testing including automated tools to catch defects such as missing alternative text for graphics, missing form field labels, missing page titles, and incorrect semantic markup. Manual accessibility tests are performed with assistive technologies including screen readers, screen magnifiers, and voice recognition software to verify compatibility.

AWS gives our developer communities curated learning content to help them create accessible experiences for customers. We provide a powerful suite of in-person and on-demand, self-paced training sessions for our staff. These sessions range from introductory content and accessibility basics to more advanced accessibility tasks, such as how to use ARIA.

Skill Development and Training

AWS offers training and certification programs that build cloud fluency across organizations. The AWS Training and Certification portal offers learning paths tailored to specific roles—from cloud practitioners to specialized architects and developers—ensuring that each professional receives relevant education for their responsibilities. For more information, refer to https://aws.amazon.com/training/

Free digital training through AWS Skill Builder subscriptions provides an entry point for those new to AWS, with hundreds of on-demand courses covering foundational concepts and service-specific implementations. As skills advance, instructor-led training delivers deeper insights through hands-on labs and expert guidance. For organizations seeking to validate their teams' expertise, AWS Certification programs offer industry-recognized credentials that benchmark skills against established standards.

Beyond formal training, AWS enables continuous learning through resources like AWS Workshops, which provide guided tutorials for implementing specific solutions, and AWS Well-Architected Labs that teach best practices for secure, high-performing architectures. The annual re:Invent conference and regional summits create opportunities for in-person learning and networking with AWS experts and community members.

Technical Support

To address common technical challenges, AWS provides self-help resources such as troubleshooting guides that offer step-by-step solutions. This includes issues such as Amazon EC2 instance connectivity problems, AWS Lambda function errors, or Amazon S3 access denials.

When self-help resources aren't sufficient, support cases raised via AWS Support Center at https://console.aws.amazon.com/support connect customers directly with AWS engineers who are experts in specific services. The case management system allows for detailed documentation of issues, secure file sharing, and ongoing communication until resolution is achieved. For Business and Enterprise customers facing critical production issues, AWS's elevated support process engages senior engineers and service teams to drive rapid resolution.

AWS has also built proactive support mechanisms that identify and address potential issues before they impact customers. Systems like AWS Health Dashboard provide personalized alerts about events that might affect infrastructure, while service teams send notifications about important changes such as runtime deprecations or security patches.

Advanced Support Features

Customers can access AWS DevOps Agent (preview), an AI-powered frontier agent that resolves and proactively prevents incidents, working alongside teams to investigate issues and engage AWS Support with one click for faster resolution. Enterprise Support includes proactive services delivered by AWS experts, helping review cloud operations health, optimise costs, and scale workloads efficiently through workload reviews, best practices workshops, and deep dives.

Technical Account Managers (TAMs) lead Well-Architected Reviews, provide architectural guidance across security, reliability, performance efficiency, cost optimisation, and operational excellence, and bring in AWS subject matter experts when needed. TAMs execute critical events with confidence through AWS Countdown, optimize cloud costs through FinOps guidance, and build team expertise through TAM-led workshops and AWS GameDays.

6 Service Continuity & Exit

6.1.1 Business Continuity

Disaster Recovery Arrangements

AWS has designed its infrastructure and services with disaster recovery at their core, providing customers with multiple options to build resilient applications that can withstand various failure scenarios. The AWS global infrastructure is divided into Regions and Availability Zones, creating a foundation for effective disaster recovery strategies.

AWS offers several disaster recovery models that customers can implement based on their recovery time objectives (RTO) and recovery point objectives (RPO):

- **Backup & Restore**: The simplest approach where data is backed up to Amazon S3 and can be restored when needed. This approach is suitable for workloads with less stringent RTO requirements.
- **Pilot Light**: Core components of the system are always running in the cloud while the rest remain turned off until needed, allowing for faster recovery than backup-restore while minimizing costs.
- **Warm Standby**: A scaled-down version of a fully functional environment is always running, ready to scale up when disaster strikes, providing even faster recovery times.
- **Multi-Site Active/Active**: The full production environment is replicated across multiple AWS Regions, providing the fastest recovery but at higher cost.

AWS provides numerous services to support DR strategies, including AWS Backup for centralized backup management, Amazon S3 for durable object storage, and AWS Elastic Disaster Recovery (formerly CloudEndure Disaster Recovery) for rapid recovery of on-premises and cloud-based applications. For more details, refer to AWS Disaster Recovery: https://aws.amazon.com/disaster-recovery/

Service Resilience Measures

AWS builds resilience into its services through a multi-layered approach that combines infrastructure design, service architecture, and operational excellence:

Infrastructure Resilience

The AWS global infrastructure includes multiple geographically isolated Regions, each containing multiple Availability Zones (AZs). These AZs are physically separated data centers with independent power, cooling, and networking, but connected with high-bandwidth, low-latency networking. This design allows customers to architect applications that automatically fail over between AZs without interruption.

AWS also provides edge locations through Amazon CloudFront and AWS Global Accelerator that bring content delivery and network optimization closer to end users, improving performance and resilience.

Service Architecture

AWS services are designed with built-in redundancy and fault tolerance:

- **Amazon S3** provides 99.999999999% (11 nines) durability by automatically storing data redundantly across multiple devices and facilities.
- **Amazon RDS** offers multi-AZ deployments that automatically replicate database updates to a standby instance in a different AZ, with automatic failover during planned maintenance or instance failure.
- **Amazon DynamoDB** automatically replicates data across multiple AZs in a Region, with global tables providing multi-region replication.
- **Amazon Route 53** offers health checks and DNS failover to route traffic away from unhealthy endpoints.

Operational Excellence

AWS implements rigorous operational practices including:

- Continuous monitoring of service health
- Automated systems that detect and mitigate issues
- Regular game days to test resilience
- Post-incident analysis to drive continuous improvement

AWS also provides detailed Service Level Agreements (SLAs) for its services, documenting service availability commitments and service credit policies. For more details, refer to AWS Global Infrastructure at https://aws.amazon.com/about-aws/global-infrastructure/

Backwards Compatibility

AWS maintains a strong commitment to backwards compatibility, helping to ensure that customer applications and workloads continue to function as AWS services evolve. Our time-bound backwards compatibility provides flexibility for customers to plan and implement upcoming changes. This compatibility stops working beyond specified cut-off dates that are announced in advance. AWS maintains backwards compatibility by auto-assigning default values to fields required in new versions, ensuring existing integrations don't fail when new mandatory fields are introduced. For example, AWS maintains backward compatibility for legacy endpoints where applicable, such as Amazon S3's support for both path-style and virtual-hosted-style URL requests.

AWS's approach is guided by several key principles:

API Stability

AWS prioritizes API stability, meaning that once an API is released, AWS works to help ensure that applications built using that API continue to function even as the underlying

service evolves. When changes are necessary, AWS typically introduces new API versions rather than modifying existing ones, allowing customers to migrate at their own pace.

Deprecation Policies

When AWS must deprecate features or services, it follows a transparent process:

- **Advance Notice**: AWS provides significant advance notice before deprecating features, typically 12+ months for major changes.
- **Migration Paths**: AWS develops and documents clear migration paths to newer alternatives.
- **Transition Tools**: AWS often provides tools to assist with transitions, such as the EC2-Classic Network Migrator for moving from EC2-Classic to VPC.

Extended Support

AWS provides extended support for older versions of software with services like Amazon EKS, Amazon RDS, Amazon Aurora, and Amazon ElastiCache. Amazon EKS supports Kubernetes versions for 14 months under standard support, followed by an additional 12 months of extended support. Amazon RDS Extended Support allows customers to continue running databases on major engine versions past the end of standard support date for up to 3 years. This extended support provides critical security updates, bug fixes, and continued technical support for legacy versions, giving organizations additional time to plan and execute upgrades at their own pace.

Service Evolution

As AWS services evolve, we follow these practices:

- **Additive Changes**: New features are added without disrupting existing functionality.
- **Versioned Services**: For significant changes, AWS may introduce new versions (e.g., Lambda runtime versions) while continuing to support older versions.
- **Compatibility Layers**: When necessary, AWS provides compatibility layers to bridge old and new implementations.

Communication Channels

AWS communicates service changes through multiple channels:

- **AWS Personal Health Dashboard**: Provides alerts about changes that might affect your specific AWS resources.
- **AWS Service Health Dashboard**: Shows the general status of AWS services.
- **Release Notes and Documentation**: Details changes, new features, and migration guidance.
- **Deprecation Announcements**: Formal notifications about feature or service deprecations.
- **AWS Blog**: Provides detailed information about significant changes and best practices for adapting.

For more details, refer to AWS service health at https://status.aws.amazon.com/

Comprehensive Business Continuity

AWS's approach to business continuity extends beyond individual service resilience to provide customers with a complete toolkit for application availability:

- **AWS Well-Architected Framework**: Provides architectural best practices for designing reliable, secure, efficient, and cost-effective systems.
- **AWS Resilience Hub**: Helps customers assess and improve application resilience by defining resilience policies and validating that applications meet availability goals.
- **AWS Fault Injection Simulator**: Allows controlled chaos engineering experiments to test application resilience.
- **AWS Business Continuity Plan**: AWS maintains its own comprehensive business continuity plans, regularly tested through simulations and exercises.

6.1.2 By combining these capabilities with AWS's global infrastructure and service-specific resilience features, customers can build applications that maintain availability even during significant disruption events. For more details about AWS Resilience refer to: https://aws.amazon.com/resilience/Exit Planning (Lot 1a Specific)

Data Access Period

AWS provides a 90-day post-closure period for data access after contract termination. During this window, customers can access billing information, contact AWS Support, and reopen their account to retrieve remaining content. AWS retains all customer content for 30 days post-termination, allowing data retrieval without transfer fees (up to 100 GB monthly, with larger transfers eligible for fee waivers through Support approval).

Exit Procedures: Clear Decision Points and Buyer Involvement

AWS provides structured exit planning procedures with a 90-day post-closure period for data access, specific termination procedures, and pricing considerations during exit. Only the AWS account root user can close an AWS account, with account closure serving as notice of termination of the AWS Customer Agreement. The account closure process requires signing in as the root user to the account settings page, reading and understanding account closure guidance, and selecting Close Account in the dialog box. AWS provides a 90-day post-closure period during which customers can view past billing information and access AWS Support. Content and services that were not deleted or terminated before account closure remain on AWS hardware during this period.

For our enterprise customers, our Technical Account Managers can help facilitate structured exit planning with defined milestones and decision points throughout the process. These typically include initial notification, data migration planning, application transition sequencing, resource decommissioning schedules, and final account closure. For more information about AWS Enterprise Support, refer to https://aws.amazon.com/premiumsupport/plans/enterprise/

Terms During Exit: Pricing and Terms During Exit Period

AWS waives data transfer out charges for eligible customers when moving outside of AWS, available to all AWS customers globally and from any AWS Region. Customers moving up to 100 GB per month can transfer their data without fees. Customers transferring more than 100 GB need to contact AWS Support to review fee waiver requests, with AWS providing credits for migrated data upon approval.

Billing for on-demand charges stops during the post-closure period and is charged at the beginning of the next month. Customers remain responsible for outstanding fees and charges before closure. Subscription charges may continue after account closure. If customers reopen their account they might be charged for the cost of running AWS services during the post-closure period. AWS does not require minimum commitments or long-term contracts, though Reserved Instances offer optional minimum commitments. AWS provides

tools and services to help migrate from AWS, minimizing vendor lock-in. Upon termination of contract, AWS will not remove customer content from systems for 30 days following the termination date and will allow content retrieval if all amounts due under the customer agreement have been paid.

Our AWS Cost Explorer and AWS Budgets services remain available during exit to help you monitor and control spending as resources are decommissioned. For customers with specific exit pricing concerns, our sales teams can help incorporate customized terms regarding pricing stability, professional services assistance, and resource commitment management during your transition period. For more information about AWS pricing, refer to https://aws.amazon.com/pricing/

Exit Plan Provision: Timeline for Providing Exit Plan

The specific timeline requirements for an exit plan vary based on the contract type and customer requirements. Our enterprise engagement model supports developing formal exit plans according to your requirements to include specific terms in your AWS contracts. When timelines are specified contractually, our Enterprise teams work with you to deliver comprehensive exit plans that include current state assessment, migration strategies for different workloads, realistic timeframes, resource requirements, and risk mitigation approaches. We can update these plans throughout the contract lifecycle to reflect your evolving AWS usage and maintain alignment with current business continuity requirements.

AWS Professional Services provides comprehensive exit support services including:

- **Exit Strategy Development**. Working with customer teams to understand requirements, constraints, and timelines
- **Technical Assessment and Documentation**. Thorough assessments of the AWS environment
- **Data Migration Planning**. Expertise in planning and executing data migrations from AWS

For more information about AWS Professional Services, refer to https://aws.amazon.com/professional-services/

7 Innovation & Technology Advancement

AWS demonstrates a strong commitment to supporting innovation and emerging technologies through several key approaches:

Commitment to Innovation

A customer-driven development is central to AWS's approach—almost 90% of the AWS technology innovation is driven by customer feedback. We start by thinking deeply about customer goals and challenges to determine the services or solutions needed to solve them. We lead by delivering new services, then rapidly iterating upon those services based on customer feedback. In 2024, AWS released over 3,000 new significant services and features, in 2023 AWS released 3,410 new services and features, up from just 80 in 2011. This sustained growth reflects AWS's commitment to delivering products and services that matter to customers.

Access to New Capabilities and Features

AWS provides customers with immediate access to the latest technologies without requiring recapitalization of investments. We offer more than 240 cloud services spanning compute, storage, networking, databases, analytics, AI/ML, IoT, security, and emerging technologies

like quantum computing. We pioneered several industry-first innovations, including serverless computing with AWS Lambda (2014), Amazon SageMaker for machine learning, and ARM-based instances powered by AWS Graviton processors (2018). AWS Support provides immediate expertise in new services so customers can use them without waiting to build internal expertise.

AWS's innovation directly impacts customer business outcomes, with services such as Amazon Bedrock, Amazon Q, AWS Transform, and Amazon SageMaker allowing businesses to drive growth and generate value. Customers that migrate workloads to AWS can reduce total cost of ownership by up to 40% while gaining access to continuously improving capabilities.

Technology Roadmap Approach

AWS's roadmap approach is characterized by flexibility and customer responsiveness. We maintain "evergreen" customer agreements that allow AWS to remain nimble and cater to customers' ever-changing demands and requests. This constant pace of innovation—much of which comes from direct customer suggestions—is one of AWS's principal advantages.

We commit to providing at least 12 months' prior notice in the rare event AWS decides to discontinue a service or material functionality, helping customers plan for any required transitions.

Evergreen/Continuously Improving Ways of Working

AWS operates with a continuous improvement philosophy across multiple dimensions:

- **Operational excellence**: AWS continually improves existing services and frequently adds new services, helping customers maintain state-of-the-art IT infrastructure
- **Incident management**: Lessons learned from previous incidents inform improvements in response plans and workload architecture, driving a continuous improvement cycle to improve workload resiliency
- **Cost optimisation**: AWS has reduced prices 161 times since launching in 2006, passing efficiency gains and economies of scale to customers
- **Sustainability**: AWS relentlessly innovates its infrastructure design, build, and operations to make progress towards net-zero carbon by 2040 and being water positive by 2030

Additionally, our development, test, and production environments follow secure software development practices with security risk reviews prior to launch, allowing continuous improvement while maintaining security standards. Our established policies define roles, responsibilities, and classifications for managing changes to production environments, with development, test, and production environments logically separated to reduce risks of unauthorised access or change. This systematic approach to change management helps ensure that innovation can be delivered safely and reliably while maintaining operational excellence.

8 Standards Compliance & Sustainability

Environmental Initiatives and Data Centre Energy Efficiency

AWS operates with a global Power Usage Effectiveness (PUE) of 1.15 in 2024, significantly outperforming both the public cloud industry average of 1.25 and on-premises data centers at 1.63. AWS has implemented advanced data center designs and innovative cooling technologies to achieve this efficiency. AWS unveiled new data center components in 2024,

offering 12% more compute power with improved availability and efficiency. Additionally, All AWS services and data centers in EMEA, Singapore, and Indonesia are ISO 50001 certified for energy management.

AWS has implemented cooling technologies including direct-to-chip liquid cooling systems that reduce mechanical energy consumption by up to 46% without increasing water usage. In 2024, AWS achieved a global data center water usage effectiveness of 0.15 liters per kilowatt hour, representing a 17% improvement from 2023 and a 40% improvement since 2021. AWS is 53% of the way toward its water positive goal and has 23 global water replenishment projects that returned 4.3 billion liters of water to communities in 2024.

Net Zero Alignment Strategy

As part of The Climate Pledge, Amazon (including AWS) is committed to reaching net-zero carbon emissions by 2040—10 years ahead of the Paris Agreement. Over 590 companies have joined this pledge. Key achievements include:

- 100% of electricity consumed by Amazon was matched with renewable energy sources in 2024, for the second consecutive year, achieving the 2025 target five years early
- Over 600 global renewable energy projects across 29 countries
- Amazon's Climate Pledge Fund investment of $2 billion into clean technologies
- 4% reduction in carbon intensity in 2024 compared to 2023

AWS supports significant carbon footprint reductions for customers migrating to the cloud. According to a 2024 Accenture study, the AWS Cloud is up to 4.1 times more efficient than on-premises environments. AWS infrastructure can reduce carbon footprint by up to 99% for optimized workloads compared to on-premises environments.

Greening Government ICT Strategy Compliance

The UK government has established the Greening Government ICT and Digital Service Strategy 2020-2025, which includes extended mandatory reporting around carbon emissions, sustainability and social value. AWS provides tools and services that help UK government departments reduce their carbon footprint by migrating to the cloud.

- The Customer Carbon Footprint Tool uses simple visualizations to show customers their historical carbon emissions, estimate emissions avoided by using AWS instead of on-premises data centers, and review forecasted emissions based on their current use.
- The AWS Sustainability Data Fabric creates a unified data management system for environmental, social, and governance of information across organisations.
- The AWS Well-Architected Framework includes a Sustainability Pillar focusing on minimizing environmental impacts of cloud workloads.
- AWS Data Exchange makes it easy to find, subscribe to, and use third-party sustainability-related data in the cloud, including data sets made available via the Amazon Sustainability Data Initiative and the Open Data Sponsorship Program

Technology Code of Practice Alignment

AWS supports the UK Government's Technology Code of Practice (TCoP) through compliance capabilities, with over 6,500 government agencies already using AWS across all classification levels. AWS supports 143 security standards and compliance certifications, including the GDPR/Data Protection Act 2018, NHS Digital Data Security and Protection Toolkit, ISO 27001/27017/27018, Cyber Essentials Plus, and NCSC Cloud Security Principles. AWS provides TCO calculators and cost optimization tools that help government

departments understand and optimize their cloud spending. AWS supports open standards and interoperability, allowing government departments to avoid vendor lock-in and enabling data portability, as required by the TCoP

ISO 27001 Certification

AWS is certified under the ISO 27001 standard. ISO 27001 is a widely adopted global security standard that outlines the requirements for information security management systems. It provides a systematic approach to managing company and customer information that is based on periodic risk assessments. 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.

AWS has established a formal programme to maintain the certification. More information regarding AWS's ISO 27001 certification can be found at http://aws.amazon.com/compliance/iso-27001-faqs/.

NCSC UK Cloud Security Principles

In 2016, National Cyber Security Centre (NCSC) UK published the Cloud Security Collection documents for public sector organisations that are considering the use of cloud services for handling information classified as OFFICIAL. The collection of guidance documents aims to help public sector organisations make informed decisions about cloud services and choose a cloud service that balances business benefits and security risks.

AWS is aligned with National Cyber Security Council Cloud Security Principles. Customers can find operational standard practices for these principles in the AWS Documentation portal here: https://docs.aws.amazon.com/config/latest/developerguide/operational-best-practices-for-ncsc.html.

We provide tools that allow customers control of access, movement, and security of their own data. These include:

- AWS Identity and Access Management (IAM), which allows customers to specify who or what can access services and resources in AWS, centrally manage fine-grained permissions, and analyse access to refine permissions across AWS
- AWS DataSync, the AWS Transfer Family and other services to move data
- A wide range of security tools and recommended practice guides to secure data.

The AWS Compliance Program helps customers to understand the controls in place at AWS to maintain security and compliance of the AWS Cloud. By tying together governance-focused, audit-friendly service features with applicable compliance or audit standards, AWS Compliance Enablers build on traditional programmes, helping customers to establish and operate in an AWS security control environment.

GDPR and processing of Personal Data

AWS offers a GDPR-compliant Data Processing Addendum (DPA), enabling customers to comply with GDPR contractual obligations. More information can be found at the following links:

- AWS GDPR Center: https://aws.amazon.com/compliance/gdpr-center/
- AWS Commitment to European Customer Data Protection: https://aws.amazon.com/compliance/eu-data-protection/
- Additional UK Addendum: https://d1.awsstatic.com/legal/aws-gdpr/UK_GDPR_Addendum_to_AWS_data_processing_addendum.pdf