G-Cloud 15

Amazon Web Services EMEA SARL, UK Branch

Consolidated Service Definition

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)

13 August 2026

Contents

List of Figures

How to use the AWS Service Definition Documents

To make it easier for customers to review AWS service content from the hundreds of individual AWS listings on the Digital Marketplace, AWS has grouped the definitions from its listed services into bundled Service Definition Documents that describe the features of each family of AWS Cloud services. The AWS service families are:

- **IaaS and PaaS (Lot 1a); SIaaS and PIaaS (Lot 2a); SaaS (Lot 2b)**
- AWS Professional Services (Lot 3)
- AWS Support (Lot 3)
- AWS Training (Lot 3)
- AWS Managed Services (Lot 3)

This **Consolidated AWS Services Service Definition** document describes the key features for each of the different AWS Cloud Services available to Customers on G-Cloud 15 in Lots 1a, 2a, and 2b.

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”.

Notwithstanding that AWS has combined its service definition into a consolidated document for ease of review by Customers, to access the options through a Call-Off Contract the Customer must reference each individual Digital Marketplace Service ID within the Call-Off Contract in order to enable that service as an option that can be procured under their G-Cloud 15 Call-Off Contract.

Lot 2a (Infrastructure Software as a Service - I-SaaS), Lot 2b (Software-as-a-Service) and Lot 3 (Cloud Support Services) services may be purchased as ancillary services to Lot 1 (IaaS and PaaS) offerings. The framework's Full Taxonomy of Services uses a Lot availability key where services are marked with "O" (Optional) to indicate they "may be purchased as ancillary service to Core Requirements".

This means that while Lot 2a, Lot 2b, and Lot 3 services are primarily available through their respective dedicated lots, buyers have the flexibility to procure these services alongside their Lot 1a cloud hosting purchases when they complement or support the core infrastructure and platform services. This ancillary purchasing mechanism allows buyers to create comprehensive cloud solutions by bundling infrastructure software and support services with their primary IaaS/PaaS deployments, streamlining procurement and ensuring integrated service delivery from a single supplier when appropriate.

AWS recommends that Buyers list all of the Digital Marketplace Service IDs for every service across all relevant lots in its Call-Off Contract to enable the option to switch between Services flexibly during the term. For a list of all AWS Digital Marketplace Service ID’s, please contact an AWS account representative at aws-ukps-frameworks@amazon.com.

Please note that the options or parameters selected by AWS on this framework are those that most closely align with our existing commercial services. AWS is happy to provide additional information about our services upon request. To find out more about AWS on G-Cloud and AWS Cloud services, visit us at AWS on G-Cloud UK.

Important Considerations

Customer Spend Monitoring

Customers are responsible for monitoring and managing their AWS spend. AWS Budgets provides the self-service financial management capabilities required for transparent, accountable public sector cloud spending, and enables customers to set custom budget thresholds, track actual and forecasted usage against those budgets, and configure alerts to help identify potential overspend early and maintain ongoing cost control. For efficient cost monitoring and setting up notifications for contract value thresholds, customers are required to sign up to AWS Budgets (https://aws.amazon.com/aws-cost-management/aws-budgets/).

Self-Service Reporting of Buyer Environment

AWS customers have several self-service options for reporting on and monitoring their environment. All these tools are accessible through the AWS Management Console, with many offering API access for programmatic integration into custom reporting solutions or business intelligence tools. See Section 4.3 Service Granularity & Measurement for more detail.

Collaborative Working Principles

The AWS Shared Responsibility Model (See Section 3 “Shared Responsibility Model”) and customer-controlled access management apply universally across all AWS services. Customers retain complete ownership and control of their AWS environment and data. Customers have full authority to grant, manage, and revoke access to their environment for any third parties, by using AWS Identity and Access Management (IAM) tools. All access is governed by customer-defined policies implementing least-privilege principles, with comprehensive logging and monitoring capabilities. This model ensures that collaborative working arrangements are entirely under customer control, with customers determining who can access their environment, what they can access, and for how long.

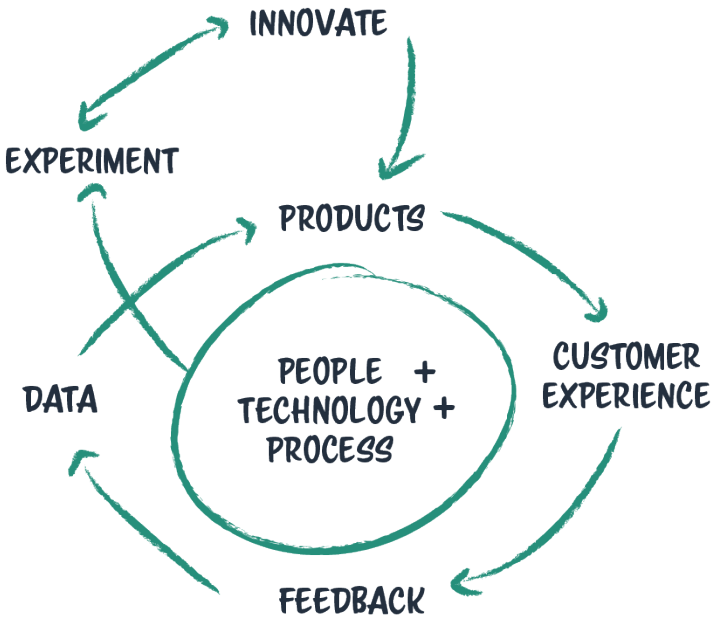

CROSS-SERVICE DEFINITIONS

1 Security Assurance

1.1 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.

Service Manual Phases Understanding

AWS supports the UK Government Service Manual phases through cloud services, proven frameworks, and ecosystem support:

- **Discovery**: Rapid prototyping and experimentation without upfront infrastructure investment.
- **Alpha**: Scalable infrastructure for testing prototypes. The AWS Cloud Adoption Framework guides cloud strategy, while the AWS Well-Architected Framework provides design principles across six pillars (operational excellence, security, reliability, performance efficiency, cost optimization, sustainability).
- **Beta**: Iterative development with scalable services, comprehensive security compliance, and expert guidance through AWS Support and Solutions Architects.
- **Live**: Production-grade infrastructure with high availability, security, and performance. AWS Partner Network provides specialized government expertise with extensive compliance programs.
- **Retirement**: Service decommissioning through data migration tools, secure deletion, and the Product Lifecycle Page for governance visibility.

The AWS Cloud Adoption Framework and AWS Well-Architected Framework parallel the GDS Service Manual approach—emphasizing iterative development, user needs, security by design, and continuous improvement. These frameworks ensure services meet GDS standards while leveraging cloud-native scalability and resilience.

1.2 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.

1.3 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

1.4 British Standard 7858:2019

Buyers selecting AWS Services and expressly requiring AWS conformity to BS7858:2019 acknowledge that AWS scopes BS7858:2019 compliance to all prospective employees with potential physical access to the 'data layer' zones within datacentres and those who are directed by the Buyer to access Buyer Data such as Technical Account Managers ("TAMS").

A list of TAMs shall be provided to the Buyer by the Supplier prior to the Start date of the Call-Off Contract and the Buyer shall only contact the listed TAMs in relation to Buyer Data during the Term of the Call-Off Contract. Buyers are obliged in accordance with the Call-Off Contract to encrypt Buyer Data when using AWS Services. Buyer should note that the Supplier does not include Supplier Staff (as defined in the Call-Off Contract) responsible for operating the AWS Services or those with logical access to encrypted Buyer Data for the purposes of its BS7858:2019 compliance.

2 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.

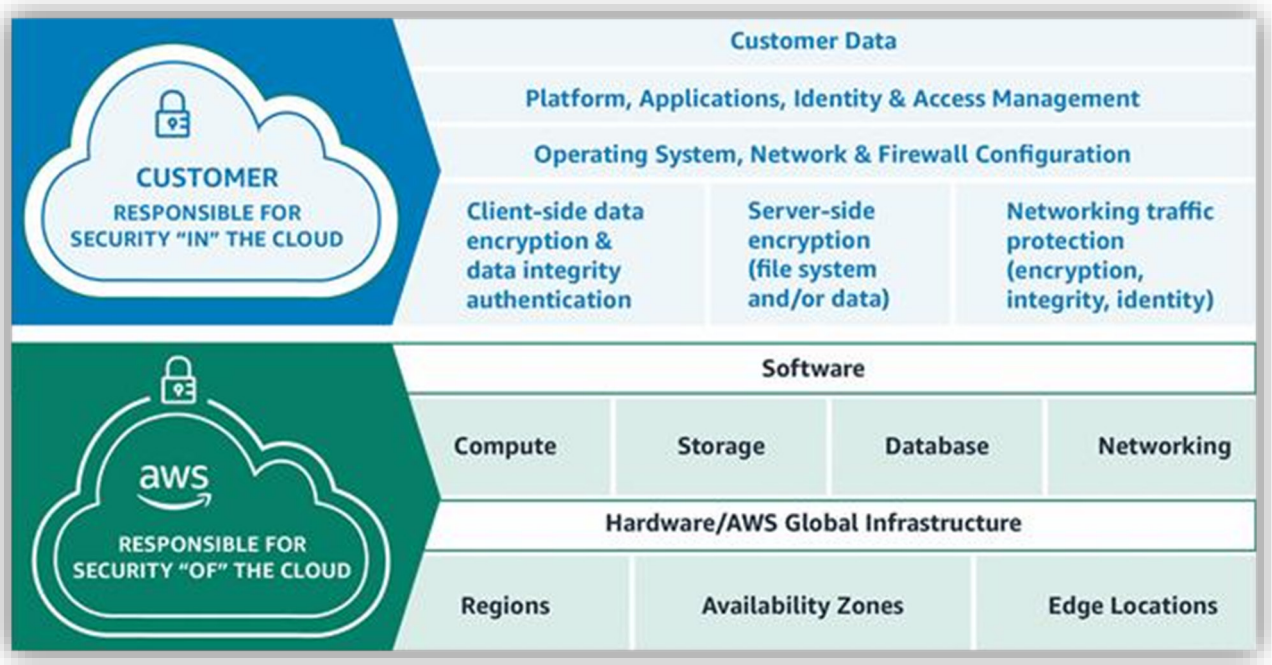

Figure 1. The Shared Responsibility Model

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.

2.1 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 centre or service. We built our data centres 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 defence, 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/.

2.2 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.

2.3 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.

2.4 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.

2.4.1 Inherited Controls

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

2.4.2 Shared Controls

Some controls apply to both the infrastructure layer (AWS responsibility) and customer layers (customer responsibility), but in 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.

2.4.3 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.

2.5 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.

2.6 Shared Responsibility and Customer Data

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

2.6.1 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.

2.6.2 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 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.

2.6.3 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.

2.6.4 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.

2.7 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

3 Technical Architecture & Service Delivery

3.1 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 centre, 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 centre 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.

3.2 Infrastructure location

AWS offers 240+ fully featured services from data centres 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.

3.3 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 centres. 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

3.4 Service Delivery & Access

3.4.1 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.

3.4.2 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, colour 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.

3.4.3 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. 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.

3.5 Service Granularity & Measurement

3.5.1 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.

3.5.2 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 analysing 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.

3.5.3 Service Level Agreements

AWS maintains Service Level Agreements (SLAs) for all its services, documenting specific uptime 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.

4 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, with transparent cost optimization tools.

The AWS Free Tier provides customers the ability to explore and try out AWS services at no cost up to specified limits for each service, operating as a no-cost trial model that enables learning and skill-building before making purchase commitments. AWS offers three distinct types of free tier offerings: Always Free services that provide perpetual access to over 30 services within specified monthly limits with no expiration, short-term trials that allow time-limited access to specific products for periods ranging from days to months, and a legacy 12-months free program for existing customers that provided free usage during the first year after account creation. For new AWS customers signing up after July 15, 2025, AWS transitioned to a credit-based Free Tier v2 program offering an initial $100 in AWS credits with the ability to earn an additional $100 by completing getting-started activities, resulting in up to $200 total credits usable across eligible services during a 6-month trial period, plus access to 30+ Always Free services with perpetual monthly limits. This credit-based program replaces traditional 12-month free offers and short-term trials while retaining Always Free offerings, with free usage calculated monthly across all regions and automatically applied to bills, enabling customers to explore AWS services, build skills, and evaluate service value prior to adoption without unexpected charges during the trial period. See http://aws.amazon.com/free/ for service-specific details.

4.1 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.

4.2 Commitment Options

AWS Savings Plans serve as a primary cost optimization tool, offering flexible pricing models that provide significant savings compared to On-Demand pricing. Further, Savings Plans offer the same savings as Reserved Instances, but unlike Reserved Instances, which require commitments to specific instance types, families, generations, and regions, Savings Plans commit to a dollar-per-hour spend rate that automatically applies to eligible usage based on the highest discount first. This added flexibility makes Savings Plans the preferred choice for businesses with evolving workloads.

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 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.

4.3 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.

4.4 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. The pay-as-you-go model helps organizations to scale resources dynamically without upfront investments, supporting the agility requirements of modern government operations.

4.5 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. Note that it is the responsibility of the buyer to select GBP as the currency in the Console.

4.6 Cost Optimization Tools and Mechanisms

AWS provides comprehensive cost optimization tools including AWS Cost Explorer, which allows viewing and analysing 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 Savings Plans. AWS Compute Optimizer uses machine learning to analyse 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. See https://aws.amazon.com/aws-cost-management/ for more detail.

4.7 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.

4.8 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. AWS Budgets ensure that customers can monitor spend effectively. This transparency aligns with G-Cloud 15 requirements for clear pricing visibility and supports competitive procurement processes.

4.9 AWS Data Transfer

Overview of Data Transfer

AWS data transfer pricing follows a tiered structure based on volume and geographic region. Data transfer into AWS from the internet is free of charge. Data transfer out from AWS to the internet (data egress) is charged based on the originating AWS region and monthly transfer volume.

UK and Europe Region Pricing: For AWS resources located in UK and European regions, data egress charges are structured as follows:

- First 100 GB per month: Free (global free tier)
- Over 10 TB per month charged according to publicly available pricing

Data transfer between AWS services within the same region (such as between Amazon S3 and Amazon EC2) is typically free of charge. However, data transfer between Availability Zones within the same region incurs charges for both incoming and outgoing traffic.

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

Global Data Egress Waiver for UK Public Sector Customers

AWS offers a **Global Data Egress Waiver (GDEW)** program that may benefit eligible UK public sector organizations, including government departments, educational institutions, and research organizations.

Program Benefits: The GDEW program waives data transfer out charges from AWS to the internet up to a maximum of 15% of total monthly AWS spending. This provides:

- **Budget predictability** for organizations with fixed procurement budgets
- **Cost savings** averaging 1-5% of monthly bills for typical usage patterns
- **Volume discounts** when using consolidated billing across multiple accounts

Eligibility for UK Public Sector: UK public sector organizations may be eligible for the GDEW program if they:

- Work in higher education, research institutions, or other eligible public sector organizations
- Are connected to AWS via a National Research and Education Network (NREN) such as JANET/Jisc (UK) or use AWS Direct Connect
- Use institutional email addresses for AWS accounts

Important Limitations: The waiver applies only to data transfer out to the internet and does **not** cover:

- Data transfers between AWS regions
- Amazon CloudFront charges
- AWS Direct Connect port speed fees (per-Gbps charges)
- VPN connection charges

How the Waiver Works: GDEW discounts appear directly on monthly bills as credits applied against Data Transfer charges, with the line-item description "Global Data Egress Waiver for DTO." If data egress costs exceed 15% of the total monthly AWS bill, customers receive a credit for the amount equal to 15% of their bill and are responsible for paying the remainder of egress charges.

Application Process for UK Public Sector Customers

Eligible UK public sector organizations can apply for the GDEW program by:

- Contacting their AWS account manager or emailing aws-data-egress@amazon.com
- Completing the GDEW application form
- Receiving an AWS Data Egress Waiver Addendum for signature via DocuSign

Organizations purchasing AWS through authorized Solution Providers may also be eligible for similar egress discount programs. UK public sector customers should consult with their AWS Solution Provider about available egress discount options.

Cost Optimization Strategies

Beyond the GDEW program, UK public sector customers can optimize data transfer costs through:

- **Architectural design**: Keeping resources within the same AWS region and Availability Zone where possible
- **Content delivery**: Using Amazon CloudFront for frequently accessed content to reduce repeated data transfers

- **Data compression**: Implementing compression techniques to reduce transfer volumes
- **VPC endpoints**: Using VPC gateway endpoints for S3 and DynamoDB to avoid internet data transfer costs

5 Operational Capabilities

5.1 Service Management

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

5.1.1 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.

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

5.1.2 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.

5.1.3 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.

5.1.4 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 behaviour 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 analyses 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 analysing cost progression toward limits, and taking actions to reduce unintentional costs.

5.2 Data Protection

5.2.1 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 centres 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.

5.2.2 Data Protection Within Supplier Network

AWS implements defence-in-depth encryption strategies across multiple layers including physical and perimeter security with transparent encryption between data centres, 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.

5.2.3 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.

6 Support & User Enablement

AWS account managers operate as part of an integrated account team to ensure customers receive the right resources at the right time throughout their cloud journey.

AWS account teams (Account Managers, SAs, Technical Account Managers etc.) provide personalised guidance to customers on using tools effectively, implementing best practices, and developing strategies to manage financial expenditure aligned with business objectives.

6.1 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. Detailed pricing available under Lot 3 of the framework.

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. Detailed pricing available under Lot 3 of the framework.

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. Detailed pricing available under Lot 3 of the framework.

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.

6.2 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. AWS 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.

6.3 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.

6.4 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.

6.5 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.

7 Service Continuity & Exit

7.1 Business Continuity

7.1.1 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/

7.1.2 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 centres 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. See https://aws.amazon.com/legal/service-level-agreements/ for up to date service-level SLA information.

For more details, refer to AWS Global Infrastructure at https://aws.amazon.com/about-aws/global-infrastructure/

7.1.3 Backward Compatibility

AWS maintains a strong commitment to backward 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/

7.1.4 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.

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/

7.2 Exit Planning (Lot 1a Specific)

7.2.1 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).

7.2.2 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/

7.2.3 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/

7.2.4 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/

8 Innovation & Technology Advancement

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

8.1 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.

8.2 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 on 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 Quick, 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.

8.3 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.

8.4 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.

9 Environmental Sustainability

9.1 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 centres at 1.63. AWS has implemented advanced data centre designs and innovative cooling technologies to achieve this efficiency. AWS unveiled new data centre components in 2024, offering 12% more compute power with improved availability and efficiency. Additionally, All AWS services and data centres 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 centre water usage effectiveness of 0.15 litres 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.

9.2 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.

9.3 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 centres, 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.

INDIVIDUAL SERVICE DEFINITIONS

Amazon API Gateway

General Information

Amazon API Gateway is a fully managed service that enables developers to create, publish, maintain, monitor, and secure APIs at any scale. It serves as the front door for applications to access backend services including AWS Lambda functions, Amazon Elastic Compute Cloud (Amazon EC2) instances, and HTTP endpoints.

Main Benefits

Amazon API Gateway helps eliminate server provisioning and infrastructure management while providing automatic scaling from few requests daily to thousands per second without warm-up limitations. Amazon API Gateway features comprehensive security with the OAuth 2.0 standard, AWS Identity and Access Management (IAM) integration, AWS WAF protection, and AWS Free Tier offering one million API calls monthly.

Key Use Cases

- **Serverless Applications:** Integration with AWS Lambda functions for automatic scaling and pay-per-use pricing
- **Microservices Architecture:** Unified access point for microservices deployed on Amazon Elastic Container Service, Amazon Elastic Kubernetes Service, or Amazon EC2 instances
- **Mobile and Web Backends:** HTTP endpoints for mobile apps and web applications with built-in CORS support
- **Real-time Applications:** Bidirectional communication for chat applications, dashboards, and multiplayer games using WebSocket APIs

Core Features

- **Multiple API Types:** Support for REST APIs, HTTP APIs, and WebSocket APIs with different feature sets
- **Integration Options:** AWS Lambda proxy and custom integrations, HTTP proxy and custom integrations, AWS service integrations, mock integrations, and private integrations
- **Security and Authorisation:** AWS IAM, AWS Lambda authorisers, Amazon Cognito user pools, OAuth 2.0, OpenID Connect, API keys, and AWS WAF integration
- **Traffic Management:** Request throttling, usage plans, quotas, caching, and burst handling with token bucket algorithm
- **Endpoint Types:** Edge-optimised (Amazon CloudFront), regional, and private endpoints with virtual private cloud integration capabilities
- **Monitoring and Observability:** Amazon CloudWatch metrics, AWS X-Ray tracing, access logging, and comprehensive API analytics
- **Developer Experience:** OpenAPI 3.0 support, SDK generation, developer portals, API documentation, and stage management
- **Data Transformation:** Request and response mapping, validation, binary media type support, and CORS configuration.

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/api-gateway/

FAQ: https://aws.amazon.com/api-gateway/faqs/

Technical Overview

Backup and Recovery Capabilities

Backup: Users can export API definitions as OpenAPI specifications, AWS CloudFormation templates stored in Amazon Simple Storage Service (Amazon S3) or Git repositories can be created manually (AWS Management Console, CLI, APIs) or automated via AWS CloudFormation with custom scheduling. Users must implement custom backup strategies using infrastructure as code practices.

Recovery: Cross-region deployment requires recreating APIs using exported definitions, blue-green deployments through stages and canary releases, and multi-region failover with Amazon Route 53 health checks. Multi-Availably Zone architecture provides regional resilience. Supports customer RPO (point-in-time configuration recovery) and RTO (sub-minute to 60 minutes) objectives.

Setup Process

Onboarding: Access through AWS Management Console, CLI, or SDKs using HTTP API quick create feature or importing OpenAPI 3.0 definitions. AWS CloudFormation templates enable infrastructure as code deployments with automated continuous integration and continuous delivery pipelines.

Offboarding: Delete APIs, stages, deployments through console, CLI, or APIs removing all associated resources. Custom domain names, VPC links require separate deletion. AWS CloudFormation stacks automate complete resource cleanup.

Data Management

- **Storage Options:** API definitions exported as OpenAPI specifications, AWS CloudFormation templates stored in Amazon S3 or Git repositories
- **Data Removal:** Delete APIs through console, CLI, or APIs removing associated resources; Amazon CloudWatch logs require separate cleanup
- **Cross-Region Replication:** Manual recreation of APIs in target AWS Regions using exported definitions; no automatic replication available
- **Encryption:** Traffic encrypted in transit; integrates with AWS WAF for additional security and API protection capabilities

Optimise Cost and Consumption

Response caching reduces backend calls and associated compute costs by serving repeated requests from cache. Request throttling with usage plans prevents unexpected cost spikes during traffic bursts. Direct AWS service integrations bypass intermediate functions for simple operations like Amazon S3 retrieval. Regional endpoint selection reduces Amazon CloudFront charges for localised traffic. Efficient API design with consolidated endpoints minimises complexity and charges.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/apigateway/latest/developerguide/welcome.html

Getting Started: https://docs.aws.amazon.com/apigateway/latest/developerguide/getting-started.html

API Reference: https://docs.aws.amazon.com/apigateway/latest/api/

Service Quotas: https://docs.aws.amazon.com/apigateway/latest/developerguide/limits.html

Amazon AppFlow

General Information

Amazon AppFlow is a fully-managed integration service that enables secure data exchange between SaaS applications and AWS services. Users can create data flows in minutes without writing code, supporting transformation, filtering, and validation with flexible scheduling options.

Main Benefits

Amazon AppFlow helps eliminate custom integration development, reducing deployment time from weeks to minutes with no code required. The fully-managed service supports over 60 native SaaS connectors, handles up to 100GB per flow, and provides enterprise-grade security with automatic encryption and scaling.

Key Use Cases

- **Data Lake Hydration:** Transfer SaaS data to Amazon Simple Storage Service (Amazon S3) for analytics and machine learning
- **Real-time Analytics:** Create event-driven flows transferring Salesforce records to Amazon Redshift for immediate analysis
- **Data Archiving:** Schedule flows for historical data transfer to Amazon S3 for long-term storage
- **Custom Integration:** Build connectors using Python and Java SDKs for private APIs and systems

Core Features

- **No-Code Integration:** Create data flows using intuitive visual interface without writing code in minutes
- **Native SaaS Connectors:** Pre-built connectors for over 60 applications including Salesforce, Google Analytics, Slack, and Zendesk
- **Data Transformations:** Map fields, concatenate data, mask sensitive information, truncate fields, and apply validations
- **Flexible Triggers:** Run flows on-demand, on-event, or on-schedule with incremental and full transfer modes
- **Enterprise Security:** Encrypt data in transit and at rest with AWS PrivateLink and AWS Key Management Service (AWS KMS) support
- **Custom Connectors:** Build connectors using Python and Java SDKs for private APIs and systems
- **AWS Integration:** Native connections to Amazon S3, Amazon Redshift, Amazon Relational Databases (Amazon RDS), AWS Glue Data Catalog, and Amazon EventBridge services
- **High Availability:** Automatic retry mechanisms, error handling, and reliability on AWS managed infrastructure.
- Pricing
- Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/appflow/

FAQ: https://aws.amazon.com/appflow/faqs/

Technical Overview

Backup and Recovery Capabilities

Backup: Amazon AppFlow does not provide native backup capabilities as a stateless managed service. Data protection relies on destination services like Amazon S3 versioning or Amazon Redshift snapshots. Users can manually backup flow configurations via AWS CloudFormation templates or API exports for recreation if needed.

Recovery: Recovery focuses on recreating flows and connections in alternate AWS Regions when necessary. The service automatically handles infrastructure failures through AWS managed resilience with automatic retry mechanisms. Supports customer RPO objectives through near real-time event-driven triggers and RTO objectives via multiple region availability for geographic distribution.

Setup Process

Onboarding: Launch via AWS Management Console with visual interface for creating connections using OAuth or API keys. Configure flows through guided setup including source selection, destination configuration, field mapping, and trigger mechanisms. Setup completes in minutes without code development.

Offboarding: Delete flows first, then remove connector profiles and connections through console or DeleteFlow API. Separately cleanup destination data in Amazon S3 buckets or Amazon Redshift tables. Remove associated AWS Identity and Access Management roles, AWS Key Management Service (AWS KMS) keys, and AWS CloudFormation stacks.

Data Management

- **Storage Options:** Native connections to Amazon S3, Amazon Redshift, and Amazon RDS PostgreSQL for comprehensive data workflows
- **Data Removal:** Delete flows first, then connector profiles and connections via console or DeleteFlow API operation
- **Cross-Region Replication:** Multiple AWS Region availability enables geographic distribution of flows for improved disaster recovery objectives
- **Encryption:** Automatic encryption in transit and at rest with AWS PrivateLink and customer-managed AWS KMS keys

Optimise Cost and Consumption

Amazon AppFlow optimises costs through incremental data transfer modes that minimise processing charges by transferring only changed records. Schedule flows during off-peak hours to reduce API usage against source application limits. Data filtering capabilities transfer only necessary records, reducing consumption. The pay-per-successful-run model avoids charges for failed attempts or idle time, helping ensure costs align with actual data transfer needs.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/appflow/latest/userguide/what-is-appflow.html

Getting Started: https://docs.aws.amazon.com/appflow/latest/userguide/getting-started.html

API Reference: https://docs.aws.amazon.com/appflow/1.0/APIReference/Welcome.html

Service Quotas: https://docs.aws.amazon.com/appflow/latest/userguide/service-quotas.html

Amazon Athena

General Information

Amazon Athena is a serverless interactive query service that enables users to analyse data directly in Amazon Simple Storage Service (Amazon S3) using standard SQL. It automatically scales to execute queries in parallel, supporting various data formats including CSV, JSON, and Parquet for business intelligence workloads.

Main Benefits

Amazon Athena helps eliminate infrastructure management overhead through its serverless architecture with automatic scaling and pay-per-query pricing. The service enables immediate analysis of Amazon S3 data without ETL processes, supports Parquet compression up to 80% versus CSV, and provides consistent performance across varying workloads while helping mitigate data silos.

Key Use Cases

- **Interactive Ad-Hoc Queries:** Run quick SQL queries on large Amazon S3 datasets for data exploration
- **Log Analysis:** Query AWS service logs, web logs, and AWS CloudTrail for security and operational insights
- **Data Lake Analytics:** Analyse staging data before warehouse loading using notebook-based analytical tools
- **Business Intelligence:** Connect with Amazon QuickSight and BI tools for creating dashboards and reports

Core Features

- **Serverless Architecture:** Zero infrastructure management with automatic scaling and pay-per-query pricing model
- **Standard SQL Support:** Full ANSI SQL compatibility with complex queries, joins, window functions, and user-defined functions
- **Multiple Data Formats:** Native support for CSV, JSON, Parquet, ORC, Avro, and Apache Iceberg table formats
- **Federated Query:** Query across multiple sources including Amazon S3, Amazon Redshift, Amazon DynamoDB, and external databases using connectors
- **Machine Learning Integration:** ML inference capabilities using Amazon SageMaker models directly within SQL queries
- **Apache Spark Support:** Run Apache Spark applications with simplified notebook experience for Python development
- **Workgroups Management:** Organise users and queries with workgroups, query result management, and cost controls
- **Security Controls:** AWS Identity and Access Management (AWS IAM) integration, encryption at rest and in transit, Amazon Virtual Private Cloud support, and fine-grained access controls
- Pricing
- Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/athena/

FAQ: https://aws.amazon.com/athena/faqs/

Technical Overview

Backup and Recovery Capabilities

Backup: AWS Glue Data Catalog metadata can be backed up and replicated across AWS Regions, while query results stored in Amazon S3 inherit S3's backup capabilities including versioning and cross-region replication. Named queries and workgroup configurations can be version-controlled and recreated using Infrastructure as Code approaches like AWS CloudFormation.

Recovery: Cross-region disaster recovery involves replicating metadata and recreating workgroups in alternate regions, with recovery procedures focusing on metadata restoration. Automatic multi- Availability Zones replication provides high availability. Supports near-zero RPO (real-time metadata replication) and RTO (minutes for Domain Name System redirection and metadata restoration).

Setup Process

Onboarding: Setup requires minimal configuration accomplished in minutes via Amazon Athena console, creating databases and defining table schemas pointing to Amazon S3 data. Workgroups can be configured for user access, query limits, and cost controls with AWS IAM integration.

Offboarding: Delete workgroups, named queries, and custom data catalogues. Manually remove query results from Amazon S3 buckets and table definitions from AWS Glue Data Catalog for complete cleanup.

Data Management

- **Storage Options:** Data stored in Amazon S3 with support for CSV, JSON, Parquet, ORC, Avro formats
- **Data Removal:** Delete workgroups, named queries, custom data catalogues, and manually remove Amazon S3 query results
- **Cross-Region Replication:** AWS Glue Data Catalog metadata and Amazon S3 query results support cross-region replication capabilities
- **Encryption:** Encryption at rest and in transit with AWS IAM integration and fine-grained access controls.

Optimise Cost and Consumption

Workgroups provide query limits and cost monitoring to prevent runaway queries. Converting data to columnar formats like Parquet or ORC reduces data scanned significantly. Effective partitioning strategies can minimise scan scope, while query result caching helps eliminate redundant processing. CTAS operations optimise data formats for long-term savings, and approximate functions control query costs through selective processing.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/athena/latest/ug/what-is.html

Getting Started: https://docs.aws.amazon.com/athena/latest/ug/getting-started.html

API Reference: https://docs.aws.amazon.com/athena/latest/APIReference/

Service Quotas: https://docs.aws.amazon.com/athena/latest/ug/service-limits.html

Amazon Aurora

General Information

Amazon Aurora is a fully managed relational database engine compatible with MySQL and PostgreSQL, delivering up to 5x MySQL and 3x PostgreSQL performance. Built for the cloud, Aurora features distributed, fault-tolerant storage that automatically scales to 128 TiB with enhanced availability.

Main Benefits

Amazon Aurora offers superior price-performance with up to 90% cost reduction compared to commercial databases while providing 5x MySQL and 3x PostgreSQL performance. The cloud-native architecture helps eliminate database administration overhead through fully managed operations, automatic scaling to 128 TiB, and innovative features like Aurora Serverless and Global Database.

Key Use Cases

- **Enterprise Applications:** High-performance databases with MySQL/PostgreSQL compatibility for mission-critical business applications
- **Web and Mobile Applications:** Scalable database solutions with automatic scaling for growing user bases
- **Data Analytics Workloads:** Business intelligence and analytics using Aurora's performance optimisation features
- **Global Applications:** Multi-Region deployment using Aurora Global Database for worldwide application support

Core Features

- **High Performance:** Up to 5x MySQL and 3x PostgreSQL throughput with Aurora Parallel Query and Optimized Reads
- **Aurora Serverless:** Automatic scaling configuration that starts, stops, and scales based on application needs
- **Aurora Global Database:** Multi-Region replication with low-latency global reads and disaster recovery capabilities
- **Storage Auto-scaling:** Distributed storage automatically grows from 10 GiB to 128 TiB without downtime
- **Multi-Availability Zone Availability:** High availability with Aurora Replicas and automatic failover across availability zones
- **Point-in-time Recovery:** Continuous incremental backups with restore capability to any second within 35 days
- **Zero-ETL Integration:** Direct integration with Amazon Redshift and Amazon SageMaker without complex data pipelines
- **Aurora Machine Learning:** Built-in ML capabilities with Amazon SageMaker and Amazon Comprehend integration
- Amazon Aurora DSQL: Standalone distributed SQL database optimized for multi-region active-active architectures.

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/rds/aurora/
FAQ: https://aws.amazon.com/rds/aurora/faqs/

Technical Overview

Backup and Recovery Capabilities

Backup: Amazon Aurora performs continuous incremental backups to Amazon Simple Storage Service (Amazon S3) without performance impact, storing data durably across multiple Availability Zones. Manual snapshots can be created on-demand (Amazon Relational Database Service (Amazon RDS) console) or automated via AWS Backup with centralised policies and cross-region replication. Integrates with AWS Backup for compliance features like Vault Lock.

Recovery: Point-in-time recovery to any second within 1-35 day retention periods, Aurora Global Database cross-Region disaster recovery, and automated failover through Aurora Replicas. Distributed fault-tolerant storage system provides high availability. Supports aggressive RPO objectives (seconds with Global Database near zero) and RTO objectives (minutes through automated failover).

Setup Process

Onboarding: Create Amazon Aurora DB clusters through Amazon RDS console by selecting Aurora MySQL or PostgreSQL engine. Configure instance classes, Amazon Virtual Private Cloud (Amazon VPC) networking, security groups, database credentials, and backup options. Push-button migration tools enable seamless migration from existing Amazon RDS databases through snapshots or replication.

Offboarding: Delete Amazon Aurora DB clusters to automatically remove associated instances and storage. Manual snapshots require explicit deletion as they persist beyond cluster deletion. Complete data removal includes deleting all snapshots, disabling automated backups, and removing exported Amazon S3 data.

Data Management

- **Storage Options:** Distributed, fault-tolerant storage that automatically scales from 10 GiB to 128 TiB without downtime
- **Data Removal:** Delete Aurora DB clusters to remove instances and storage; manual snapshots require explicit deletion
- **Cross-Region Replication:** Aurora Global Database enables multi-Region replication with low-latency reads and disaster recovery capabilities
- **Encryption:** Built-in encryption capabilities with AWS managed keys for data protection at rest and in transit

Optimise Cost and Consumption

Amazon Aurora Serverless v2 automatically scales based on demand and can scale to zero during idle periods, reducing unused capacity costs. Aurora I/O-Optimized configuration specifically reduces I/O costs for high-throughput workloads. Reserved Instances provide significant discounts for predictable workloads. Storage auto-scaling ensures payment only for used storage capacity. Database cloning creates cost-effective copies for testing without duplicating storage resources.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/
Getting Started:
https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/CHAP_GettingStartedAurora.html
API Reference: https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/ProgrammingGuide.html
Service Quotas: https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/CHAP_Limits.html

Amazon Bedrock

General Information

Amazon Bedrock is a fully managed service providing access to high-performing foundation models from leading AI companies through a unified API. It enables developers to build generative AI applications with comprehensive security, privacy, and responsible AI features without managing infrastructure.

Main Benefits

Amazon Bedrock helps eliminate AI infrastructure complexity while providing unparalleled model choice from leading providers through a single API. Key advantages include up to 90% cost reduction through prompt caching, up to 75% savings with model distillation, and serverless architecture with automatic scaling and enterprise-grade security.

Key Use Cases

- **Content Generation:** Create marketing materials, reports, and creative writing with AI-powered document processing
- **Conversational AI:** Build intelligent chatbots and customer service applications for enhanced support experiences
- **Code Development:** Generate, optimise, and scan code for vulnerabilities in software development workflows
- **Data Analysis:** Extract insights from documents, images, videos, and audio files automatically.

Core Features

- **Foundation Model Choice:** Access industry-leading models from Amazon, Anthropic, AI21 Labs, Cohere, Meta, and others through unified API
- **Model Customization:** Fine-tuning, continued pretraining, and model distillation capabilities for domain-specific adaptation and optimisation
- **Knowledge Bases and RAG:** Retrieval-augmented generation with support for multimodal data and structured data retrieval capabilities
- **Managed Agents:** Execute complex tasks and coordinate with each other for sophisticated multi-agent collaborative workflows
- **Guardrails:** Built-in safeguards against harmful content with automated reasoning and contextual grounding security checks
- **Data Automation:** Automated processing and insight extraction from documents, images, videos, and audio files seamlessly
- **Model Evaluation:** Automatic and human-based evaluation capabilities with Large language model (LLM)-as-a-Judge functionality for performance assessment
- **Service Tiers:** Priority, Standard, and Flex tiers to optimise performance and cost for different workload requirements

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/bedrock/

FAQ: https://aws.amazon.com/bedrock/faqs/

Technical Overview

Backup and Recovery Capabilities

Backup: Amazon Bedrock operates as a managed service where AWS handles infrastructure resilience with built-in redundancy. Custom models and configurations are automatically maintained by AWS across multiple Availability Zones (AZ). Customer data in Knowledge Bases relies on underlying services like Amazon Simple Storage Service (Amazon S3), with AWS Backup integration available for storage components.

Recovery: Built-in disaster recovery through AWS's multi-AZ architecture provides automatic failover capabilities and high availability. The service maintains redundancy for quick recovery of custom models and configurations. Supports customer RPO objectives through consistent API availability and RTO objectives via multi-Region deployment with proper error handling.

Setup Process

Onboarding: Setup requires AWS account with proper AWS Identity and Access Management (AWS IAM) permissions for Amazon Bedrock. Users can start through console playgrounds, APIs, SDKs, or step-by-step tutorials. Configuration involves creating IAM roles, requesting model access, and selecting appropriate foundation models for specific use cases.

Offboarding: Users can delete custom models, Knowledge Bases, agents, and associated resources through console or APIs. Data in integrated services like Amazon S3 must be deleted separately by customers.

Data Management

- **Storage Options:** Integrates with Amazon S3 for Knowledge Bases and Amazon DynamoDB for state management with AWS-managed redundancy
- **Data Removal:** Delete custom models, Knowledge Bases, and agents via console or APIs; Amazon S3 data requires separate deletion
- **Cross-Region Replication:** Multi-AZ deployment with AWS global infrastructure supporting multiple Regions for disaster recovery strategies
- **Encryption:** Enterprise-grade security with complete data privacy and no sharing of customer inputs with third-party providers.

Optimise Cost and Consumption

Amazon Bedrock provides comprehensive cost optimisation through service tiers enabling workload-specific pricing, with Flex tier offering significant discounts for batch processing. Prompt caching reduces costs up to 90% by storing repeated prefixes. Model Distillation creates smaller models up to 75% less expensive with minimal accuracy loss. Intelligent Prompt Routing automatically selects cost-effective models, reducing expenses up to 30%. Provisioned Throughput offers predictable pricing for consistent workloads.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/bedrock/latest/userguide/what-is-bedrock.html
Getting Started: https://docs.aws.amazon.com/bedrock/latest/userguide/getting-started.html
API Reference: https://docs.aws.amazon.com/bedrock/latest/APIReference/welcome.html
Service Quotas: https://docs.aws.amazon.com/bedrock/latest/userguide/quotas.html

Amazon Braket

General Information

Amazon Braket is a fully managed quantum computing service providing researchers and developers access to quantum computers from multiple providers including IonQ, Rigetti, and QuEra. It enables quantum algorithm development, testing on simulators, and execution on various quantum hardware types.

Main Benefits

Amazon Braket helps eliminate traditional quantum computing barriers by providing cost-effective, on-demand access to multiple quantum hardware providers without upfront commitments. The hardware-agnostic approach enables algorithm testing across different quantum technologies with single line code changes, while consumption-based pricing makes quantum computing accessible to organisations of all sizes.

Key Use Cases

- **Quantum Algorithm Research:** Design and optimise quantum algorithms for computational problems beyond classical reach
- **Hybrid Computing:** Develop variational quantum algorithms combining classical and quantum resources using PennyLane
- **Quantum Simulation:** Simulate quantum systems and molecular behaviour for drug discovery and materials science
- **Optimisation Problems:** Solve complex challenges in financial portfolio management, logistics, and machine learning.

Core Features

- **Multi-Provider Hardware Access:** On-demand access to quantum computers from AQT, IonQ, IQM, Rigetti, and QuEra providers
- **Quantum Circuit Simulators:** Four simulator options including local simulator, SV1, DM1, and TN1 for testing circuits
- **Hardware-Agnostic SDK:** Braket SDK for designing quantum algorithms with multiple programming languages and framework support
- **Hybrid Jobs:** Managed execution environment for quantum-classical hybrid algorithms with embedded simulators and scaling capabilities
- **Braket Direct:** Dedicated device access through reservations with expert advice and experimental quantum computing capabilities
- **Jupyter Notebook Environment:** Fully managed notebooks pre-installed with sample algorithms, resources, and quantum development tools
- **AWS Service Integration:** Deep integration with Amazon Simple Storage Service (Amazon S3), Amazon SageMaker, AWS Identity and Access Management (AWS IAM), Amazon CloudWatch, and Amazon EventBridge for comprehensive workflows
- **PennyLane Framework:** Integration with PennyLane for developing variational quantum algorithms in NISQ era applications

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/braket/

FAQ: https://aws.amazon.com/braket/faqs/

Technical Overview

Backup and Recovery Capabilities

Backup: Amazon Braket lacks native backup capabilities as it functions as a stateless quantum computing execution service. Quantum computation results are stored in customer-owned Amazon S3 buckets, which support versioning, cross-Region replication, and integration with AWS Backup for automated scheduling. Customers must implement backup strategies for quantum algorithms and notebook code using standard AWS services.

Recovery: Recovery relies on re-running quantum circuits rather than restoring previous states due to the service's stateless design. Multi-Region availability across US East, US West, and Europe provides geographic distribution. Supports customer RPO objectives through Amazon S3 cross-Region replication for near-zero data loss and RTO objectives limited by quantum hardware queue times and circuit execution duration.

Setup Process

Onboarding: Enable Amazon Braket in AWS account and create managed Jupyter notebook instances through AWS Management Console. Configure AWS IAM permissions for service access and Amazon S3 bucket permissions. Start with local simulators, progress to on-demand simulators, then access quantum hardware.

Offboarding: Stop and delete notebook instances through Amazon Braket console. Delete Amazon S3 buckets and associated data following standard procedures. Terminate running hybrid jobs and delete created AWS IAM roles and policies.

Data Management

- **Storage Options:** Quantum computation results stored in customer-owned Amazon S3 buckets with full lifecycle control
- **Data Removal:** Delete Amazon S3 buckets and data following standard procedures, stop notebook instances via Braket console
- **Cross-Region Replication:** Supported through Amazon S3 cross-Region replication capabilities for quantum computation results and algorithms
- **Encryption:** Inherits Amazon S3 encryption capabilities for quantum computation results stored in customer-controlled buckets.

Optimize Cost and Consumption

Amazon Braket provides spending limits for quantum processing units enabling proactive cost control and budget management. Near real-time cost tracking within the Braket SDK offers immediate visibility into quantum task expenses. TN1 simulator with low shot counts optimises initial algorithm testing costs. Amazon Braket Direct reserved device access provides hourly rates for high-volume users requiring dedicated quantum hardware access. Cost estimation tools integrated within development notebooks enable precise budget planning before quantum task execution.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/braket/latest/developerguide/what-is-braket.html

Getting Started: https://docs.aws.amazon.com/braket/latest/developerguide/braket-get-started.html

API Reference: https://docs.aws.amazon.com/braket/latest/APIReference/index.html

Service Quotas: https://docs.aws.amazon.com/braket/latest/developerguide/braket-quotas.html

Amazon Chime SDK

General Information

Amazon Chime SDK provides real-time communication components enabling developers to add messaging, audio, video, screen sharing, and telephony capabilities to applications. The SDK eliminates infrastructure complexity through pay-for-use pricing and supports multiple platforms including JavaScript, iOS, Android, Windows, and C++.

Note: While the Amazon Chime service has been discontinued, the Amazon Chime SDK remains fully available and supported.

Main Benefits

Amazon Chime SDK eliminates infrastructure complexity and upfront costs through pay-per-use pricing with elastic scaling. Supports up to 250 participants, 4K screen sharing, 100,000-member messaging channels, and broadcasting to 10,000 attendees. Features modular architecture enabling developers to implement only required capabilities while integrating seamlessly with AWS AI services.

Key Use Cases

- **Video Conferencing:** Custom meeting platforms with audio, video, and screen sharing for business collaboration
- **Telemedicine Solutions:** Remote patient consultations with healthcare providers through secure video communications
- **Online Education:** Interactive learning environments with real-time audio, video, and content sharing capabilities
- **Customer Support:** Click-to-call functionality and contact centre solutions with PSTN integration

Core Features

- **WebRTC Media Sessions:** Real-time audio, video up to 1080p, screen sharing up to 4K supporting 250 participants
- **PSTN Audio:** Serverless voice applications with inbound/outbound calling, call recording, and phone number provisioning
- **Large-Scale Messaging:** Up to 100,000 members per channel with message history, moderation tools, and push notifications
- **Speech Enhancement:** Amazon Voice Focus noise reduction, echo reduction, and live transcription with Amazon Transcribe
- **Call Analytics:** Speaker search, voice tone analysis, real-time transcription, and AWS AI integrations for insights
- **Media Pipelines:** Media capture, streaming to Kinesis, live connector for broadcasting, and concatenation capabilities
- **Multi-Platform Support:** Client libraries for JavaScript, iOS, Android, Windows, and C++ platforms for diverse applications
- **Interactive Broadcasting:** Stream meetings to up to 10,000 attendees with live connector and media replication capabilities

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/chime/chime-sdk/

FAQ: https://aws.amazon.com/chime/chime-sdk/faq/

Technical Overview

Backup & Recovery Capabilities

Backup: Most Amazon Chime SDK resources are ephemeral by design and don't require traditional backup. Persistent data like call recordings stored in Amazon S3 can leverage AWS Backup for automated backup policies (console, CLI, APIs). Message data and call analytics can be backed up using standard AWS services with configurable retention policies.

Recovery: Automatic failover between Availability Zones provides minimal recovery time for active communications, with cross-region strategies available for S3-stored recordings. Built-in infrastructure redundancy handles failures automatically. Supports customer RPO (S3 cross-region replication and backup policies) and RTO (automatic scaling and infrastructure redundancy) objectives.

Setup Process

Onboarding: Requires AWS account and programming skills. Create AppInstances for messaging, configure server/client applications for meetings, and create SIP rules for PSTN audio. Select appropriate client libraries (JavaScript, iOS, Android, Windows, C++), integrate with AWS SDKs, and configure IAM permissions for secure access.

Offboarding: Delete resources through console, CLI, or APIs. Remove AppInstances, channels, meeting resources, call recordings from S3, voice analytics data via APIs, and PSTN resources. Service-linked roles removed after resource deletion for clean offboarding.

Data Management

- **Storage Options:** Amazon S3 stores call recordings, analytics data, and attachments with high durability and availability
- **Data Removal:** Delete resources via console, CLI, or APIs including AppInstances, channels, recordings, and analytics data
- **Cross-Region Replication:** Available for persistent S3-stored recordings and analytics data to meet specific RPO requirements
- **Encryption:** Enterprise-grade security with end-to-end encryption built on AWS secure, scalable, and reliable global infrastructure

Optimize Cost and Consumption

Amazon Chime SDK's modular architecture enables selective implementation of required media modalities, optimizing costs by avoiding unused features. Choose standard or high-definition sessions based on participant needs to balance quality and expense. Media replication strategically broadcasts to multiple attendees through single sessions rather than separate instances. Efficient session management prevents idle resource costs through automatic scaling. Configure S3 retention policies for recordings and analytics data to control storage expenses while leveraging regional deployment flexibility for optimal pricing.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/chime-sdk/latest/dg/

Getting Started: https://docs.aws.amazon.com/chime-sdk/latest/dg/getting-started.html

API Reference: https://docs.aws.amazon.com/chime-sdk/latest/APIReference/

Service Quotas: https://docs.aws.amazon.com/general/latest/gr/chime-sdk.html

Amazon CloudFront

General Information

Amazon CloudFront is a global content delivery network that accelerates distribution of static and dynamic web content through worldwide edge locations. It routes user requests to the lowest-latency edge location, delivering cached content immediately or retrieving it from origin servers for optimal performance.

Main Benefits

Amazon CloudFront delivers global performance through hundreds of terabits of deployed capacity and connections to thousands of telecom carriers. It provides automatic distributed denial of service (DDoS) protection, sub-millisecond edge function execution, free data transfer between AWS services, and comprehensive compliance certifications including Payment Card Industry Data Security Standard Level 1, Health Insurance Portability and Accountability Act, and System and Organization Controls (SOC).

Key Use Cases

- **Static Website Content Delivery:** Accelerates images, stylesheets, and JavaScript using Amazon Simple Storage Service (Amazon S3) as origin
- **Video Streaming:** Supports video on demand and live streaming with edge caching of media fragments
- **API and Application Acceleration:** Enhances dynamic content delivery through intelligent routing and protocol optimisations
- **Global Content Distribution:** Delivers websites, applications, and APIs worldwide with low latency performance.

Core Features

- **Global Edge Network:** Worldwide network of point of presences, Regional Edge Caches, and embedded points of presence
- **Origin Shield:** Additional caching layer that reduces origin traffic and improves cache hit ratios
- **CloudFront Functions:** Lightweight JavaScript functions with sub-millisecond execution for high-scale, latency-sensitive CDN customisations
- **Lambda@Edge:** Serverless compute for complex operations running closer to viewers in Node.js or Python
- **Security Integration:** Built-in AWS Shield DDoS protection, AWS WAF integration, field-level encryption, and SSL/TLS support
- **Real-time Monitoring:** Amazon CloudWatch integration for metrics, standard and real-time logging with comprehensive request tracking
- **Multiple Origin Support:** Supports Amazon S3, MediaPackage, Load Balancers, Amazon Elastic Compute Cloud instances, and custom HTTP servers
- **Origin Failover:** Automatic traffic routing to backup origins when primary origins become unavailable.

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/cloudfront/

FAQ: https://aws.amazon.com/cloudfront/faqs/

Technical Overview

Backup and Recovery Capabilities

Backup: Amazon CloudFront does not provide traditional backup capabilities as it focuses on caching and distribution rather than data storage. Content backup is managed at the origin level through Amazon S3 versioning, database backups, or application server backups. Amazon CloudFront automatically caches content across its global edge network for performance optimization.

Recovery: Amazon CloudFront supports disaster recovery through automatic origin failover that switches to backup origins during primary origin failures. The globally distributed architecture maintains content availability across hundreds of edge locations during regional outages. Supports aggressive RPO objectives (near-zero for cached content) and RTO objectives (seconds to minutes through near-instantaneous origin switching).

Setup Process

Onboarding: Setup involves creating distributions via AWS Management Console or CLI by specifying origin servers. Configure cache behaviours, security settings, and origin access control. AWS provides tutorials for console-based setup and secure static website deployment.

Offboarding: Disable distributions first, wait for propagation, then delete using distribution ETag. Clean up associated resources including behaviours S3 buckets, SSL certificates, Amazon CloudFront Functions, Lambda@Edge functions, and AWS Identity and Access Management roles to avoid ongoing charges.

Data Management

- **Storage Options:** Caching at edge locations, Regional Edge Caches, and Origin Shield for performance optimisation
- **Data Removal:** Disable distributions, wait for propagation, then delete with ETag and clean associated resources
- **Cross-Region Replication:** Multi-origin support across different AWS Regions with automatic origin failover capabilities
- **Encryption:** Field-level encryption, SSL/TLS support, and integration with AWS Certificate Manager for secure delivery.

Optimise Cost and Consumption

Amazon CloudFront Origin Shield reduces origin server costs by collapsing requests and improving cache hit ratios through centralised caching. Intelligent cache policy configuration maximises cache hit ratios, reducing expensive origin requests. Price class selection limits delivery to specific geographic regions, excluding more expensive edge locations. Regional Edge Caches decrease operational burden on origins while flat-rate pricing plans provide predictable monthly costs with included DDoS protection.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/
Getting Started: https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/GettingStarted.html
API Reference: https://docs.aws.amazon.com/cloudfront/latest/APIReference/
Service Quotas: https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/cloudfront-limits.html

Amazon CloudWatch

General Information

Amazon CloudWatch is a monitoring and observability service that provides real-time visibility into AWS resources and applications. It collects metrics, monitors logs, sets alarms, and offers dashboards for system-wide operational insights across your infrastructure.

Main Benefits

Amazon CloudWatch offers native integration with all AWS services, offering automatic metric collection without additional setup or separate monitoring infrastructure. Unlike third-party solutions, it delivers unified visibility across compute, storage, database, and network resources with cross-account observability, reducing operational overhead.

Key Use Cases

- **Infrastructure Monitoring:** Monitor Amazon Elastic Compute Cloud (Amazon EC2) instances, Amazon Relational Database Service (Amazon RDS) databases, and Amazon Elastic Kubernetes Service (Amazon EKS) clusters with real-time performance metrics
- **Application Performance Monitoring:** Track application health, latency, error rates, and request volumes with automatic detection
- **Log Management and Analysis:** Centralise, store, search, and analyse logs using Amazon CloudWatch Logs Insights capabilities
- **Cross-Account Observability:** Monitor applications spanning multiple AWS accounts from centralised monitoring with unified dashboards

Core Features

- **Metrics and Alarms:** Collect performance metrics with customisable alarms that trigger automated actions when thresholds are breached
- **Dashboards:** Create unified visualisations of metrics and logs with customisable dashboards shared across accounts and AWS Regions
- **Application Signals:** Automatically detect and monitor application performance indicators like latency and error rates without manual instrumentation
- **CloudWatch Logs:** Log management with storage, search, analysis, and real-time processing using multiple query languages
- **Container Insights:** Specialised monitoring for containerised applications on Amazon ECS, Amazon EKS, and self-managed Kubernetes clusters
- **Database Insights:** Real-time database performance monitoring with SQL query analysis and load troubleshooting for AWS databases
- **Synthetics:** Proactive endpoint and API monitoring through configurable canaries that simulate user behaviour and detect issues
- **Network Monitoring:** Internet Monitor and Network Flow Monitor for analysing network performance and identifying connectivity bottlenecks.

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/cloudwatch/
FAQ: https://aws.amazon.com/cloudwatch/faqs/

Technical Overview

Backup and Recovery Capabilities

Backup: Amazon CloudWatch serves as the monitoring layer for backup operations performed by AWS Backup, Amazon Elastic Block Store snapshots, and Amazon RDS automated backups rather than providing direct backup capabilities. Metric and log data are automatically replicated across multiple Availability Zones with built-in durability and configurable retention policies up to 15 months for metrics.

Recovery: Cross-Region log centralisation enables backup of log data to alternate AWS Regions for compliance and disaster recovery. Amazon CloudWatch monitors replication health and integrates with AWS Elastic Disaster Recovery for visibility into recovery operations. Supports customer RPO objectives through backup frequency monitoring and RTO objectives via automated alarm-triggered disaster recovery procedures.

Setup Process

Onboarding: Getting started requires an AWS account and basic AWS Identity and Access Management permissions. AWS services automatically send metrics to Amazon CloudWatch console at no additional cost. Install CloudWatch agent on Amazon EC2 instances for detailed metrics. Enable advanced features like Application Signals through console with pre-built configurations.

Offboarding: Disable Amazon CloudWatch agents and remove custom metrics publishing. Delete dashboards, alarms, and log groups to stop charges. Configure log retention policies and clear data while helping ensure compliance requirements.

Data Management

- **Storage Options:** Metrics automatically expire after 15 months with built-in durability across multiple Availability Zones
- **Data Removal:** Configure retention policies for logs and delete custom metrics, alarms, dashboards, and log groups immediately
- **Cross-Region Replication:** Cross-Region log centralisation enables backup of log data to alternate Regions for compliance purposes
- **Encryption:** Built-in encryption and secure data transmission through HTTPS endpoints and Amazon Virtual Private Cloud endpoint connectivity options.

Optimise Cost and Consumption

Amazon CloudWatch provides tiered pricing for vended logs with volume discounts for custom metrics, reducing costs at scale. Log retention policies automatically delete old data, while log sampling and filtering minimise ingestion volumes. Cross-account observability consolidates monitoring infrastructure across multiple accounts. Cost analysis tools identify high-cost log groups and metrics, and alternative destinations like Amazon Simple Storage Service and Amazon Kinesis Data Firehose offer lower storage costs.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/AmazonCloudWatch/latest/monitoring/WhatIsCloudWatch.html
Getting Started: https://docs.aws.amazon.com/AmazonCloudWatch/latest/monitoring/GettingSetup.html
API Reference: https://docs.aws.amazon.com/AmazonCloudWatch/latest/APIReference/
Service Quotas: https://docs.aws.amazon.com/AmazonCloudWatch/latest/monitoring/cloudwatch_limits.html

Amazon Cognito

General Information

Amazon Cognito is a customer identity and access management service providing user authentication, authorisation, and management for web and mobile applications. It enables secure sign-up, sign-in, access control, and AWS resource access through user pools and identity pools with scalable identity stores.

Main Benefits

Amazon Cognito helps eliminate authentication system complexity with fully managed identity solutions scaling to millions of users. Enterprise-grade security includes Health Insurance Portability and Accountability Act (HIPAA), Payment Card Data Security Standards (PCI DSS), System and Organization Controls (SOC), and International Organization for Standardization (ISO) compliance certifications. Monthly active user pricing with AWS Free Tiers (10,000 MAUs) makes it cost-effective for startups and enterprises alike.

Key Use Cases

- **User Authentication and Authorisation:** Secure sign-up, sign-in, and access control for web and mobile applications
- **Social and Enterprise Federation:** Authentication through Google, Facebook, Apple, or enterprise SAML/OIDC identity providers
- **AWS Resource Access Control:** Temporary AWS credentials for authenticated users accessing Amazon Simple Storage Service, Amazon DynamoDB, and AWS Lambda
- **Machine-to-Machine Authentication:** OAuth 2.0 client credentials flow for secure API-to-API authentication with tokens.

Core Features

- **User Pools:** User directories handling sign-up, sign-in, and profile management with customisable authentication workflows
- **Identity Pools:** Authorisation service granting temporary AWS credentials to authenticated or anonymous users for resource access
- **Managed Login:** No-code, branded web-based sign-in interface with customisable UI and multiple authentication method support
- **Multi-Factor Authentication:** Short Message Service, email OTP, authenticator apps, and WebAuthn passkeys for enhanced security protection
- **Federation Support:** Integration with social providers and SAML/OIDC-based enterprise identity systems for authentication
- **Lambda Triggers:** Customisable user workflows and authentication flows using serverless functions for complete extensibility
- **Machine-to-Machine Authentication:** OAuth 2.0 client credentials flow for secure API-to-API authentication with token customisation
- **Risk-Based Authentication:** Adaptive authentication using device, location, and behavioural analysis with compromised credential protection.

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/cognito/
FAQ: https://aws.amazon.com/cognito/faqs/

Technical Overview

Backup and Recovery Capabilities

Backup: Amazon Cognito does not provide built-in backup capabilities for user data or configurations and does not integrate with AWS Backup service. The service relies on AWS infrastructure redundancy for data protection. Customers must implement their own backup strategies for user pool configurations and export user data if needed.

Recovery: Recovery relies on multi- Availability Zone (AZ) architecture within AWS Regions providing automatic failover and high availability. Cross-Region disaster recovery requires manual user pool configuration replication and data migration. Supports customer RPO objectives through infrastructure redundancy and RTO objectives via automatic regional failover, though cross-region recovery involves longer RTOs.

Setup Process

Onboarding: Setup through AWS Console with quick setup wizard offering use case-specific recommendations and framework integration instructions. Choose managed login for no-code implementation, SDK integration for custom interfaces, or AWS Amplify for full-stack applications.

Offboarding: User data removal via DeleteUser API or console interface permanently deletes accounts and profiles. Complete user pool and identity pool deletion through console or APIs removes all contained data with bulk deletion support.

Data Management

- **Storage Options:** Scalable identity stores in user pools and identity pools with multi-AZ data replication
- **Data Removal:** DeleteUser API and console interface permanently delete accounts with bulk deletion and export capabilities
- **Cross-Region Replication:** Not supported natively; customers must manually implement cross-Region disaster recovery strategies and configuration replication
- **Encryption:** Enterprise-grade security with HIPAA, PCI DSS, SOC, and ISO compliance certifications for data protection.

Optimise Cost and Consumption

Amazon Cognito provides Monthly Active User pricing model optimisation through strategic user pool architecture and feature tier selection. Implementing efficient token refresh mechanisms reduces unnecessary API calls, while linking federated users to local profiles minimises federation costs. Using ListUsers for bulk operations and proper user lifecycle management prevents billing for inactive users. Service Quotas console monitoring enables consumption tracking and spending optimisation.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/cognito/latest/developerguide/what-is-amazon-cognito.html

Getting Started: https://docs.aws.amazon.com/cognito/latest/developerguide/cognito-guided-setup.html

API Reference: https://docs.aws.amazon.com/cognito-user-identity-pools/latest/APIReference/

Service Quotas: https://docs.aws.amazon.com/cognito/latest/developerguide/quotas.html

Amazon Comprehend Medical

General Information

Amazon Comprehend Medical is a HIPAA-eligible natural language processing service that extracts medical insights from unstructured clinical text. It identifies medical entities, detects PHI, and links findings to standardized ontologies like ICD-10-CM and SNOMED CT.

Main Benefits

Delivers industry-leading accuracy with fully managed, HIPAA-eligible NLP that eliminates custom model development. Features pre-trained healthcare-specific models, automatic scaling, pay-as-you-use pricing with no upfront commitments, and 8.5-million-character free tier for evaluation.

Key Use Cases

- **Patient Case Management:** Extract structured data from clinical notes for improved care coordination
- **Clinical Research:** Analyse medical text to identify patient cohorts and treatment patterns
- **Medical Billing:** Automate coding and claim validation using standardized medical ontology linking
- **Population Health Analytics:** Process clinical documents to identify health trends and risk factors

Core Features

- **Medical Named Entity Extraction:** Identifies conditions, medications, dosages, tests, treatments, and procedures with confidence scores
- **PHI Detection:** Identifies and helps redact sensitive patient information for HIPAA compliance
- **Medical Ontology Linking:** Links entities to ICD-10-CM, RxNorm, and SNOMED CT for standardized coding
- **Batch Processing:** Supports asynchronous batch operations for collections of documents stored in S3
- **Real-time Processing:** Enables synchronous operations for individual document analysis with immediate results
- **Confidence Scoring:** Provides scores for detected entities to enable thresholding and human review workflows
- **HIPAA Eligibility:** Fully managed service designed for healthcare compliance and data security requirements
- **Automatic Scaling:** Handles varying workloads without manual intervention or infrastructure management

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/comprehend/medical/

FAQ: https://aws.amazon.com/comprehend/medical/faqs/

Technical Overview

Backup & Recovery Capabilities

Backup: Amazon Comprehend Medical processes data transiently without persistent storage. Customers implement backup strategies for input data and results using S3 Cross-Region Replication, versioning, and lifecycle policies. AWS Backup can backup S3 buckets containing clinical documents and analysis results.

Recovery: Multi-AZ architecture provides rapid failover within regions, cross-region workflows with Route 53 failover routing, and API-based architecture enables quick restoration. Supports customer RPO (low objectives through transient processing eliminating data loss risks) and RTO (sub-hour objectives via multi-region architectures with pre-configured workflows).

Setup Process

Onboarding: Access via AWS Management Console, CLI, or SDKs with IAM user setup and appropriate permissions. Configure S3 buckets for batch processing with proper access policies. Test using DetectEntitiesV2 API with sample clinical text and 8.5 million character free tier.

Offboarding: Stop active batch analysis jobs, delete stored results from S3 output buckets, and revoke IAM permissions. Standard AWS resource management procedures enable complete cleanup automation.

Data Management

- **Storage Options:** Transient processing with no persistent storage; results stored in customer-specified S3 buckets
- **Data Removal:** Service doesn't retain data beyond processing; customers delete S3 results and revoke permissions
- **Cross-Region Replication:** Supported via S3 Cross-Region Replication for input documents and analysis results backup
- **Encryption:** HIPAA-eligible usage requires encrypted connections and proper data handling procedures for compliance

Optimize Cost and Consumption

Batch processing delivers better unit economics for high-volume scenarios compared to synchronous operations. Confidence threshold filtering reduces processing of low-value text segments while maintaining quality standards. Text preprocessing workflows eliminate unnecessary non-clinical content processing. Strategic API operation selection between entity detection and ontology linking optimizes costs based on specific requirements. Regional selection and data locality optimization further reduce overall processing expenses.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/comprehend-medical/latest/dev/comprehendmedical-welcome.html

Getting Started: https://docs.aws.amazon.com/comprehend-medical/latest/dev/comprehendmedical-gettingstarted.html

API Reference: https://docs.aws.amazon.com/comprehend-medical/latest/api/Welcome.html

Service Quotas: https://docs.aws.amazon.com/comprehend-medical/latest/dev/comprehendmedical-quotas.html

Amazon Connect

General Information

Amazon Connect is a cloud-based omnichannel contact centre-as-a-service (CCaaS) solution that enables businesses to deliver customer experiences across voice, chat, short message service (SMS), email, and video. It provides AI-powered capabilities including conversational IVR, chatbots, agent assistance, and real-time analytics with consumption-based pricing.

Main Benefits

Amazon Connect helps eliminate infrastructure management while providing virtually unlimited AI capabilities and integration with over 200 AWS services. Organisations achieve 30-60% cost savings compared to traditional solutions, with studies showing $4.3 million in cloud technology savings and $4.6 million in agent labour savings through improved productivity.

Key Use Cases

- **Customer Service and Support:** Handle inquiries across voice, chat, email, SMS with intelligent routing
- **Sales and Lead Generation:** Manage outbound campaigns with predictive dialling and ML-powered optimisation
- **Self-Service Automation:** Deploy AI-powered chatbots using Amazon Lex for authentication and transactions
- **Technical Support and IT Help Desk:** Provide multi-tier support with skill-based routing integration.

Core Features

- **Omnichannel Communication:** Voice, chat, email, SMS, video calling, and task management in unified platform
- **AI-Powered Self-Service:** Conversational interactive voice response (IVR) and chatbots with Amazon Lex for natural language interactions
- **Agent Workspace:** Integrated desktop with unified customer profiles, step-by-step guides, and real-time assistance
- **Contact Routing:** Skills-based routing, queue management, and intelligent distribution based on agent capabilities
- **Real-time Analytics:** Live dashboards, historical reports, conversational analytics, and sentiment analysis capabilities
- **Workforce Management:** ML-powered forecasting, capacity planning, agent scheduling, and adherence tracking tools
- **Outbound Campaigns:** High-volume predictive dialling with answering machine detection and campaign optimisation
- **Generative AI Assistant:** Amazon Connect provides real-time agent recommendations and post-contact summaries powered by generative AI.

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/connect/
FAQ: https://aws.amazon.com/connect/faqs/

Technical Overview

Backup and Recovery Capabilities

Backup: Call recordings stored in Amazon Simple Storage Service (Amazon S3) designed for 99.999999999% durability and optional cross-region replication. Flow configurations can be exported manually (AWS Management Console) or automated via API-based programmatic export. Contact records automatically replicated across Availability Zones with streaming capabilities for external backup systems.

Recovery: Automatic failover between Availability Zones (AZ) with recovery in minutes, cross-Region disaster recovery with 15–30-minute RTOs, and Infrastructure as Code approaches for rapid environment recreation. Multi-AZ active-active architecture provides fault tolerance. Supports near-zero RPO (automatic multi-AZ replication) and aggressive RTO (automated configuration replication) objectives.

Setup Process

Onboarding: Create instances via AWS Management Console with six-step process. Configure telephony options, routing profiles, and drag-and-drop flow designer for customer experiences. Setup completed in hours with guided wizards and tutorials.

Offboarding: Delete instances through AWS Management Console releasing phone numbers to inventory. All settings, data, and reports become permanently inaccessible. Manually delete Amazon CloudWatch log groups and Amazon S3 call recordings separately.

Data Management

- **Storage Options:** Amazon S3 for call recordings and reports designed for 99.999999999% durability, and Amazon DynamoDB for real-time data
- **Data Removal:** Instance deletion makes all settings permanently inaccessible, 30-day GDPR retention for scheduling data
- **Cross-Region Replication:** Supported through Amazon S3 Cross-Region Replication for call recordings and additional data protection
- **Encryption:** Enterprise-grade security and compliance built on AWS infrastructure with encryption and access controls.

Optimise Cost and Consumption

Strategic AWS Region selection delivers 20–40% telephony cost savings through optimised phone number and inbound usage pricing. Callback functionality reduces Public Switch Telephone Network charges by minimising customer hold times. Self-service automation using Amazon Lex decreases agent labour costs for routine inquiries. Voice-to-chat redirection enables agents to handle multiple simultaneous conversations, improving operational efficiency. Amazon S3 lifecycle policies automatically transition older call recordings to cheaper storage tiers, optimising long-term storage costs.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/connect/latest/adminguide/
Getting Started: https://docs.aws.amazon.com/connect/latest/adminguide/amazon-connect-get-started.html
API Reference: https://docs.aws.amazon.com/connect/latest/APIReference/
Service Quotas: https://docs.aws.amazon.com/connect/latest/adminguide/amazon-connect-service-limits.html

Amazon Data Firehose

General Information

Amazon Data Firehose is a fully managed service for delivering real-time streaming data to destinations like Amazon Simple Storage Service (Amazon S3), Amazon Redshift, and Amazon OpenSearch Service. It automatically buffers, transforms, and compresses data without requiring custom applications or resource management.

Main Benefits

Eliminates infrastructure management complexity with zero provisioning requirements and automatic scaling. Offers up to 99.9% SLA uptime guarantee with at-least-once delivery semantics and 24-hour retry logic. Integrates with over 30 AWS services and provides serverless consumption-based pricing model with no upfront commitments.

Key Use Cases

- **Data Lakes and Warehouses:** Loading real-time application data into Amazon S3 and Amazon Redshift for analytics
- **Security Monitoring:** Streaming security logs to platforms like Splunk for real-time threat detection
- **Machine Learning Pipelines:** Feeding streaming data into ML platforms for real-time inference and training
- **Business Intelligence:** Delivering streaming business metrics to data warehouses for real-time dashboards.

Core Features

- **Fully Managed Service:** No infrastructure management required with automatic scaling and provisioning
- **Multiple Data Sources:** Supports Amazon Kinesis Data Streams, Amazon Managed Streaming for Apache Kafka (Amazon MSK), Direct PUT, and AWS service logs
- **Data Transformation:** Real-time data transformation using AWS Lambda functions before delivery to destinations
- **Format Conversion:** Convert data to columnar formats like Apache Parquet and ORC for optimised analytics
- **Dynamic Partitioning:** Automatically partition data in Amazon S3 based on content for improved query performance
- **Multiple Destinations:** Deliver to Amazon S3, Amazon Redshift, Amazon OpenSearch, Splunk, Snowflake, and HTTP endpoints
- **Buffer Management:** Configurable buffer sizes and intervals for optimised delivery batches and throughput
- **Error Handling:** Automatic retry logic and error logging with failed record backup capabilities

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/firehose/

FAQ: https://aws.amazon.com/firehose/faqs/

Technical Overview

Backup and Recovery Capabilities

Backup: Amazon Data Firehose provides automatic backup by storing source data in Amazon S3 concurrently with destination delivery. Failed record backup automatically saves undelivered records in designated Amazon S3 error folders. The service buffers data for up to 24 hours during delivery failures for DirectPut sources.

Recovery: Recovery relies on at-least-once delivery semantics with automatic retry logic for up to 24 hours. Cross-Region disaster recovery involves destination-level replication and recreating Firehose streams in alternate AWS Regions. Supports customer RPO objectives as low as seconds through real-time streaming and RTO objectives through automatic scaling and quick stream recreation.

Setup Process

Onboarding: Create delivery streams through AWS Management Console, CLI, or SDKs. Configure data sources, optional transformations, destinations, and buffer settings using the console wizard. No infrastructure provisioning required, enabling data streaming within minutes.

Offboarding: Stop data ingestion, allow buffered data delivery, then delete streams via console or APIs. Manual deletion required for Amazon S3 objects, error buckets, and source data backups depending on destination configurations.

Data Management

- **Storage Options:** Delivers to Amazon S3, Amazon Redshift, Amazon OpenSearch, Splunk, Snowflake, and HTTP endpoints with automatic buffering
- **Data Removal:** Manual deletion required for Amazon S3 objects, error buckets, and source backups per destination
- **Cross-Region Replication:** No native cross-Region replication; achieved through destination-level replication and stream recreation
- **Encryption:** Server-side encryption with AWS Key Management and AWS Identity and Access Management-based access control for data security

Optimise Cost and Consumption

Optimise record sizes above the 5 KB minimum billing increment for Direct PUT sources to reduce processing charges. Use PutRecordBatch API for batch processing to minimise per-request costs. Configure buffer sizes and intervals to optimise delivery batches and reduce destination charges. Enable format conversion to columnar formats like Parquet to decrease storage costs and improve query performance. Apply compression algorithms to reduce data transfer expenses while maintaining throughput efficiency.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/firehose/latest/dev/index.html
Getting Started: https://docs.aws.amazon.com/firehose/latest/dev/what-is-this-service.html
API Reference: https://docs.aws.amazon.com/firehose/latest/APIReference/Welcome.html
Service Quotas: https://docs.aws.amazon.com/firehose/latest/dev/limits.html

Amazon DataZone

General Information

Amazon DataZone is a comprehensive data management service that enables organisations to catalogue, discover, share, and govern data across AWS, on-premises, and third-party sources. It provides fine-grained access controls and automated workflows for secure data collaboration while breaking down organisational data silos.

Main Benefits

Amazon DataZone delivers unified data governance combining cataloguing, discovery, sharing, and governance in a single service, helping eliminate complexity from managing multiple separate tools. It provides AI-powered automated metadata enrichment, supports up to 1 million data assets per domain, and offers consumption-based pricing with no user subscription fees.

Key Use Cases

- **Breaking Down Data Silos:** Enable business teams to build catalogues while maintaining access control across boundaries
- **Data Mesh Architecture:** Support decentralised ownership with federated governance and secure data sharing workflows
- **Automated Discovery:** Reduce manual entry by automating metadata generation with AI-powered recommendations for cataloguing
- **Cross-Account Governance:** Share data across multiple AWS accounts with centralised governance and domain organisation.

Core Features

- **Business Data Catalog:** Centralised catalogue with AI-generated metadata, business context descriptions, and automated asset recommendations
- **Governed Data Sharing:** Subscription and fulfillment workflows with automated permission management for AWS Glue and Amazon Redshift
- **Projects and Environments:** Collaborative workspaces with role-based access controls and AWS service integration through blueprints
- **Data Products:** Collections of interconnected data assets designed for specific business use-cases with simplified administration
- **Fine-grained Access Control:** Row and column-level filtering using AWS Lake Formation Data Cell Filters and Amazon Redshift views
- **Multi-account Integration:** Support for data sharing across multiple AWS accounts with centralised governance and organisation
- **Browser-based Data Portal:** Web-based interface enabling team collaboration and self-service data access across organisational boundaries
- **Native AWS Integration:** Connectivity with AWS Glue, Amazon Redshift, Amazon Athena, Amazon SageMaker, Amazon Quick, and AWS Lake Formation services

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/datazone/
FAQ: https://aws.amazon.com/datazone/faqs/

Technical Overview

Backup and Recovery Capabilities

Backup: Amazon DataZone does not provide dedicated backup capabilities as it manages metadata and access controls rather than storing actual data. Metadata is stored in AWS managed infrastructure with built-in redundancy and durability. Customers must implement backup strategies for underlying data sources using AWS Backup for supported services or native backup features of Amazon Redshift and Amazon Relational Database Service.

Recovery: Recovery relies on AWS infrastructure resilience and does not integrate directly with AWS Elastic Disaster Recovery. The service maintains metadata consistency and access control definitions that can be quickly restored when underlying data sources recover. Recovery time depends primarily on underlying data services rather than Amazon DataZone itself. Supports customer objectives through metadata catalog preservation of data lineage, business context, and access permissions.

Setup Process

Onboarding: Configure AWS Identity and Access Management Identity Center and AWS Lake Formation permissions, then create domains through the AWS Management Console. Create publishing projects, environment profiles, and environments using built-in blueprints for data lakes, warehouses, or Amazon SageMaker. Sample scripts and quickstart workflows available for AWS Glue and Amazon Redshift sources.

Offboarding: Delete assets from project inventories, unpublish data from catalogues, and remove subscription grants. Domain administrators can delete entire domains to remove all associated data while following AWS data retention policies.

Data Management

- **Storage Options:** Metadata stored in AWS managed infrastructure with built-in redundancy and durability capabilities
- **Data Removal:** Delete assets from inventories, unpublish from catalogue, remove subscription grants, delete projects and environments
- **Cross-Region Replication:** Regional service relying on AWS infrastructure resilience without explicit cross-Region replication features
- **Encryption:** Secure data access with fine-grained controls using AWS Lake Formation Data Cell Filters and compliance features.

Optimise Cost and Consumption

Amazon DataZone optimises costs through automated metadata generation that reduces manual cataloguing overhead and operational expenses. Data mesh architecture implementation enables decentralised ownership models that reduce centralised infrastructure requirements. Automated data lifecycle policies can be implemented for underlying storage services to manage long-term retention costs. The service's data discovery and sharing capabilities reduce duplicate data storage across organisational teams, maximising resource use efficiency.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/datazone/latest/userguide/what-is-datazone.html
Getting Started: https://docs.aws.amazon.com/datazone/latest/userguide/getting-started.html
API Reference: https://docs.aws.amazon.com/datazone/latest/APIReference/Welcome.html
Service Quotas: https://docs.aws.amazon.com/datazone/latest/userguide/datazone-limits.html

Amazon Detective

General Information

Amazon Detective is a security investigation service that analyses and identifies root causes of security findings in AWS environments. It automatically collects log data and uses ML to create behaviour graphs with interactive visualisations for efficient security investigations.

Main Benefits

Amazon Detective helps eliminate manual log correlation through automated behaviour graph construction, reducing investigation time with ML insights. It offers integration with Amazon GuardDuty and AWS Security Hub, requires no complex configuration, and provides visualisations within two hours of enabling.

Key Use Cases

- **Security Incident Investigation:** Analyse Amazon GuardDuty findings and AWS Security Hub alerts to determine root cause and scope
- **Threat Hunting:** Proactively search for indicators of compromise and suspicious activities across AWS environments
- **Malware Analysis:** Investigate Amazon Elastic Compute Cloud instances for potential malware compromises and understand attack chains
- **Multi-Account Security Monitoring:** Centralise security investigations across multiple AWS accounts through behaviour graphs.

Core Features

- **Behaviour Graph Construction:** Creates linked datasets using ML and graph theory to visualise entity relationships and behaviour
- **Detective Investigations:** Automated investigations using ML models and threat intelligence to analyse AWS Identity and Access Management users and roles
- **Finding Groups:** Consolidates related security findings and events into single navigable timelines for holistic incident analysis
- **Interactive Visualisations:** Provides guided visualisations and generative AI insights to help analysts understand attack patterns
- **Multi-Account Management:** Supports administrator and member account structures with AWS Organizations integration for centralised security monitoring
- **Automatic Data Collection:** Automatically ingests data from Amazon CloudTrail, VPC Flow Logs, and Amazon Elastic Kubernetes Service audit logs
- **GuardDuty Integration:** Ingests findings from Amazon GuardDuty and AWS Security Hub for investigation workflows.

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/detective/

FAQ: https://aws.amazon.com/detective/faqs/

Technical Overview

Backup and Recovery Capabilities

Backup: Amazon Detective does not provide traditional backup capabilities and does not integrate with AWS Backup service. As an analytical service that processes log data to create behaviour graphs, Amazon Detective automatically maintains up to one year of aggregated data for analysis purposes but offers no customer-controlled backup functionality.

Recovery: Amazon Detective does not support disaster recovery or integrate with AWS Elastic Disaster Recovery. As a regional service, Amazon Detective operates independently in each AWS Region and requires separate enablement per region. Amazon Detective does not help customers meet RPO or RTO objectives as it serves as a security investigation tool rather than a data protection service.

Setup Process

Onboarding: Enable Amazon Detective through AWS Management Console, CLI, or API in desired AWS Regions. Detective automatically creates behaviour graphs, begins 30-day free trial, and starts data ingestion with visualisations available within 2 hours. Minimal configuration required.

Offboarding: Disable Amazon Detective through console, API, or CLI to permanently delete behaviour graphs and data. AWS provides Python scripts for bulk account management across multiple AWS Regions during cleanup.

Data Management

- **Storage Options:** Automatically maintains up to one year of aggregated data for analysis in Region-specific behaviour graphs
- **Data Removal:** Permanently and irreversibly deletes behaviour graphs and all associated data when disabling Detective service
- **Cross-Region Replication:** Regional service requiring separate enablement per AWS Region with no cross-Region data replication capabilities
- **Encryption:** AWS manages underlying infrastructure resilience and data durability within each AWS Region as managed service

Optimise Cost and Consumption

Amazon Detective provides tiered pricing that automatically reduces per-GB costs as data volume increases, with significant reductions for high-volume customers. The Usage page enables data volume monitoring to control ingestion costs. Selective member account management allows precise control over which accounts contribute data to behaviour graphs. Regional service enablement controls help optimise costs by disabling Amazon Detective in unused Regions. Automatic data aggregation reduces storage requirements compared to raw log retention.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/detective/latest/userguide/index.html
Getting Started: https://docs.aws.amazon.com/detective/latest/userguide/detective-setup.html
API Reference: https://docs.aws.amazon.com/detective/latest/APIReference/Welcome.html
Service Quotas: https://docs.aws.amazon.com/detective/latest/userguide/regions-limitations.html

Amazon DocumentDB (with MongoDB compatibility)

General Information

Amazon DocumentDB is a fully managed, MongoDB-compatible document database service that stores, queries, and indexes JSON-like documents. It supports existing MongoDB drivers and tools while providing enterprise-grade reliability, automatic scaling, and built-in security features.

Main Benefits

Amazon DocumentDB delivers twice the throughput of traditional managed MongoDB services while providing up to 90% cost savings through serverless configurations. Six-way data replication helps ensures enterprise-grade reliability, while seamless MongoDB compatibility enables migration with existing applications and tools.

Key Use Cases

- **Content Management Systems:** Managing user-generated content with flexible schema requirements for various content types
- **E-commerce Applications:** Handling product catalogues, customer profiles, and order management with complex nested data structures
- **Mobile and Web Applications:** Storing user profiles, session data, and metadata with rapid development cycles
- **IoT Applications:** Collecting and analysing sensor data, device telemetry, and time-series information from devices

Core Features

- **MongoDB Compatibility:** Compatible with MongoDB 3.6, 4.0, 5.0, and 8.0 APIs supporting existing drivers and tools
- **Serverless Configuration:** Automatic capacity scaling with Amazon DocumentDB Capacity Units providing up to 90% cost savings for variable workloads
- **Elastic Clusters:** Virtually limitless scale supporting millions of reads/writes with petabytes of storage capacity across multiple shards
- **Vector Search:** HNSW and IVFFlat indexing methods for semantic search with 30x faster parallel index building
- **ACID Transactions:** Multi-document ACID transactions helping ensure data consistency across operations for complex business requirements
- **Point-in-Time Recovery:** Restore clusters to any second within 1-35 day backup retention with continuous incremental backups
- **Global Clusters:** Multi-Region clusters with primary and up to five secondary AWS Regions for disaster recovery
- **Advanced Query Capabilities:** Aggregation pipelines, geospatial querying, text search, and complex array operations for sophisticated data analysis.

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/documentdb/

FAQ: https://aws.amazon.com/documentdb/faqs/

Technical Overview

Backup and Recovery Capabilities

Backup: Automated continuous incremental backups stored in Amazon Simple Storage Service with 1-35 day retention can be created manually (AWS Management Console, CLI) or automated via AWS Backup with centralised management. Cross-Region snapshot copying and sharing capabilities included.

Recovery: Point-in-time recovery to any second, near-instant crash recovery under 30 seconds, and Global Cluster cross-Region failover. Six-way data replication across Availability Zones (AZ) helps ensure fault tolerance. Supports customer RPO (second-level granularity) and RTO (sub-30 second recovery) objectives.

Setup Process

Onboarding: Create clusters through AWS Management Console, CLI, or CloudFormation templates. Configure Amazon Virtual Private Cloud (Amazon VPC) settings with multi-AZ subnets, instance types, and security groups. Getting started guide enables cluster creation in under five minutes.

Offboarding: Delete clusters through AWS Management Console or CLI, removing all instances and data. Delete individual documents, collections, or databases using MongoDB commands. Manual snapshots require separate deletion for complete data removal.

Data Management

- **Storage Options:** Log-structured storage with separated compute and storage, supporting up to 64TB per instance-based cluster
- **Data Removal:** Delete documents, collections, or databases using MongoDB commands; cluster deletion removes all instances and data
- **Cross-Region Replication:** Global Clusters support primary and up to five secondary AWS Regions with automatic cross-Region data replication
- **Encryption:** Built-in encryption at rest and in transit with AWS Key Management Service integration and VPC isolation security.

Optimise Cost and Consumption

Serverless configurations with Amazon DocumentDB Capacity Units automatically scale based on demand, delivering up to 90% cost savings for variable workloads. Database Savings Plans provide up to 35% savings with 1-year commitments. I/O-Optimised storage configurations offer predictable pricing for intensive applications. Storage auto-scaling optimises resource use while per-second billing with 10-minute minimums enables precise cost control.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/documentdb/latest/developerguide/what-is.html

Getting Started: https://docs.aws.amazon.com/documentdb/latest/developerguide/get-started-guide.html

API Reference: https://docs.aws.amazon.com/documentdb/latest/developerguide/api-reference.html

Service Quotas: https://docs.aws.amazon.com/documentdb/latest/developerguide/limits.html

Amazon DynamoDB

General Information

Amazon DynamoDB is a fully managed, serverless NoSQL database service delivering single-digit millisecond performance at any scale. It supports key-value and document data models with automatic infrastructure management, built-in security, and AWS integration for modern applications.

Main Benefits

Amazon DynamoDB delivers consistent single-digit millisecond performance at any scale with zero infrastructure management. It provides up to 99.999% availability through global tables, automatic scaling without capacity planning, and up to 76% cost savings through optimised pricing models and serverless architecture.

Key Use Cases

- **Mobile and Web Applications:** Session management, user profiles, and gaming leaderboards for millions of users
- **Gaming Applications:** Player data, game state management, and global leaderboards with multi-region support
- **IoT Applications:** Time-series data collection and device management for sensor networks and IoT systems
- **Financial Services:** Transaction processing, fraud detection, and trading platforms requiring ACID transactions.

Core Features

- **Serverless Architecture:** Zero infrastructure management with automatic scaling and pay-per-use pricing model
- **Single-digit Millisecond Performance:** Consistent low-latency performance at any scale with optional microsecond latency via DAX
- **Global Tables:** Multi-active, multi-Region replication for globally distributed applications designed for 99.999% availability
- **ACID Transactions:** Support for complex business logic with atomicity, consistency, isolation, and durability
- **DynamoDB Streams:** Real-time change data capture for event-driven architectures and data processing workflows
- **Point-in-time Recovery:** Continuous backups with recovery to any second within the last 35 days
- **Auto Scaling:** Automatic capacity adjustment based on traffic patterns for cost optimisation and performance
- **Secondary Indexes:** Global and local secondary indexes for flexible query patterns and data access.
- DynamoDB Accelerator (DAX): In-memory caching for DynamoDB with microsecond response times.

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/dynamodb/
FAQ: https://aws.amazon.com/dynamodb/faqs/

Technical Overview

Backup and Recovery Capabilities

Backup: Continuous backups with point-in-time recovery for up to 35 days with per-second granularity, and on-demand backups for long-term retention stored redundantly across multiple Availability Zones. Backups can be created manually (AWS Console, CLI, APIs) or automated via AWS Backup with centralised policies and cross-region copying.

Recovery: Point-in-time recovery with RPO up to 1 second and RTO in minutes, global tables providing active-active multi-Region replication with automatic failover and designed for 99.999% availability. Multi-Availability Zone architecture enables sub-minute recovery from infrastructure failures. Supports near-zero RTO (active-active architecture) and zero RPO (strong consistency mode) objectives.

Setup Process

Onboarding: Start immediately through AWS Management Console, CLI, NoSQL Workbench, or SDKs without infrastructure provisioning. Create tables with primary keys, choose on-demand or provisioned capacity modes, and configure encryption, streams, or global tables. AWS provides tutorials and migration tools.

Offboarding: Delete tables through AWS Management Console, CLI, or APIs to permanently remove data. Create backups or export to Amazon Simple Storage Service before deletion. Remove global table replicas first, then clean up associated streams, Amazon DynamoDB Accelerator (DAX) clusters, and backups separately.

Data Management

- **Storage Options:** Standard and Standard-IA table classes with automatic scaling and pay-per-use pricing models
- **Data Removal:** Delete tables through Console, CLI, or APIs with permanent data removal and backup options
- **Cross-Region Replication:** Global tables provide active-active multi-Region replication designed for 99.999% availability and automatic failover
- **Encryption:** Server-side encryption at rest and in transit with AWS Key Management Service integration for comprehensive security.

Optimise Cost and Consumption

Amazon DynamoDB's auto scaling automatically adjusts provisioned capacity based on traffic patterns to prevent over-provisioning. Standard-IA table class reduces storage costs by up to 60% for infrequently accessed data. Reserved capacity provides up to 76% savings for predictable workloads, while Database Savings Plans offer additional discounts with usage commitments. On-Demand mode scales to zero when unused, optimising costs for unpredictable traffic patterns.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/Introduction.html
Getting Started:
https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/GettingStartedDynamoDB.html
API Reference: https://docs.aws.amazon.com/amazondynamodb/latest/APIReference/Welcome.html
Service Quotas: https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/ServiceQuotas.html

Amazon Elastic Block Store (Amazon EBS)

General Information

Amazon Elastic Block Store (Amazon EBS) provides scalable, high-performance block storage for Amazon Elastic Compute Cloud (Amazon EC2) instances. Amazon EBS volumes function as persistent storage devices that can be attached to instances for databases, file systems, and applications. The service offers multiple-volume types optimised for different performance and cost requirements.

Main Benefits

Amazon EBS is designed to deliver 99.8% to 99.999% durability with persistent storage that survives instance failures. The service enables dynamic scaling of performance independently from capacity without downtime, offers up to 20% cost savings through gp3 migration, and provides Fast Snapshot Restore for instant data recovery without initialisation latency.

Key Use Cases

- **Database Storage:** Primary storage for relational and NoSQL databases including MySQL, PostgreSQL, MongoDB, and Oracle
- **File Systems and Applications:** Boot volumes for Amazon EC2 instances, application data storage, and distributed applications
- **Big Data and Analytics:** Storage for data warehouses, Hadoop clusters, log processing, and MapReduce workloads
- **Virtual Desktops and Development:** Persistent storage for VDI deployments, development and testing environments.

Core Features

- **Multiple Volume Types:** Six volume types including gp3, gp2, io2 Block Express, io1, st1, and sc1
- **EBS Snapshots:** Point-in-time incremental backups stored in Amazon Simple Storage Service (Amazon S3) for volume creation and data restoration
- **Elastic Volumes:** Dynamically modify volume size, performance, and type in production without service interruption
- **Encryption:** Data-at-rest and data-in-transit encryption using AWS Key Management Service (AWS KMS) with AES-256 encryption support
- **Multi-Attach:** Attach single Amazon EBS volume to multiple Amazon EC2 instances simultaneously within same Availability Zone (AZ)
- **Fast Snapshot Restore:** Enables instant data restoration from snapshots without initialisation latency penalty
- **Cross-Region Replication:** Copy snapshots across AWS Regions for geographic redundancy and disaster recovery scenarios
- **Automated Lifecycle Management:** Amazon Data Lifecycle Manager automates snapshot creation, retention, and deletion policies.

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/ebs/

FAQ: https://aws.amazon.com/ebs/faqs/

Technical Overview

Backup and Recovery Capabilities

Backup: Amazon EBS snapshots provide incremental, point-in-time backups stored in Amazon S3 with cross-AZ replication can be created manually (AWS Management Console, CLI, APIs) or automated via Data Lifecycle Manager with 15-minute scheduling frequencies. AWS Backup offers centralised management with cross-Region copy capabilities and compliance reporting.

Recovery: Fast Snapshot Restore enables instant volume initialisation without latency penalties, cross-Region snapshot copying for disaster recovery, and AWS Elastic Disaster Recovery integration for complete application recovery. Multi-Attach supports high-availability architectures. Supports customer RPO (minutes via frequent snapshots) and RTO (minutes via Fast Snapshot Restore) objectives.

Setup Process

Onboarding: Create volumes through Amazon EC2 console, AWS CLI, or APIs, then attach to instances within the same Availability Zone. Configure encryption settings, establish backup policies, and implement AWS Identity and Access Management permissions. Best practices include right-sizing volumes and implementing regular snapshot schedules.

Offboarding: Detach volumes from Amazon EC2 instances and permanently delete them. Create final snapshots for data retention before deletion. AWS automatically handles encryption key cleanup, with charges stopping immediately upon deletion.

Data Management

- **Storage Options:** Six volume types including gp3, gp2, io2 Block Express, io1, st1, and sc1
- **Data Removal:** Detach volumes from instances and permanently delete them; supports secure data sanitisation procedures
- **Cross-Region Replication:** Snapshots can be copied across AWS Regions for geographic redundancy and disaster recovery
- **Encryption:** Data-at-rest and data-in-transit encryption using AWS KMS with AES-256 encryption supporting managed keys.

Optimise Cost and Consumption

Amazon EBS Snapshot Archive delivers 75% cost reduction for long-term snapshot storage while Amazon Data Lifecycle Manager automates snapshot retention policies to prevent indefinite accumulation. Independent IOPS and throughput scaling on gp3 and io2 volumes help prevent over-provisioning, and right-sizing tools match volumes to actual usage patterns. Volume cloning enables efficient testing environments, and automated auditing identifies unattached volumes for deletion.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/ebs/latest/userguide/index.html

Getting Started: https://docs.aws.amazon.com/ebs/latest/userguide/ebs-getting-started.html

API Reference: https://docs.aws.amazon.com/AWSEC2/latest/APIReference/API_Operations_Amazon_Elastic_Block_Store.html

Service Quotas: https://docs.aws.amazon.com/ebs/latest/userguide/ebs-resource-quotas.html

Amazon Elastic Compute Cloud (EC2)

General Information

Amazon Elastic Compute Cloud (EC2) provides secure, resizable compute capacity in the cloud through virtual servers called instances. Users can launch and manage instances with complete control over operating systems and applications. Amazon EC2 offers over 750 instance types optimised for different workload requirements.

Main Benefits

Amazon EC2 delivers unmatched flexibility with over 750 instance types and multiple processor options. Spot Instances provide up to 90% cost savings, while AWS Graviton instances offer 40% better price performance. Integration with over 200 AWS services creates a comprehensive cloud environment with consumption-based pricing.

Key Use Cases

- **Web Application Hosting:** Run web applications, APIs, and microservices with scalable compute capacity
- **High-Performance Computing:** Scientific computing, financial modelling, batch processing, and simulation workloads requiring powerful processors
- **Enterprise Applications:** Business-critical applications, databases, enterprise resource planning systems, and development environments with customisable configurations
- **Big Data Analytics:** Data processing, ML training, real-time analytics, and distributed computing workloads.

Core Features

- **Instance Types:** Over 750 instance types across 7 categories with various processor, memory, storage configurations
- **Amazon Machine Images:** Pre-configured templates containing operating systems, applications, and software for quick instance deployment
- **Elastic Block Store:** High-performance block storage volumes with multiple types and encryption capabilities persisting independently
- **Auto Scaling:** Automatically adjusts instance numbers based on demand, maintaining availability while optimising costs
- **Security Groups:** Virtual firewalls controlling inbound and outbound traffic at protocol and port access levels
- **Elastic Load Balancing:** Distributes incoming traffic across multiple instances ensuring high availability and fault tolerance
- **Spot Instances:** Access spare Amazon EC2 capacity at up to 90% discount for fault-tolerant workloads
- **Enhanced Networking:** High packet-per-second performance with low latency and SR-IOV Elastic Fabric Adapter support.

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/ec2/

FAQ: https://aws.amazon.com/ec2/faqs/

Technical Overview

Backup and Recovery Capabilities

Backup: Amazon Elastic Block Store (Amazon EBS) snapshots and Amazon Machine Images (AMIs) create point-in-time backups stored in Amazon Simple Storage Service designed to provide 99.999999999% durability, created manually through console/CLI or automated via AWS Backup with hourly scheduling. Cross-Region replication and centralised backup management with retention policies are supported.

Recovery: AMIs enable rapid instance recreation, Auto Scaling groups provide automatic replacement within minutes, and AWS Elastic Disaster Recovery offers continuous replication. Multi-Availability Zone deployments provide built-in fault tolerance. Supports customer RPO (hourly snapshots) and RTO (minutes replacement) objectives.

Setup Process

Onboarding: Launch instances through AWS Management Console wizard, CLI, SDKs, or AWS CloudFormation. Configure AMI selection, instance types, security groups, and key pairs. Use launch templates for consistent deployments and Auto Scaling for production workloads.

Offboarding: Terminate instances to delete instance store data and detach Amazon EBS volumes. Delete Amazon EBS volumes, snapshots, AMIs, key pairs, and security groups. AWS CloudFormation templates automate complete resource cleanup.

Data Management

- **Storage Options:** Amazon EBS persistent block storage with multiple volume types and instance store temporary storage
- **Data Removal:** Terminate instances to delete instance store data, separately delete EBS volumes and snapshots
- **Cross-Region Replication:** Amazon EBS snapshots can be copied across AWS Regions, AWS Backup supports cross-Region replication
- **Encryption:** EBS volumes support encryption capabilities with multiple encryption options for data protection.

Optimise Cost and Consumption

Spot Instances deliver up to 90% savings for fault-tolerant workloads, while Savings Plans provide up to 72% discounts for consistent usage commitments. AWS Graviton instances offer 40% better price performance. Auto Scaling automatically adjusts capacity to match demand, preventing over-provisioning. Amazon EC2 Instance Type Finder and AWS Compute Optimiser recommend optimal configurations based on workload patterns, while hibernation pauses instances to preserve memory state and reduce costs.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/concepts.html

Getting Started: https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/EC2_GetStarted.html

API Reference: https://docs.aws.amazon.com/AWSEC2/latest/APIReference/

Service Quotas: https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ec2-resource-limits.html

Amazon Elastic Container Registry (Amazon ECR)

General Information

Amazon Elastic Container Registry (Amazon ECR) is a fully managed container image registry service that provides secure, scalable storage for Docker and OCI-compatible artifacts. Amazon ECR supports both private and public repositories with deep AWS integration for streamlined containerised application workflows.

Main Benefits

Amazon ECR helps eliminate infrastructure management overhead and is designed to deliver 99.9% uptime SLA and 99.999999999% durability through Amazon Simple Storage Service (Amazon S3)-backed storage. Native integration with Amazon Elastic Compute Cloud, Amazon Elastic Kubernetes Service, and AWS Fargate streamlines container workflows, while enterprise-grade security includes vulnerability scanning, encryption, and AWS Identity and Access Management (AWS IAM) -based access controls for container image management.

Key Use Cases

- **Container Image Storage:** Centralised repository for storing and distributing Docker images within organisations
- **CI/CD Integration:** Automated building, testing, and deployment of containerised applications with various tools
- **Microservices Architecture:** Supporting containerised microservices deployments across AWS container orchestration platforms
- **Security and Compliance:** Image vulnerability scanning, access control, and managed signing for enhanced security.

Core Features

- **Lifecycle Policies:** Automated cleanup of unused images based on defined rules for image age or tag criteria
- **Image Scanning:** Vulnerability scanning for container images with scan-on-push capability using Clair CVE database
- **Cross-Region Replication:** Automatic replication of images across AWS Regions for high availability and reduced latency
- **Pull Through Cache:** Cache repositories from upstream registries with periodic synchronisation to help ensure up-to-date images
- **Managed Signing:** Automatic cryptographic signature generation for pushed images using AWS Signer for enhanced security
- **Repository Creation Templates:** Define settings for automatically created repositories including encryption, policies, and tags
- **Encryption:** Image encryption at rest using AWS Key Management Service (AWS KMS) and in-transit encryption for secure image storage
- **IAM Integration:** Fine-grained access control using AWS IAM roles and policies for repository and image access.

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/ecr/

FAQ: https://aws.amazon.com/ecr/faqs/

Technical Overview

Backup and Recovery Capabilities

Backup: Cross-Region replication designed for 99.999999999% durability through Amazon S3-backed storage can be created manually (AWS Management Console, CLI, SDKs) or automated via registry-level configuration with continuous synchronisation. Amazon ECR APIs enables programmatic copying to multiple AWS Regions or accounts for additional backup strategies.

Recovery: Multi-Availability Zone architecture provides automatic failover within minutes, cross-Region replication enables rapid disaster recovery, and AWS Application Recovery Controller orchestrates complete application recovery. Multi-Region simultaneous replication helps ensure high availability. Supports customer RPO (seconds to minutes through real-time replication) and RTO (minutes via automated failover) objectives.

Setup Process

Onboarding: Create repositories using AWS Management Console, CLI, or SDKs, then authenticate Docker CLI with get-authorisation-token command. Configure encryption, policies, and tags using repository creation templates. Build images locally, tag with Amazon ECR uniform resource identifier, and push using standard Docker commands.

Offboarding: Delete individual images using the AWS Management Console, CLI, or SDKs removing tags and metadata. Repository deletion requires force flag for permanent removal. Lifecycle policies automate cleanup with cross-Region replicated images requiring separate regional deletion.

Data Management

- **Storage Options:** Amazon S3-backed storage designed for 99.999999999% durability supporting private and public repositories for Docker/OCI images
- **Data Removal:** Delete images via AWS Management Console, CLI, or SDKs; repository deletion with force flag permanently removes all data
- **Cross-Region Replication:** Automatic replication across AWS Regions configured at registry level for high availability and disaster recovery
- **Encryption:** Image encryption at rest using AWS KMS and in-transit encryption for secure container image storage.

Optimise Cost and Consumption

Lifecycle policies automatically delete unused images based on configurable age, count, or tag criteria to prevent storage cost accumulation. Pull-through cache repositories reduce bandwidth costs by caching upstream registries with periodic synchronisation. Image layer deduplication optimises storage through consolidated repositories. Virtual Private Cloud endpoints help eliminate internet gateway charges for private network access. Amazon ECR usage reports enable regular auditing to identify optimisation opportunities, while strategic regional placement reduces data transfer costs between container workloads and repositories.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/AmazonECR/latest/userguide/what-is-ecr.html

Getting Started: https://docs.aws.amazon.com/AmazonECR/latest/userguide/getting-started-cli.html

API Reference: https://docs.aws.amazon.com/AmazonECR/latest/APIReference/

Service Quotas: https://docs.aws.amazon.com/AmazonECR/latest/userguide/service-quotas.html

Amazon Elastic Container Service (Amazon ECS)

General Information

Amazon Elastic Container Service (Amazon ECS) is a fully managed container orchestration service that helps deploy, manage, and scale containerised applications. It eliminates infrastructure management overhead while providing flexible compute options including serverless AWS Fargate and Amazon Elastic Compute Cloud (Amazon EC2) instances.

Main Benefits

Amazon ECS helps eliminate infrastructure management overhead while providing flexible compute options and seamless AWS service integration. Features like AWS Fargate Spot deliver up to 70% cost savings, sub-minute RTO for multi-Availability Zone (AZ) deployments, and zero-downtime blue/green deployments, enabling teams to focus on applications rather than infrastructure complexity.

Key Use Cases

- **Web Applications and APIs:** Long-running stateless applications requiring high availability and automatic scaling
- **Batch Processing:** Containerised batch jobs for data processing, ETL operations, and computational workloads
- **Microservices Architecture:** Breaking applications into smaller, independently deployable services with service discovery
- **Hybrid Deployments:** Running containerised workloads across AWS and on-premises infrastructure using Availability Zone ECS Anywhere.

Core Features

- **Task Definitions:** Blueprint specifying container configurations, resource requirements, and networking for applications
- **Auto Scaling:** Automatic scaling of container instances and services based on demand patterns
- **Service Connect:** Service discovery and connectivity with traffic monitoring for inter-service communication
- **Blue/Green Deployments:** Built-in deployment strategy enabling safer software updates with zero downtime
- **Multiple Launch Types:** Support for AWS Fargate serverless, Amazon EC2 instances, and Amazon ECS Managed Instances
- **ECS Anywhere:** Extends Amazon ECS capabilities to run containers on on-premises infrastructure
- **Express Mode:** Simplified deployment automatically provisioning all required infrastructure components with single command
- **Clusters:** Logical grouping of compute resources where containerised applications run and scale.

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/ecs/
FAQ: https://aws.amazon.com/ecs/faqs/

Technical Overview

Backup and Recovery Capabilities

Backup: Amazon ECS lacks native backup capabilities but integrates with AWS Backup for associated resources like Amazon Elastic Block Store (Amazon EBS) volumes. Container images stored in Amazon ECR support cross-Region replication. Task definitions and service configurations can be version-controlled and backed up through Infrastructure as Code (IaC) tools like AWS CloudFormation and CDK.

Recovery: Multi-Region deployments enable quick recovery by launching new tasks in healthy Regions. Multi-AZ deployments provide fault tolerance with automatic health checks. Supports customer RPO objectives through real-time data replication and RTO objectives with sub-minute recovery for Availability Zone failures and minute-level recovery for cross-Region deployments.

Setup Process

Onboarding: Express Mode provides fastest setup by automatically provisioning infrastructure with a single command. Traditional setup involves creating clusters, task definitions, and services through AWS Management Console, CLI, or Infrastructure as Code tools like CloudFormation and CDK.

Offboarding: Delete services, task definitions, and clusters through console or API. Amazon ECS automatically scales services to zero before removal and provides automated cleanup for stopped tasks and unused Docker images.

Data Management

- **Storage Options:** Container images stored in Amazon ECR; Amazon EBS volumes for persistent data with container instances
- **Data Removal:** Delete services, task definitions, clusters through console or API; automated cleanup for stopped tasks
- **Cross-Region Replication:** Container images in ECR support cross-Region replication; multi-Region deployments for disaster recovery
- **Encryption:** Integrates with AWS encryption services through Amazon Virtual Private Cloud, AWS Identity and Access Management, and underlying compute infrastructure security

Optimise Cost and Consumption

AWS Fargate Spot delivers up to 70% savings for fault-tolerant workloads through interruptible pricing. Amazon ECS Managed Instances automatically right-size and consolidate tasks while removing idle instances. Auto Scaling matches capacity with demand patterns to prevent over-provisioning. Express Mode shares Application Load Balancers across multiple services. Compute Optimiser provides specific recommendations for AWS Fargate services, and Savings Plans optimise predictable workloads.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/AmazonECS/latest/developerguide/
Getting Started: https://docs.aws.amazon.com/AmazonECS/latest/developerguide/Welcome.html
API Reference: https://docs.aws.amazon.com/AmazonECS/latest/APIReference/
Service Quotas: https://docs.aws.amazon.com/AmazonECS/latest/developerguide/service-quotas.html

Amazon Elastic File System (Amazon EFS)

General Information

Amazon Elastic File System (Amazon EFS) provides serverless, fully elastic file storage that automatically scales from gigabytes to petabytes without disrupting applications. It offers scalable, shared storage for AWS compute services using NFSv4 protocols, helping eliminate the need to provision or manage storage capacity and performance.

Main Benefits

Amazon EFS is designed to deliver 99.999999999% durability with automatic multi-Availability Zone (AZ) replication and up to 219% ROI. It reduces storage costs by up to 97% through intelligent tiering, provides sub-millisecond latency, supports thousands of concurrent connections, and helps eliminate infrastructure management overhead with serverless scaling.

Key Use Cases

- **Big Data and Analytics:** Processing large datasets with concurrent access from multiple compute instances
- **Content Management:** Sharing website content, media files, and documents across web servers with low latency
- **Media Processing Workflows:** Supporting video encoding, image processing, and media operations requiring shared file access
- **Application Development:** Providing shared storage for development environments, continuous integration and continuous deployment, pipelines, and testing frameworks.

Core Features

- **Elastic Scaling:** Automatically scales from gigabytes to petabytes without disrupting applications or capacity planning
- **Multiple Storage Classes:** Standard, Infrequent Access, and Archive classes with automatic lifecycle management capabilities
- **Performance Modes:** General Purpose and Max I/O modes for different latency and throughput requirements
- **Throughput Modes:** Elastic, Provisioned, and Bursting throughput modes for optimal workload performance optimisation
- **Encryption:** Data encryption at rest and in transit using AWS Key Management Service (AWS KMS) for security
- **Cross-Region Replication:** Disaster recovery with continuous synchronisation and 15-minute RPO for business continuity
- **Access Points:** Secure, fine-grained access control to specific directories and files within file systems
- **AWS Backup Integration:** Automated backup and restore capabilities with configurable retention policies and incremental backups.

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/efs/

FAQ: https://aws.amazon.com/efs/faqs/

Technical Overview

Backup and Recovery Capabilities

Backup: Amazon EFS performs incremental backups designed for 99.999999999% durability and automatic multi-Availability Zone replication can be created manually (Amazon EFS console, CLI, API) or automated via AWS Backup with daily scheduling. Backup rates reach 2,000 files per second without affecting operations.

Recovery: Amazon EFS enables point-in-time recovery with restore rates of 1,500 files per second, cross-Region replication for disaster recovery, and continuous availability during single Availability Zone outages. Regional file systems provide automatic fault tolerance. Supports customer RPO (15-minute cross-Region synchronisation) and RTO (fast elastic scaling) objectives.

Setup Process

Onboarding: Amazon EFS deployment uses Amazon Elastic Compute Cloud launch wizard for one-click creation, Amazon EFS console for custom configurations, or AWS CLI/API for programmatic setup. Quick Create automatically configures Regional file systems with recommended performance settings. Requires virtual private cloud setup and security groups for Network File System (NFS) protocol mounting.

Offboarding: File system deletion through Amazon EFS console, CLI, or API requires unmounting from all instances first, then deleting mount targets in each Availability Zone. AWS recommends creating backups before irreversible deletion. Access points and replication configurations require separate removal.

Data Management

- **Storage Options:** Standard, Infrequent Access, and Archive classes with automatic lifecycle management and intelligent tiering
- **Data Removal:** Irreversible file system deletion via console, CLI, or API after unmounting and removing mount targets
- **Cross-Region Replication:** Continuous synchronisation with 15-minute RPO for disaster recovery and configurable failover scenarios
- **Encryption:** Data encryption at rest and in transit using AWS KMS with comprehensive key management.

Optimise Cost and Consumption

Amazon EFS Intelligent-Tiering automatically moves files between storage classes based on access patterns, reducing costs by up to 97%. Lifecycle policies transition rarely accessed data to Archive storage tiers. One Zone file systems provide cost savings for development environments. Elastic throughput mode automatically scales performance to prevent over-provisioning charges during low-usage periods. Shared file systems across multiple applications maximise throughput use and resource efficiency.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/efs/latest/ug/whatisefs.html

Getting Started: https://docs.aws.amazon.com/efs/latest/ug/getting-started.html

API Reference: https://docs.aws.amazon.com/efs/latest/ug/api-reference.html

Service Quotas: https://docs.aws.amazon.com/efs/latest/ug/limits.html

Amazon Elastic Kubernetes Service (Amazon EKS)

General Information

Amazon Elastic Kubernetes Service (Amazon EKS) is a fully managed Kubernetes service that helps remove the need to install, operate, or maintain your own Kubernetes control plane. Amazon EKS automatically manages control plane availability, scalability, and patching across multiple Availability Zones (AZs) for high availability.

Main Benefits

Amazon EKS helps eliminate operational complexity of self-managed Kubernetes while maintaining full compatibility, enabling teams to focus on application development rather than infrastructure management. With automatic control plane scaling, patching, and multi-AZ high availability, Amazon EKS delivers enterprise-grade security through deep AWS integrations and supports cost optimisation with up to 90% savings using Spot Instances.

Key Use Cases

- **Microservices Architecture:** Building and deploying cloud-native applications using microservices patterns with Kubernetes service discovery
- **Machine Learning Workloads:** Running ML training and inference workloads with popular frameworks and GPU support
- **Batch Processing:** Executing cost-effective batch processing and big data workloads using Amazon Elastic Compute Cloud (Amazon EC2) Spot Instances
- **Hybrid Deployments:** Deploying applications across cloud and on-premises environments using Amazon EKS Hybrid Nodes

Core Features

- **Managed Control Plane:** Fully managed Kubernetes control plane with automatic scaling, patching, and multi-AZ high availability
- **EKS Auto Mode:** Serverless compute automatically managing both control plane and worker nodes with automated patching
- **Multiple Compute Options:** Support for Amazon EC2, AWS Fargate, Nitro, Graviton instances, Spot Instances, and managed node groups
- **Advanced Networking:** Amazon Virtual Private Cloud native networking, IPv6 support, load balancing integration, and application networking capabilities
- **Security and Compliance:** Amazon EKS Pod Identity, AWS Identity and Access Management role-based access control, encryption at rest/transit, and container image signature verification
- **Hybrid Deployment Support:** Amazon EKS Anywhere, Hybrid Nodes, AWS Outposts integration, and Amazon EKS Connector for external clusters
- **GitOps Capabilities:** Fully managed cluster features including Argo CD, AWS Controllers for Kubernetes, and Kube Resource Orchestrator
- **Observability Integration:** Amazon Managed Prometheus, Amazon CloudWatch Container Insights, logging, and cost allocation tagging support
- Amazon EKS Anywhere: Create and operate Kubernetes clusters on on-premises infrastructure.

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/eks/

FAQ: https://aws.amazon.com/eks/faqs/

Technical Overview

Backup and Recovery Capabilities

Backup: Amazon EKS integrates with AWS Backup to capture Kubernetes manifests, secrets, config maps, and persistent volumes from Amazon Elastic Block Store (Amazon EBS), Amazon Elastic File System (Amazon EFS), and Amazon Simple Storage Service (Amazon S3). Backups can be created on-demand or automated via AWS Backup with configurable scheduling. Features include incremental backups, lifecycle management, encryption, cross-Region replication, and immutable backup vaults for compliance.

Recovery: Amazon EKS creates composite recovery points enabling full cluster restoration or selective namespace recovery across AWS Regions and accounts. Multi-AZ control plane architecture provides high availability. Supports RPO objectives through automated backup scheduling and RTO objectives via cross-Region replication and active-active deployments.

Setup Process

Onboarding: Deploy clusters using eksctl command-line utility, AWS Management Console, AWS CLI, or infrastructure-as-code tools like AWS CloudFormation and Terraform. Amazon EKS Auto Mode automatically manages control plane and worker nodes. Choose between AWS Fargate serverless containers or Amazon EC2-based managed node groups.

Offboarding: Delete workloads and persistent volumes first, then remove node groups and terminate clusters via AWS Management Console, CLI, or eksctl. Explicitly delete persistent data from Amazon EBS, Amazon EFS, or Amazon S3 resources to prevent orphaned charges.

Data Management

- **Storage Options:** Amazon EBS and Amazon EFS for persistent storage, Amazon S3 integration, and Amazon Elastic Container Registry for container images
- **Data Removal:** Delete workloads and persistent volumes first, then remove clusters via AWS Management Console, CLI, or eksctl
- **Cross-Region Replication:** Supports cross-Region and cross-account backup copies, Amazon EBS snapshots, Amazon EFS replication, and Amazon S3 replication
- **Encryption:** Provides encryption at rest and in transit with AWS Backup encryption and immutable vaults

Optimise Cost and Consumption

Amazon EKS Auto Mode provides serverless compute with automatic rightsizing to eliminate idle node costs. Amazon EC2 Spot Instances integration delivers up to 90% savings for batch workloads, while Graviton-based instances offer superior price-performance ratios. Cluster Autoscaler and Vertical Pod Autoscaler optimise resource use dynamically. Cost allocation tagging enables granular tracking through Kubecost integration, and topology-aware routing minimises cross-AZ data transfer expenses for comprehensive cost management.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/eks/latest/userguide/index.html

Getting Started: https://docs.aws.amazon.com/eks/latest/userguide/getting-started.html

API Reference: https://docs.aws.amazon.com/eks/latest/APIReference/Welcome.html

Service Quotas: https://docs.aws.amazon.com/eks/latest/userguide/service-quotas.html

Amazon Elastic MapReduce (Amazon EMR)

General Information

Amazon Elastic MapReduce (Amazon EMR) is a managed big data platform that simplifies running frameworks like Apache Hadoop, Spark, and Hive on AWS. It automates cluster provisioning, configuration, and scaling for cost-effective processing and analysis of vast data amounts.

Main Benefits

Amazon EMR help eliminate big data infrastructure complexity while delivering significant cost savings through per-second billing and up to 90% savings with Spot instances. It provides AWS-optimised open-source frameworks, enterprise-grade security, and automatic scaling for performance and operational efficiency.

Key Use Cases

- **ETL Operations:** Moving and transforming data between AWS data stores for analytics preparation
- **Machine Learning:** Running ML algorithms at scale using Spark MLlib for model training and inference
- **Log Processing:** Processing large volumes of web logs and clickstream data for business insights
- **Real-time Stream Processing:** Analysing streaming data from Kinesis for real-time insights and decision making.

Core Features

- **Managed Scaling:** Automatically scales clusters up and down based on workload demands with configurable policies
- **Multiple Deployment Options:** Offers Amazon EMR on Amazon Elastic Compute Cloud (Amazon EC2), EMR Serverless, EMR on Amazon Elastic Kubernetes Service (Amazon EKS), and EMR on AWS Outposts
- **EMR Studio:** Web-based integrated development environment with Jupyter notebook support and collaborative features for data scientists
- **Open-Source Framework Support:** Pre-configured support for Apache Spark, Hadoop, Hive, Flink, Presto, HBase, and other frameworks
- **Instance Fleet Management:** Supports mixing On-Demand and Spot instances across multiple instance types and Availability Zones
- **Security Integration:** Built-in integration with AWS Identity and Access Management (AWS IAM), Amazon Virtual Private Cloud (Amazon VPC), encryption at rest and in transit, and AWS Lake Formation
- **Data Store Flexibility:** Native integration with Amazon Simple Storage Service (Amazon S3), Hadoop Distributed File System (HDFS), Amazon DynamoDB, and other data stores with optimised connectors
- **AWS Service Integration:** Deep integration with Amazon S3, AWS Glue, Amazon Kinesis, Amazon CloudWatch, and Amazon SageMaker for comprehensive data workflows.

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/emr/
FAQ: https://aws.amazon.com/emr/faqs/

Technical Overview

Backup and Recovery Capabilities

Backup: Amazon EMR integrates with AWS Backup for Amazon Elastic Block Store (Amazon EBS) volumes and relies on Amazon S3 storage designed for 99.999999999% durability. Amazon EBS snapshots can be created manually (AWS Management Console /CLI) or automated via backup policies with configurable scheduling. Amazon S3 provides versioning and cross-Region replication capabilities.

Recovery: Clusters can be rapidly recreated using infrastructure templates, with recovery timeframes from minutes to hours. Amazon EMR supports backup and restore, pilot light, and warm standby disaster recovery patterns. Supports customer RPO (near-zero data loss through Amazon S3 replication) and RTO (rapid cluster provisioning and automation scripts) objectives.

Setup Process

Onboarding: Launch clusters using Amazon EMR console's quick launch option or comprehensive tutorials. Choose between Amazon EC2, Serverless, Amazon EKS, or AWS Outposts deployment options. Amazon EMR provides templates and guided workflows for common use cases with AWS Identity and Access Management (AWS IAM) and VPC configuration.

Offboarding: Terminate active clusters through console or CLI to stop charges. Delete Amazon S3 buckets, Amazon CloudWatch log groups, security configurations, and Amazon EMR-created AWS IAM roles across all Regions for complete cleanup.

Data Management

- **Storage Options:** Native integration with Amazon S3, HDFS, Amazon DynamoDB with optimised connectors for efficient data access
- **Data Removal:** HDFS data automatically deleted when clusters terminate; Amazon S3 data requires explicit deletion
- **Cross-Region Replication:** Supported through Amazon S3 cross-Region replication and multi-AZ storage for disaster recovery
- **Encryption:** Built-in encryption at rest and in transit with AWS IAM and VPC integration

Optimise Cost and Consumption

Amazon EMR provides per-second billing with one-minute minimums for precise resource consumption tracking. Spot instance integration delivers up to 90% savings through intelligent instance fleet management across multiple pools. Managed scaling automatically adjusts cluster size based on workload demands preventing over-provisioning. Amazon EMR Serverless automatically starts and stops resources eliminating idle time charges. Auto-termination policies prevent forgotten clusters from accruing costs while cost allocation tagging enables detailed spending optimisation.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/emr/latest/ManagementGuide/emr-what-is-emr.html
Getting Started: https://docs.aws.amazon.com/emr/latest/ManagementGuide/emr-gs.html
API Reference: https://docs.aws.amazon.com/emr/latest/APIReference/Welcome.html
Service Quotas: https://docs.aws.amazon.com/emr/latest/ManagementGuide/emr-service-limits-what-are.html

Amazon Elastic VMware Service (Amazon EVS)

General Information

Amazon Elastic VMware Service (Amazon EVS) enables organisations to run VMware Cloud Foundation environments directly within Amazon Virtual Private Cloud (Amazon VPC) on Amazon Elastic Compute Cloud (Amazon EC2) bare metal instances. The service provides a ready-to-use Run VMware Cloud Foundation (VCF) environment including ESXi, vCenter Server, NSX-T, and vSAN for rapid deployment without extensive re-architecture.

Main Benefits

Amazon EVS enables lift-and-shift migration without workload modification while preserving existing VMware expertise and tools. Deploy fully functional VCF environments in hours with native integration to over 200 AWS services. Supports Bring Your Own License (BYOL) licensing, elastic scaling, and eliminates infrastructure complexity while providing enterprise-grade reliability.

Key Use Cases

- **VMware Workload Migration:** Lift-and-shift on-premises VMware virtual machines without modification or refactoring
- **Hybrid Cloud Operations:** Extend VMware environments to AWS maintaining consistent operational tools and processes
- **Disaster Recovery:** Implement recovery solutions using vSphere Site Recovery Manager integrated with AWS services
- **Cloud Modernisation:** Use AWS's broad service portfolio while preserving existing VMware investment and expertise.

Core Features

- **VMware Cloud Foundation Integration:** Full VCF 5.2.1 stack including ESXi, vCenter Server, NSX-T, and vSAN
- **Rapid Deployment:** Deploy fully functional VCF environments in hours with guided workflows and automation
- **Self-Management Control:** Full administrative access to optimise virtualisation stack and integrate third-party solutions
- **Flexible Storage Options:** Local vSAN storage and external options like Amazon FSx for NetApp ONTAP
- **Hybrid Connectivity:** AWS Direct Connect and Virtual Private Networks (VPN) connectivity with VMware HCX support for migration
- **AWS Service Integration:** Native integration with over 200 AWS services while maintaining VMware tools
- **Elastic Scaling:** Add or remove hosts based on capacity requirements to match workload demands
- **BYOL Support:** BYOL options to use existing VMware license investments and reduce costs

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/evs/

FAQ: https://aws.amazon.com/evs/faqs/

Technical Overview

Backup and Recovery Capabilities

Backup: Amazon EVS supports VMware native backup solutions, third-party integrations, and AWS services including vSAN-based local backup and Amazon FSx for NetApp ONTAP with native backup capabilities. Customers can use Amazon Simple Storage Service for long-term data retention and cross-Region backup strategies through AWS Management Console or third-party tools.

Recovery: VMware Site Recovery Manager enables orchestrated disaster recovery with cross-Region replication and multi-Availability Zone configurations for high availability. NetApp BlueXP disaster recovery and AWS Elastic Disaster Recovery provide comprehensive protection. Supports customer RPO objectives (vSphere replication for low RPO) and RTO objectives (scalable recovery infrastructure for rapid provisioning).

Setup Process

Onboarding: Deploy through AWS Management Console using guided workflows after completing prerequisites including VCF licenses, Amazon EC2 capacity reservations, and VPC configuration with Route Server. Automatically provisions VCF management appliances and ESXi hosts, retrieving credentials from AWS Secrets Manager.

Offboarding: Delete all hosts first, then environment via AWS Management Console or CLI. Service automatically removes VCF appliances, cleans VPC resources, deletes AWS Secrets Manager credentials, and removes Route Server components for complete cleanup.

Data Management

- **Storage Options:** Local vSAN storage and external options like Amazon FSx for NetApp ONTAP
- **Data Removal:** Automatic cleanup of VCF appliances, VPC resources, and AWS Secrets Manager credentials during offboarding
- **Cross-Region Replication:** Supported for disaster recovery with Site Recovery Manager and AWS services integration
- **Encryption:** VMware native encryption capabilities through vSAN and NSX-T with AWS integration support.

Optimise Cost and Consumption

Amazon EVS offers flexible consumption models with one-year and three-year commitment options for predictable cost savings. BYOL capabilities use existing VMware investments while AWS Instance Savings Plans and Reserved Instances optimise underlying Amazon EC2 infrastructure costs. Elastic scaling enables adding or removing hosts based on capacity requirements to prevent over-provisioning, while flexible storage options between local vSAN and external Amazon FSx optimise storage expenses.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/evs/latest/userguide/index.html

Getting Started: https://docs.aws.amazon.com/evs/latest/userguide/getting-started.html

API Reference: https://docs.aws.amazon.com/evs/latest/APIReference/index.html

Service Quotas: https://docs.aws.amazon.com/evs/latest/userguide/service-quotas-evs.html

Amazon ElastiCache

General Information

Amazon ElastiCache is a fully managed in-memory data store service supporting Valkey, Memcached, and Redis OSS engines. It provides sub-millisecond response times by caching frequently accessed data, reducing database loads and improving application performance with automatic scaling capabilities.

Main Benefits

Amazon ElastiCache delivers sub-millisecond response times supporting millions of operations per second, eliminating operational overhead of self-managed caching solutions. Designed for 99.99% availability SLA, automatic failover, and enterprise-grade security including HIPAA and PCI DSS compliance, it provides reliable high-performance caching for mission-critical applications.

Key Use Cases

- **Session Caching:** Store user session data with high performance and low latency access
- **Database Caching:** Cache query results to reduce database load and improve response times
- **Real-Time Analytics:** Power gaming leaderboards requiring sub-millisecond response times for competitive applications
- **AI Applications:** Enable semantic caching and memory mechanisms for chatbots and AI assistants.

Core Features

- **Multi-Engine Support:** Support for Valkey, Memcached, and Redis OSS engines with full compatibility
- **ElastiCache Serverless:** Fully managed serverless caching that automatically scales without capacity planning
- **Multi-AZ Failover:** High availability across multiple Availability Zones (AZs) with automatic failover capability
- **Read Replicas:** Up to five read replicas per primary node for improved performance
- **Global Datastore:** Cross-Region replication for disaster recovery and global data access
- **Data Tiering:** Cost optimisation using SSDs for least frequently used data storage
- **Vector Search:** Low-latency vector search capabilities for AI and ML applications
- **Amazon Virtual Private Cloud (VPC) Support:** Network isolation within Amazon VPC with security group controls.

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/elasticache/

FAQ: https://aws.amazon.com/elasticache/faqs/

Technical Overview

Backup and Recovery Capabilities

Backup: Automatic and manual snapshots stored in Amazon Simple Storage Service (Amazon S3) with configurable retention periods can be created manually (AWS Management Console, CLI, API) or automated via daily scheduling. Supports up to 24 manual backups per day for serverless caches and 20 for node-based clusters.

Recovery: Point-in-time recovery capabilities, Global Datastore cross-Region replication, and multi-AZ automatic failover with RTO under 1 minute. Read replica promotion enables disaster recovery scenarios. Supports customer RPO (under 1 second with Global Datastore) and RTO (minutes for cross-Region scenarios) objectives.

Setup Process

Onboarding: Deploy via AWS Management Console, CLI, or APIs after configuring AWS Identity and Access Management permissions. Choose serverless or node-based clusters, select engine (Valkey, Memcached, Redis OSS), configure VPC settings and security groups. AWS Free Tier offers 750 hours for testing.

Offboarding: Delete clusters through AWS Management Console, CLI, or API with automatic final backups for replication groups. Clean up subnet groups, parameter groups, and security groups. Manually delete Amazon S3 backups to avoid ongoing charges.

Data Management

- **Storage Options:** In-memory storage with data tiering using SSDs for least frequently used data
- **Data Removal:** Data permanently removed from memory upon cluster deletion via AWS Management Console, CLI, or API
- **Cross-Region Replication:** Global Datastore enables cross-Region replication for disaster recovery and global access
- **Encryption:** Encryption at rest and in transit with AWS Key Management Service integration for data security.

Optimise Cost and Consumption

Amazon ElastiCache Serverless automatically scales with pay-per-use pricing based on data storage and compute consumption. Data tiering on R6gd nodes moves infrequently accessed data to cost-effective SSD storage. Reserved instances deliver up to 72% savings, while Database Savings Plans provide 35% discounts. The Valkey engine offers 20-33% lower pricing, and rightsizing tools identify underused resources for optimal capacity matching.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/AmazonElastiCache/latest/dg/index.html

Getting Started: https://docs.aws.amazon.com/AmazonElastiCache/latest/dg/GettingStarted.html

API Reference: https://docs.aws.amazon.com/AmazonElastiCache/latest/APIReference/

Service Quotas: https://docs.aws.amazon.com/general/latest/gr/elasticache-service.html

Amazon EventBridge

General Information

Amazon EventBridge is a serverless event bus service that connects application components through events, enabling scalable event-driven architectures. It provides event buses for routing, EventBridge Pipes for point-to-point integrations, and EventBridge Scheduler for task scheduling.

Main Benefits

Amazon EventBridge helps eliminate custom integration code with native connections to over 200 AWS services and over 45 SaaS partners. It provides exactly-once processing guarantees, automatic scaling, and reduced operational overhead. The AWS Free Tier includes 5 million custom events monthly, significantly reducing development and maintenance costs.

Key Use Cases

- **Application Integration:** Decoupling microservices and distributed systems by routing events between application components
- **IT Automation:** Automating operational tasks and workflows by responding to infrastructure changes in real-time
- **SaaS Integration:** Connecting over 45 SaaS partner applications to AWS services without custom connection code
- **Real-time Stream Processing:** Processing and routing streaming data with content-based filtering across multiple targets.

Core Features

- **Event Buses:** Serverless routers that receive events from multiple sources and deliver to multiple targets
- **EventBridge Pipes:** Point-to-point integrations for processing events with advanced transformations and enrichment capabilities
- **EventBridge Scheduler:** Serverless scheduler for creating and managing tasks using cron and rate expressions
- **Schema Registry:** Fully managed service for defining, discovering, and sharing event schemas with code generation
- **API Destinations:** Low-code integrations for sending events to external HTTP endpoints via REST API calls
- **Event Replay:** Archive and replay past events for debugging, error recovery, and reprocessing scenarios
- **Global Endpoints:** Improve application availability by replicating events across multiple AWS regions for disaster recovery
- **Content-Based Filtering:** Advanced filtering capabilities with built-in retry mechanisms and exactly-once processing guarantees.

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/eventbridge/

FAQ: https://aws.amazon.com/eventbridge/faqs/

Technical Overview

Backup and Recovery Capabilities

Backup: Amazon EventBridge processes events in real-time without traditional persistence. The Event Archive feature captures and stores events for replay purposes, serving as a backup mechanism for event history. Archive creation is managed through AWS Management Console, CLI, or APIs.

Recovery: Event Replay enables recovery from specific points in time when combined with Event Archive, allowing reprocessing of archived events. Global Endpoints provide automatic failover between regions during outages with high availability within AWS Regions. Supports low RPO through near real-time cross-Region replication and RTO through automatic failover capabilities.

Setup Process

Onboarding: Start immediately through AWS Management Console, CLI, or SDKs without provisioning infrastructure. Create event buses, define rules with event patterns, and specify targets. Schema Registry discovers event structures and generates code bindings.

Offboarding: Delete event buses, rules, targets, and associated resources through AWS Management Console, CLI, or APIs. Remove AWS Identity and Access Management (AWS IAM) roles, disconnect API destinations, clean up partner event sources, and delete Schema Registry schemas.

Data Management

- **Storage Options:** Events processed in real-time without persistence; Event Archive feature stores events for replay scenarios
- **Data Removal:** Events not persisted by default; delete archived events, AWS IAM policies, and Schema Registry schemas via console/APIs
- **Cross-Region Replication:** Global Endpoints automatically replicate events across multiple AWS Regions for disaster recovery and high availability
- **Encryption:** Standard AWS encryption capabilities apply; integrates with AWS PrivateLink for secure connectivity and data protection

Optimise Cost and Consumption

Amazon EventBridge optimises costs through precise event filtering that reduces unnecessary processing and efficient event patterns that minimise rule evaluations. Event transformation at the source reduces data transfer costs compared to target-side transformations. The pay-per-event model scales costs with actual usage, avoiding over-provisioning expenses. Using Amazon EventBridge as a central hub reduces networking costs compared to point-to-point integrations.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/eventbridge/latest/userguide/

Getting Started: https://docs.aws.amazon.com/eventbridge/latest/userguide/eb-tutorial-get-started.html

API Reference: https://docs.aws.amazon.com/eventbridge/latest/APIReference/Welcome.html

Service Quotas: https://docs.aws.amazon.com/eventbridge/latest/userguide/eb-quota.html

Amazon FSx Services

General Information

Amazon FSx is a fully managed file storage service offering four specialised file system types optimised for different workloads. It provides high-performance storage with sub-millisecond latencies and AWS integration without hardware provisioning or management overhead.

Main Benefits

Amazon FSx delivers exceptional performance with up to 1000 GB/s throughput and sub-millisecond latencies while achieving 30-80% cost savings through data deduplication. It is designed for 99.99% availability, helps eliminate infrastructure management overhead, and provides integration with existing applications and AWS services.

Key Use Cases

- **High Performance Computing:** FSx for Lustre delivers 1000 GB/s throughput for ML and scientific computing
- **Windows Workloads:** Support for business applications, SQL Server, SharePoint with Active Directory integration
- **Content Management:** High-throughput file access for video processing, media workflows, and content repositories
- **Data Analytics:** Big data analytics, ML training, and data lake architectures with enterprise performance.

Core Features

- **High Performance Storage:** Sub-millisecond latencies, up to 1000 GB/s throughput, hundreds of thousands of IOPS
- **Multi-Protocol Support:** SMB 2.0-3.1.1, NFS 3.0-4.2, and iSCSI protocols with cross-platform compatibility
- **Data Management:** Data deduplication, compression, intelligent tiering, user quotas, snapshots, and automated lifecycle management
- **High Availability:** Multi-Availability Zone (AZ) deployments with 99.99% SLA, automatic failover, and disaster recovery capabilities
- **Security and Compliance:** Encryption at rest and in transit, Active Directory integration, Amazon Virtual Private Cloud (Amazon VPC) isolation
- **AWS Integration:** Native integration with Amazon Simple Storage Service, Amazon Elastic Compute Cloud, Amazon Elastic Container Service, Amazon Elastic Kubernetes Service, AWS Backup, and AWS DataSync services
- **Four File System Types:** Lustre for HPC, Windows File Server, NetApp ONTAP, OpenZFS options
- **Fully Managed Operations:** No hardware provisioning, automatic software updates, patching, and built-in high availability/

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/fsx/windows/

FAQ: https://aws.amazon.com/fsx/windows/faqs/

Technical Overview

Backup and Recovery Capabilities

Backup: Incremental, encrypted backups stored across multiple Availability Zones can be created manually (AWS Management Console, CLI, API) or automated via AWS Backup with daily scheduling and configurable retention (0-90 days). Cross-Region backup copying and centralised management provide comprehensive protection.

Recovery: Point-in-time restore creates new file systems from any backup, NetApp SnapMirror enables real-time replication, and multi-AZ deployments offer automatic failover. Supports customer RPO (seconds with synchronous replication) and RTO (minutes with automatic failover) objectives.

Setup Process

Onboarding: Create file systems through AWS Management Console, CLI, or API after configuring VPC and Active Directory prerequisites. Select file system type, storage capacity, throughput levels, and network configuration. Setup typically takes minutes with immediate access to enterprise features.

Offboarding: Delete file systems via AWS Management Console, CLI, or API with automatic final backup creation. All data is permanently removed and charging stops. Remove associated resources like mount targets and security groups for complete cleanup.

Data Management

- **Storage Options:** SSD and HDD storage with data deduplication, compression, intelligent tiering, and automated lifecycle management
- **Data Removal:** Delete via AWS Management Console, CLI, or API with automatic final backups and permanent data removal
- **Cross-Region Replication:** NetApp SnapMirror for real-time replication and AWS DataSync for scheduled cross-Region synchronisation
- **Encryption:** Encryption at rest and in transit with all backups encrypted and stored across Availability Zones.

Optimise Cost and Consumption

Data deduplication automatically reduces storage costs by up to 30-80% through duplicate block elimination. Intelligent tiering moves infrequently accessed data to lower-cost storage tiers without performance impact. User and group quotas prevent storage overuse while flexible throughput provisioning enables right-sizing performance capacity. Single-AZ deployments provide up to 40% cost savings for non-critical workloads, and HDD storage options offer cost-effective capacity for less performance-intensive applications.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/fsx/

Getting Started: https://docs.aws.amazon.com/fsx/latest/WindowsGuide/getting-started.html

API Reference: https://docs.aws.amazon.com/fsx/latest/APIReference/

Service Quotas: https://docs.aws.amazon.com/fsx/latest/WindowsGuide/limits.html

Amazon FSx Services – Amazon FSx for Lustre

Feature of Amazon FSx

General Information

Amazon FSx for Lustre is a fully managed, high-performance Lustre file system service delivering sub-millisecond latencies and terabytes per second throughput. It supports ML, HPC, and media workloads with native Amazon Simple Storage Service (Amazon S3) integration.

Main Benefits

Amazon FSx for Lustre delivers the only fully managed Lustre file system in the cloud, providing up to 70% better price-performance than other cloud-based solutions. It helps eliminate operational overhead while delivering sub-millisecond latencies, terabytes per second throughput, and millions of IOPS.

Key Use Cases

- **ML and AI:** Training large-scale models requiring high-throughput data processing and low-latency storage access
- **High Performance Computing:** Scientific computing, weather forecasting, genomic analysis requiring parallel file system performance
- **Media and Entertainment:** Video processing, rendering, content creation workflows requiring fast file I/O and bandwidth
- **Electronic Design Automation:** Semiconductor design and verification workloads with high metadata operations on small files.

Core Features

- **Multiple Deployment Options:** Scratch file systems for temporary storage and persistent systems with data replication
- **Storage Classes:** SSD for high IOPS, HDD for throughput-intensive workloads, Intelligent-Tiering for automatic cost optimisation
- **Amazon S3 Integration:** Native integration with Amazon S3 buckets for data import/export and automatic synchronisation
- **Scalable Metadata Performance:** User-provisioned metadata IOPS up to 15x higher than automatic allocation for small files
- **High Performance:** Sub-millisecond latencies, terabytes per second throughput, and millions of IOPS for demanding workloads
- **Elastic Fabric Adapter Support:** Enhanced network performance for Message Passing Interface applications and tightly coupled HPC workloads
- **Data Compression:** LZ4 compression support to reduce storage consumption and optimise costs while maintaining performance
- **POSIX Compliance:** Works with existing Linux-based applications without modifications, supporting standard file system operations.

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/fsx/lustre/

FAQ: https://aws.amazon.com/fsx/lustre/faqs/

Technical Overview

Backup and Recovery Capabilities

Backup: Incremental, file-system-consistent backups with high durability can be created manually (AWS Management Console, CLI, API) or automated via AWS Backup with daily scheduling. Configurable retention periods from 0-90 days with cross-Region and cross-account backup copying capabilities.

Recovery: Point-in-time recovery creates new file systems with identical configurations, cross-Region disaster recovery with logically air-gapped vaults, and Amazon S3 integration for additional data durability. Supports customer RPO (automatic daily backups with real-time Amazon S3 synchronisation) and RTO (rapid restoration with lazy data loading) objectives.

Setup Process

Onboarding: Create file systems through AWS Management Console, CLI, or API. Configure deployment type (scratch or persistent), storage class (SSD, HDD, or Intelligent-Tiering), capacity, and throughput. Optional Amazon S3 data repository linking and metadata IOPS provisioning available.

Offboarding: Delete file systems via AWS Management Console, CLI, or API with permanent data removal. Export data to Amazon S3 repositories or backup using AWS Backup beforehand. Unmount all clients before deletion.

Data Management

- **Storage Options:** SSD for high IOPS workloads, HDD for throughput-intensive tasks, Intelligent-Tiering for automatic cost optimisation
- **Data Removal:** Delete file systems via AWS Management Console, CLI, or API with permanent data removal after unmounting clients
- **Cross-Region Replication:** Supported through AWS Backup cross-Region copying and Amazon S3 integration for geographic data distribution
- **Encryption:** AWS Key Management Service integration for encryption key management and data protection at rest and in transit.

Optimise Cost and Consumption

Intelligent-Tiering storage class automatically moves infrequently accessed data to lower-cost tiers based on access patterns. LZ4 data compression reduces storage consumption while maintaining performance. Scratch file systems provide cost-effective temporary storage for short-term processing workloads. User-provisioned metadata IOPS optimisation allows independent scaling based on workload requirements. Incremental backups minimise backup storage expenses through efficient data management.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/fsx/latest/LustreGuide/what-is.html

Getting Started: https://docs.aws.amazon.com/fsx/latest/LustreGuide/getting-started.html

API Reference: https://docs.aws.amazon.com/fsx/latest/APIReference/

Service Quotas: https://docs.aws.amazon.com/fsx/latest/LustreGuide/limits.html

Amazon FSx Services – Amazon FSx for NetApp ONTAP

Feature of Amazon FSx

General Information

Amazon FSx for NetApp ONTAP provides fully managed, high-performance shared file storage using NetApp's ONTAP file system. It supports multi-protocol access through Network File System (NFS), Server Message Block (SMB), Internet Small Computer System Interface (iSCSI), and Non-Volatile Memory Express (NVMe) with automatic data-tiering and enterprise-grade data management features.

Main Benefits

Delivers sub-millisecond latencies with petabyte-scale capacity and tens of GBps throughput. Achieves up to 65% storage savings through compression, deduplication, and automatic tiering. Helps eliminate capital expenditure while providing familiar NetApp ONTAP functionality through fully managed AWS service delivery.

Key Use Cases

- **Enterprise File Shares:** High-performance multi-protocol access for Linux, Windows, and macOS clients
- **Database Workloads:** Microsoft SQL Server, Oracle, SAP applications requiring iSCSI and NVMe-over-TCP storage
- **Hybrid Cloud Deployments:** Integration with on-premises NetApp systems using SnapMirror and FlexCache
- **High-Performance Applications:** Sub-millisecond latency storage for demanding enterprise workloads requiring fast access.

Core Features

- **Multi-Protocol Access:** Supports NFS, SMB, iSCSI, and NVMe-over-TCP protocols for broad application compatibility
- **Elastic Capacity Pool Tiering:** Automatic data tiering between high-performance SSD and cost-optimised capacity pool storage
- **Storage Efficiency:** Block-level compression, deduplication, and compaction achieving up to 65% storage savings
- **High Availability:** Multi-Availability Zone (AZ) and Single-AZ deployment options with automatic failover capabilities for resilience
- **NetApp Integration:** Native support for SnapMirror replication, FlexCache caching, and NetApp management tools
- **Enterprise Security:** Encryption at rest and in transit, Active Directory integration, antivirus scanning
- **Instantaneous Snapshots:** Up to 1,023 snapshots per volume for point-in-time recovery and data protection
- **Sub-Millisecond Latency:** High-performance SSD storage delivering enterprise-grade performance with petabyte-scale capacity.

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://docs.aws.amazon.com/prescriptive-guidance/latest/fsx-ontap-enterprise-deployment.html
FAQ: https://docs.aws.amazon.com/prescriptive-guidance/latest/fsx-ontap-enterprise-deployment/faq.html

Technical Overview

Backup and Recovery Capabilities

Backup: Incremental backups stored across multiple Availability Zones for durability can be created manually (AWS Management Console, CLI, API) or automated via AWS Backup with daily scheduling and 0-90 day retention. Native ONTAP snapshots support up to 1,023 snapshots per volume.

Recovery: Instantaneous snapshots enable point-in-time recovery, SnapMirror provides cross-Region replication, and FlexCache offers distributed caching. Multi-AZ deployments provide automatic failover capabilities. Supports customer RPO (near-zero with snapshots, within 24 hours with backups) and RTO (minutes with automatic failover) objectives.

Setup Process

Onboarding: Quick Create option in Amazon FSx console provisions Multi-AZ or Single-AZ file systems with default settings. Configure throughput capacity, SSD IOPS provisioning, and capacity pool tiering policies during or after creation.

Offboarding: Delete file systems through AWS Management Console, CLI, or API after data backup. Sequential deletion of volumes, storage virtual machines, then file system with automated cleanup procedures for associated resources.

Data Management

- **Storage Options:** High-performance SSD and cost-optimised capacity pool storage with automatic tiering between tiers
- **Data Removal:** Delete through AWS Management Console, CLI, or API with sequential removal of volumes, SVMs, and file systems
- **Cross-Region Replication:** Supported via NetApp SnapMirror technology for automated replication to secondary sites and AWS Regions
- **Encryption:** Comprehensive encryption at rest and in transit with enterprise-grade security and access auditing capabilities.

Optimise Cost and Consumption

Amazon Elastic capacity pool tiering automatically moves infrequently accessed data to lower-cost storage while maintaining fast access. Storage efficiency features including compression, deduplication, and compaction achieve up to 65% storage savings. Thin provisioning allows over-provisioning volumes without consuming actual storage until data is written. Flexible throughput and IOPS provisioning enables right-sizing for specific performance requirements without over-provisioning resources.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/fsx/latest/ONTAPGuide/what-is-fsx-ontap.html
Getting Started: https://docs.aws.amazon.com/fsx/latest/ONTAPGuide/getting-started.html
API Reference: https://docs.aws.amazon.com/fsx/latest/APIReference/
Service Quotas: https://docs.aws.amazon.com/fsx/latest/ONTAPGuide/limits.html

Amazon FSx Services – Amazon FSx for OpenZFS

Feature of Amazon FSx

General Information

Amazon FSx for OpenZFS is a fully managed file storage service that enables migration from on-premises ZFS or Linux-based file servers to AWS without changing application code. It delivers high-performance Network File System (NFS) storage with powerful OpenZFS data management capabilities.

Main Benefits

Amazon FSx for OpenZFS delivers exceptional performance with up to 2 million IOPS and 21 GBps throughput while providing sub-millisecond latencies. Intelligent-Tiering automatically optimises costs up to 85% compared to SSD storage, while helping eliminate operational overhead of managing on-premises file servers.

Key Use Cases

- **Media and Entertainment:** High-performance video editing, rendering, and content distribution with low-latency access
- **Scientific Research:** Genomics, financial analytics, and machine learning workloads with high-throughput data processing
- **Software Development:** Containerised applications on Amazon Elastic Kubernetes Service with shared storage and instant snapshots
- **Database Applications:** High-performance Oracle workloads with customisable record sizes and compression optimisation.

Core Features

- **High Performance:** Up to 2 million IOPS and 21 GBps throughput with sub-millisecond latencies
- **OpenZFS Data Management:** Transparent compression, continuous integrity verification, snapshots, thin provisioning, and copy-on-write capabilities
- **Intelligent-Tiering Storage:** Automatically moves data between access tiers based on patterns with 85% cost savings
- **Multi-Protocol Access:** NFS protocol support for Linux, Windows, macOS clients plus Amazon S3 Access Points
- **Volume Management:** Multiple volumes per file system with independent capacity, compression, and user quotas
- **Data Protection:** Automatic daily backups, point-in-time snapshots, cross-Region replication, and AWS Backup integration
- **Deployment Options:** Multi-Availability Zone (AZ) high availability, Single-AZ high availability, and Single-AZ standard deployments available
- **Security Compliance:** Encryption at rest and transit, Amazon Virtual Private Cloud isolation, AWS Identity and Access Management integration, and ISO/PCI/SOC/HIPAA compliance.

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/fsx/openzfs/

FAQ: https://aws.amazon.com/fsx/openzfs/faqs/

Technical Overview

Backup and Recovery Capabilities

Backup: Incremental, crash-consistent backups with automatic daily backups retained up to 90 days and user-initiated backups retained indefinitely can be created manually (AWS Management Console, CLI, API) or automated via AWS Backup with configurable scheduling. Supports cross-Region backup copying and centralised policy management.

Recovery: Point-in-time snapshots enable file-level recovery, instant clone creation, and cross-Region/cross-account replication. Multi-AZ deployments provide automatic failover within 60 seconds. Supports customer RPO (near-zero with synchronous replication) and RTO (under 60 seconds with multi-AZ) objectives.

Setup Process

Onboarding: Deploy via AWS Management Console Quick create for rapid setup or Standard create for customised configurations. Supports AWS CLI/API and AWS CloudFormation templates for infrastructure-as-code automation. Use AWS DataSync for data migration from on-premises ZFS systems.

Offboarding: Terminate Amazon Elastic Compute Cloud instances, delete file system via AWS Management Console /CLI/API. Explicitly delete all automatic and user-initiated backups as they persist independently. Use AWS DataSync for pre-deletion data migration to alternative storage.

Data Management

- **Storage Options:** Intelligent-Tiering automatically moves data between Frequent, Infrequent, and Archive tiers with 85% cost savings
- **Data Removal:** Delete file system via Console/CLI/API, then explicitly delete all automatic and user-initiated backups separately
- **Cross-Region Replication:** On-demand data replication supports cross-Region and cross-account synchronisation with snapshot-based data transfer
- **Encryption:** Provides encryption at rest and in transit with AWS Key Management Service integration for key management.

Optimise Cost and Consumption

Intelligent-Tiering automatically optimises storage costs by moving data between Frequent, Infrequent, and Archive tiers based on access patterns, delivering up to 85% savings. Built-in transparent compression reduces storage consumption while thin provisioning helps ensure payment only for actual data stored. Flexible throughput capacity enables right-sizing performance to workload requirements, and instant snapshots with zero-copy clones minimise development environment storage overhead.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/fsx/latest/OpenZFSGuide/what-is-fsx.html

Getting Started: https://docs.aws.amazon.com/fsx/latest/OpenZFSGuide/getting-started.html

API Reference: https://docs.aws.amazon.com/fsx/latest/APIReference/

Service Quotas: https://docs.aws.amazon.com/fsx/latest/OpenZFSGuide/limits.html

Amazon FSx Services – Amazon FSx for Windows File Server

Feature of Amazon FSx

General Information

Amazon FSx for Windows File Server provides fully managed Microsoft Windows file storage with enterprise-grade features and SMB protocol compatibility. It delivers high-performance shared storage with Active Directory integration and advanced file management capabilities for Windows workloads.

Main Benefits

Amazon FSx helps eliminate infrastructure complexity while reducing operational burden by up to 80%. It delivers up to 12 GB/s throughput and 400,000 IOPS with sub-millisecond latencies, plus data deduplication achieving 50-60% storage savings and designed for 99.99% availability through Multi-Availability Zone (AZ) deployments.

Key Use Cases

- **Enterprise Applications:** Shared file storage for Windows applications with Active Directory integration
- **SQL Server High Availability:** Continuously available file shares for Always On Failover Cluster Instances
- **Virtual Desktop Infrastructure:** File storage for Amazon WorkSpaces and Workspaces Applications (Formerly AppStream 2.0) environments
- **Home Directories:** Centralised storage for user profiles and personal folders with quota management.

Core Features

- **Active Directory Integration:** Native integration with AWS Managed Microsoft AD and self-managed Active Directory for authentication
- **Multi-AZ Deployment:** High availability across multiple Availability Zones with automatic failover and standby file servers
- **Data Deduplication:** Automatic duplicate data elimination delivering typical savings of 50-60% for general-purpose file shares
- **Performance Scaling:** Independent scaling with up to 12 GB/s throughput and 400,000 IOPS with sub-millisecond latencies
- **File Server Resource Manager:** Advanced file management including quotas, file screening, classification, and storage reporting capabilities
- **Shadow Copies:** Point-in-time snapshots enabling end-user file restore and version recovery functionality
- **SMB Protocol Support:** Industry-standard Server Message Block protocol compatibility supporting thousands of concurrent compute instances
- **Hybrid Connectivity:** On-premises access via AWS Direct Connect, Site-to-Site VPN, and FSx File Gateway.

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/fsx/windows/

FAQ: https://aws.amazon.com/fsx/windows/faqs/

Technical Overview

Backup and Recovery Capabilities

Backup: Incremental, file-system-consistent backups stored in Amazon Simple Storage Service designed for 99.999999999% durability can be created manually (AWS Management Console, CLI, API) or automated via automatic daily backups with configurable retention up to 90 days. AWS Backup integration provides centralised management with cross-Region copying and lifecycle policies.

Recovery: Point-in-time recovery creates new file systems with lazy-loading for faster access, multi-AZ automatic failover within 30 seconds, and cross-Region disaster recovery via backup copying. Shadow copies enable end-user file restore. Supports customer RPO (near-zero with multi-AZ, minutes with frequent backups) and RTO (under 30 seconds for failover) objectives.

Setup Process

Onboarding: Launch via AWS Management Console, CLI, or API with console wizard guidance. Requires Amazon Virtual Private Cloud configuration and Microsoft Active Directory setup (AWS Managed or self-managed). File system creation takes 30-45 minutes with immediate administrative access.

Offboarding: Delete file system through AWS Management Console, CLI, or API with permanent data removal. Backup important data before deletion and update application dependencies. Cleanup includes removing Domain Name System aliases and Directory Service connections.

Data Management

- **Storage Options:** SSD and HDD storage types with automatic data deduplication delivering 50-60% savings
- **Data Removal:** Permanent file system deletion via AWS Management Console, CLI, or API with optional final backups
- **Cross-Region Replication:** Supported through backup copying and AWS DataSync for geographic disaster recovery scenarios
- **Encryption:** Built-in encryption at rest and in transit with incremental encrypted backups.

Optimise Cost and Consumption

Automatic data deduplication reduces storage costs by 50-60% through duplicate file elimination. Independent throughput capacity scaling enables performance optimisation without over-provisioning storage. HDD storage options provide cost savings for less performance-sensitive workloads. Single-AZ deployments offer approximately 40% savings for non-critical environments. User and folder-level quotas control storage consumption and prevent cost overruns through granular resource management.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/fsx/latest/WindowsGuide/what-is.html

Getting Started: https://docs.aws.amazon.com/fsx/latest/WindowsGuide/getting-started.html

API Reference: https://docs.aws.amazon.com/fsx/latest/APIReference/welcome.html

Service Quotas: https://docs.aws.amazon.com/fsx/latest/WindowsGuide/limits.html

Amazon GuardDuty

General Information

Amazon GuardDuty is an intelligent threat detection service that continuously monitors AWS accounts, workloads, and data for malicious activity using AI, ML, and threat intelligence. It provides fully managed security monitoring across multiple data sources including AWS CloudTrail logs, Amazon Virtual Private Cloud (Amazon VPC) Flow Logs, and container workloads.

Main Benefits

Amazon GuardDuty helps eliminate infrastructure investment and operational overhead with one-click deployment and automatic scaling without performance impact. Advanced AI/ML Extended Threat Detection identifies sophisticated multi-stage attacks that traditional security tools miss, while supporting up to 50,000 member accounts through AWS Organizations with comprehensive threat intelligence integration.

Key Use Cases

- **Account-Level Threat Detection:** Detecting unauthorised access, credential compromise, and anomalous API activities across AWS accounts
- **Container and Kubernetes Security:** Monitoring Amazon Elastic Kubernetes Service (Amazon EKS) clusters for container-based exploits and unauthorised access through audit logs
- **Data Protection and Compliance:** Monitoring Amazon Simple Storage Service (Amazon S3) buckets for data exfiltration attempts, unauthorised policy changes, and malware detection
- **Database Security Monitoring:** Analysing Amazon Aurora and RDS login activity for potential access threats and unauthorised activities

Core Features

- **Foundational Threat Detection:** Continuously monitors Amazon CloudTrail management events, VPC Flow Logs, and DNS query logs using ML
- **Amazon S3 Protection:** Monitors Amazon S3 buckets for potential security risks including data exfiltration attempts and unauthorised policy changes
- **Amazon EKS Protection:** Analyses Kubernetes audit logs from Amazon EKS clusters to detect container-based exploits and unauthorised activities
- **Runtime Monitoring:** Provides deeper visibility into threats by analysing operating system-level events on Amazon Elastic Compute Cloud (Amazon EC2) instances and containers
- **Malware Protection:** Scans Amazon Elastic Block Store volumes, Amazon EC2 instances, Amazon S3 objects, and AWS Backup data for malware without requiring agents
- **Amazon RDS Protection:** Analyses login activity for Amazon Aurora and RDS databases to detect potential access threats
- **AWS Lambda Protection:** Monitors Lambda function network activity through VPC Flow Logs to detect threats in serverless environments
- **Extended Threat Detection:** Uses AI/ML to correlate security signals across data sources and identify multi-stage attack sequences.

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/guardduty/
FAQ: https://aws.amazon.com/guardduty/faqs/

Technical Overview

Backup and Recovery Capabilities

Backup: Amazon GuardDuty does not provide traditional backup capabilities as it is a security monitoring service. However, findings can be exported to Amazon S3 for indefinite retention beyond the default 90-day period via AWS Management Console or automated through Amazon EventBridge integration for long-term security event archival.

Recovery: Amazon GuardDuty operates regionally and must be enabled in each desired region. Service restoration involves re-enabling Amazon GuardDuty through the AWS Management Console, with immediate monitoring resumption. Enhanced disaster recovery security through AWS Backup integration and continuous threat detection during recovery operations. Does not directly contribute to customer RPO/RTO objectives.

Setup Process

Onboarding: Enable GuardDuty through AWS Management Console with one-click deployment requiring no additional infrastructure. Multi-account environments use AWS Organizations or invitation-based management. Extended Threat Detection activates by default with immediate foundational monitoring.

Offboarding: Suspend Amazon GuardDuty to retain findings or completely disable to permanently delete all configurations. Member accounts must be disassociated before offboarding. Malware Protection for Amazon S3 requires separate management during termination.

Data Management

- **Storage Options:** 90-day finding retention period with export capability to Amazon S3 for indefinite long-term storage
- **Data Removal:** Suspend to retain findings or completely disable to permanently delete all findings and configurations
- **Cross-Region Replication:** Regional service requiring separate enablement in each desired region with no cross-Region replication support
- **Encryption:** Findings encrypted in transit and at rest using AWS managed encryption for security compliance.

Optimise Cost and Consumption

Amazon GuardDuty offers selective protection plan enablement to optimise costs based on specific workload requirements. Volume-based discounts apply to VPC Flow Logs and DNS query log analysis. Trusted IP lists reduce false positives and associated processing costs. The usage monitoring dashboard provides granular cost visibility broken down by individual data sources and protection plans. Malware Protection for Amazon S3 includes incremental scanning that reduces costs by analysing only changed data.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/guardduty/latest/ug/
Getting Started: https://docs.aws.amazon.com/guardduty/latest/ug/guardduty_settingup.html
API Reference: https://docs.aws.amazon.com/guardduty/latest/APIReference/
Service Quotas: https://docs.aws.amazon.com/guardduty/latest/ug/guardduty_limits.html

Amazon Inspector

General Information

Amazon Inspector is an automated vulnerability management service that continuously scans AWS workloads for software vulnerabilities and unintended network exposure. It automatically discovers and assesses Amazon Elastic Compute Cloud (Amazon EC2) instances, container images, AWS Lambda functions, and code repositories with risk-based prioritisation.

Main Benefits

Amazon Inspector helps eliminate operational overhead of traditional vulnerability management by providing fully managed, automated scanning with native AWS integration. It offers contextualised risk scoring considering network accessibility and environment-specific factors, plus hybrid agent-based and agentless scanning for maximum coverage without infrastructure requirements.

Key Use Cases

- **Continuous Vulnerability Scanning:** Automatically assess Amazon EC2 instances, AWS Lambda functions, and container images without manual intervention
- **DevSecOps Integration:** Integrate vulnerability scanning into CI/CD pipelines to identify security issues before production deployment
- **Multi-Account Security Management:** Centrally manage vulnerability assessments across multiple AWS accounts using AWS Organizations capabilities
- **Compliance and Governance:** Meet security compliance requirements by continuously monitoring workloads and generating comprehensive security reports.

Core Features

- **Automated Discovery and Scanning:** Continuously discovers and scans Amazon EC2 instances, AWSLambda functions, Amazon Elastic Container Registry (Amazon ECR) container images, and code repositories without manual configuration
- **Inspector Risk Score:** Provides contextualised risk scoring based on common vulnerabilities and exposures information, network accessibility, and environment-specific factors for prioritisation
- **Multi-Scan Methods:** Supports both agent-based scanning using SSM Agent and agentless scanning using Amazon Elastic Block Store snapshots for maximum coverage flexibility
- **AWS Organisations Integration:** Enables centralised management across multiple accounts with delegated administrator capabilities and organisation-wide policy enforcement
- **Security Hub Integration:** Automatically publishes findings to AWS Security Hub and Amazon EventBridge for centralised monitoring and automated response workflows
- **Code Security Scanning:** Scans application source code, dependencies, and Infrastructure as Code for vulnerabilities using Static application security testing (SAST), Strong customer authentication, and infrastructure as code analysis
- **CIS Benchmark Assessments:** Performs Center for Internet Security benchmark assessments for operating system security configuration compliance
- **Software Bill of Materials (SBOM) Export:** Exports Software Bill of Materials data for compliance reporting and integration with developer tools.

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/inspector/
FAQ: https://aws.amazon.com/inspector/faqs/

Technical Overview

Backup and Recovery Capabilities

Backup: Amazon Inspector does not provide traditional backup capabilities as it generates transient findings data rather than storing critical business data. Users can export findings and SBOM data to Amazon Simple Storage Service (Amazon S3) for retention purposes manually through the AWS Management Console or APIs. The service automatically manages finding lifecycle, closing findings when vulnerabilities are remediated and recreating them during rescans.

Recovery: Inspector does not support traditional disaster recovery mechanisms and requires reactivation in alternate regions during regional failures. Scanning functionality resumes automatically once resources are available without requiring data restoration. Does not directly contribute to RPO and RTO objectives but supports security resilience through continuous vulnerability identification and rapid assessment of disaster recovery environments.

Setup Process

Onboarding: Activate Amazon Inspector with single-click through AWS Management Console or APIs, automatically enabling all scan types. Creates service-linked roles automatically and begins resource discovery immediately. Multi-account environments use AWS Organizations with delegated administrator capabilities for automated deployment.

Offboarding: Deactivating removes all scan settings, suppression rules, and findings data from your account. Export findings to Amazon S3 before deactivation to retain reports. Deactivation through AWS Management Console or API results in immediate cessation of scanning and billing.

Data Management

- **Storage Options:** Regional findings storage with automatic lifecycle management; export findings and SBOM data to Amazon S3 for retention
- **Data Removal:** Deactivating removes all scan settings, suppression rules, and findings data; automatically deletes findings after three days
- **Cross-Region Replication:** Not supported; findings are region-specific and require reactivation in alternate AWS Regions during regional failures
- **Encryption:** Use AWS infrastructure encryption standards; findings data encrypted at rest and in transit through AWS services.

Optimise Cost and Consumption

Amazon Inspector optimises costs through resource exclusion tagging to prevent scanning non-critical instances. Hybrid scanning mode automatically selects the most cost-effective method based on instance characteristics. Suppression rules filter acceptable findings, reducing scanning overhead. Selective scan type activation enables only required capabilities, while automatic finding lifecycle management closes remediated vulnerabilities to minimise storage costs.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/inspector/latest/user/index.html
Getting Started: https://docs.aws.amazon.com/inspector/latest/user/getting_started_tutorial.html
API Reference: https://docs.aws.amazon.com/inspector/v2/APIReference/Welcome.html
Service Quotas: https://docs.aws.amazon.com/inspector/latest/user/quotas.html

Amazon Interactive Video Service (Amazon IVS)

General Information

Amazon Interactive Video Service (Amazon IVS) is a managed live-video streaming service that enables developers to create interactive video experiences and distribute content at scale. It offers low-latency streaming under five seconds and real-time streaming under 300ms with comprehensive multi-platform SDK support.

Main Benefits

Amazon IVS delivers industry-leading performance with ultra-low latency under 300ms for real-time streaming and sub-5-second latency for standard streaming, powered by proven Twitch technology. The service reduces transcoding costs by up to 75% through multitrack video capabilities while supporting millions of viewers with automatic scaling.

Key Use Cases

- **Interactive Live Streaming:** Creating engaging experiences with real-time chat, polls, and Q&A alongside video streaming
- **Real-Time Collaboration:** Supporting collaborative live streams, host-vs-host events, and live video auctions with multiple participants
- **Community Building:** Building social communities with chat rooms, monetising content, and hosting collaborative audio experiences
- **Gaming and Entertainment:** Supporting live gaming streams, educational content delivery, and enterprise communications with interactive features

Core Features

- **Low-Latency Streaming:** Delivers video with latency under five seconds supporting millions of viewers with automatic scaling
- **Real-Time Streaming:** Provides ultra-low latency under 300ms for up to 10,000 viewers and 12 hosts
- **Multi-Platform SDKs:** Broadcast and Player SDKs for iOS, Android, and web with RTMP, RTMPS, SRT support
- **Interactive Chat:** Built-in chat APIs with message delivery, moderation capabilities, and EventBridge integration for real-time interactions
- **Adaptive Bitrate Streaming:** Pre-configured transcoding with multiple quality levels and automatic bitrate adjustment based on viewer conditions
- **Recording and VOD:** Automatic recording to Amazon Simple Storage Service (Amazon S3) with individual participant and composite recording options available
- **Server-Side Composition:** Offloads composition and broadcasting of stages to IVS-managed service with customisable layout options
- **Multitrack Video:** Enables encoding multiple video quality levels from devices, reducing server-side transcoding costs by 75%.

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/ivs/

FAQ: https://aws.amazon.com/ivs/faqs/

Technical Overview

Backup and Recovery Capabilities

Backup: Recorded live streams are automatically stored in Amazon S3 with Cross-Region Replication, versioning, and lifecycle policies. Channel configurations can be backed up manually through infrastructure-as-code templates and API exports or automated via AWS APIs with configurable scheduling.

Recovery: Service provides rapid recovery through global infrastructure with automatic failover capabilities, built-in high availability, and Regional redundancy for content delivery. Supports customer RPO objectives through Amazon S3 recording with configurable frequency and RTO objectives through immediate availability and global edge locations.

Setup Process

Onboarding: Setup involves creating channels through AWS Management Console, CLI, or SDKs with AWS Identity and Access Management permissions. Optional configuration includes Amazon S3 recording, chat functionality, and service quotas. Complete setup process takes minutes with immediate streaming capability.

Offboarding: Termination requires deleting channels, stream keys, and recording configurations through AWS Management Console or APIs. Amazon S3 recorded content follows standard lifecycle policies. Automated cleanup processes available using AWS APIs.

Data Management

- **Storage Options:** Automatic recording to Amazon S3 with individual participant and composite recording options available
- **Data Removal:** Delete channels, stream keys, and configurations through AWS Management Console or APIs with automated cleanup processes
- **Cross-Region Replication:** Supported through Amazon S3 Cross-Region Replication, versioning, and lifecycle policies for recorded content backup
- **Encryption:** Utilises standard AWS encryption capabilities through Amazon S3 integration for stored recorded video content.

Optimise Cost and Consumption

Multitrack Video capability enables client-side encoding of multiple quality levels, reducing transcoding costs by up to 75%. Channel type selection optimises transcoding requirements through Basic, Standard, and Advanced configurations. Amazon S3 lifecycle policies and intelligent tiering reduce recording storage costs. Geographic Region selection leverages tiered pricing structures for output optimisation. Efficient stage management practices control participant hours and real-time streaming consumption.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/ivs/latest/LowLatencyUserGuide/
Getting Started: https://docs.aws.amazon.com/ivs/latest/LowLatencyUserGuide/getting-started.html
API Reference: https://docs.aws.amazon.com/ivs/latest/LowLatencyAPIReference/Welcome.html
Service Quotas: https://docs.aws.amazon.com/general/latest/gr/ivs.html

Amazon Keyspaces (for Apache Cassandra)

General Information

Amazon Keyspaces is a fully managed, serverless Apache Cassandra-compatible database service that helps eliminate infrastructure management while maintaining full API compatibility. It automatically scales tables based on application traffic and supports applications requiring single-digit-millisecond latency with virtually unlimited throughput.

Main Benefits

Amazon Keyspaces is designed to deliver 99.999% availability SLA with serverless architecture that helps eliminate infrastructure management overhead. It provides single-digit-millisecond latency performance, automatic scaling for thousands of requests per second, and cost reductions up to 75% through flexible capacity modes and pay-per-use billing.

Key Use Cases

- **Low Latency Applications:** Single-digit-millisecond performance for industrial equipment maintenance, trade monitoring, and fleet management
- **Open-Source Development:** Build applications using Cassandra APIs and drivers across Java, Python, .NET, and Go
- **Cassandra Migration:** Move existing Cassandra workloads to AWS without managing infrastructure or servers
- **High-Throughput Workloads:** Support thousands of requests per second with virtually unlimited throughput and storage.

Core Features

- **Point-in-Time Recovery:** Continuous backups with restoration to any second within the last 35 days
- **Multi-Region Replication:** Replicate data across multiple AWS Regions for improved availability and resiliency
- **Change Data Capture Streams:** Record row-level change events from tables in near-real time for processing
- **Serverless Architecture:** Automatic scaling without infrastructure management, paying only for resources used
- **Flexible Capacity Modes:** On-demand pay-per-request billing or provisioned capacity modes for predictable workloads
- **Time to Live:** Automatically expire data from tables based on configured values to reduce storage
- **User-Defined Types:** Define custom data structures representing real-world data hierarchies for complex applications
- **Enterprise Security:** Built-in encryption at rest and in transit with AWS-managed or customer-managed keys.

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/keyspaces/

FAQ: https://aws.amazon.com/keyspaces/faqs/

Technical Overview

Backup and Recovery Capabilities

Backup: Point-in-Time Recovery (PITR) creates continuous backups automatically with durable storage across multiple Availability Zones. PITR can be enabled per table via AWS Management Console, AWS CLI, or CQL editor with automated backup retention and cleanup management.

Recovery: Restoration to any second within the last 35 days, multi-Region replication for geographic failover, and automatic failover across Availability Zones. Built-in resilience with 99.999% availability SLA. Supports customer RPO (near-zero through continuous backups) and RTO (rapid scaling without infrastructure provisioning delays) objectives.

Setup Process

Onboarding: Create keyspaces and tables using AWS Management Console, AWS CLI, CQL editor, or AWS SDKs without deploying infrastructure. Configure AWS Identity and Access Management permissions and optional cqlsh-expansion for command-line access. AWS CloudShell recommended for initial setup and testing.

Offboarding: Delete tables and keyspaces using AWS Management Console, AWS CLI, or CQL commands. Cancel Change Data Capture streams, disable auto-scaling policies, and remove Amazon CloudWatch alarms with automatic cleanup of backups and associated resources.

Data Management

- **Storage Options:** Serverless architecture with virtually unlimited throughput and storage capacity, automatic scaling based on traffic
- **Data Removal:** Delete tables and keyspaces via AWS Management Console, AWS CLI, or CQL commands with automatic cleanup
- **Cross-Region Replication:** Multi-Region replication supported across AWS Regions for improved availability and disaster recovery capabilities
- **Encryption:** Built-in encryption at rest and in transit using AWS-managed or customer-managed encryption keys.

Optimise Cost and Consumption

Amazon Keyspaces provides dual capacity modes with on-demand pay-per-request billing for variable workloads and provisioned mode for predictable usage patterns. Application Auto Scaling automatically adjusts provisioned capacity based on usage targets to prevent over-provisioning. Time to Live automatically expires data from tables based on configured values, reducing storage costs. Right-sizing tools identify unused resources through performance metrics optimisation for additional savings opportunities.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/keyspaces/latest/devguide/what-is-keyspaces.html

Getting Started: https://docs.aws.amazon.com/keyspaces/latest/devguide/getting- started.html

API Reference: https://docs.aws.amazon.com/keyspaces/latest/APIReference/

Service Quotas: https://docs.aws.amazon.com/keyspaces/latest/devguide/quotas.html

Amazon Kinesis Data Streams

Feature of Amazon Kinesis

General Information

Amazon Kinesis Data Streams is a fully managed, serverless streaming service that enables real-time collection and processing of large volumes of data. It captures gigabytes per second from sources like clickstreams, transactions, and logs with sub-second latency for immediate analytics and response.

Main Benefits

Designed to deliver enterprise-grade reliability with 99.95% availability and 99.999% durability through automatic multi- Availability Zone (AZ) replication. Provides sub-second latency processing with instant scaling up to 10 GiB/s, while On-demand Advantage mode offers up to 60% cost savings for high-throughput workloads.

Key Use Cases

- **Real-time Analytics:** Processing streaming data for live dashboards, anomaly detection, and dynamic pricing
- **Log Processing:** Accelerated intake and processing of IT infrastructure, application, and clickstream data
- **Metrics and Reporting:** Continuous monitoring and reporting of business metrics and KPIs in real-time
- **Stream Processing:** Building sophisticated data processing pipelines for machine learning and advanced analytics.

Core Features

- **Serverless Architecture:** Fully managed service helps eliminate infrastructure management overhead with automatic scaling capabilities
- **Multiple Capacity Modes:** Choose between On-demand Standard, On-demand Advantage, and provisioned modes for optimal throughput
- **Enhanced Fan-out:** Dedicated throughput of 2 MB/sec per consumer enables parallel data processing without degradation
- **Server-side Encryption:** Automatic encryption at rest using AWS Key Management Service (AWS KMS) keys for comprehensive data protection
- **High Availability:** Data synchronously replicated across three Availability Zones helping ensure 99.95% availability and 99.999% durability
- **Warm Throughput:** Pre-configured throughput capacity up to 10 GiB/s for instant scaling in On-demand Advantage mode
- **Extended Data Retention:** Configurable retention periods from 24 hours up to 365 days for historical data access
- **Large Record Support:** Support for records up to 10 MiB in size accommodating diverse data requirements.

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/kinesis/data-streams/

FAQ: https://aws.amazon.com/kinesis/data-streams/faqs/

Technical Overview

Backup and Recovery Capabilities

Backup: Data protection through configurable retention periods (24 hours to 365 days) with 99.999% durability via automatic synchronous replication across three Availability Zones. Extended retention enables historical data access and replay scenarios for up to one year without traditional backup mechanisms.

Recovery: Point-in-time recovery through configurable retention periods, automatic failover mechanisms during Availability Zone failures, and cross-Region disaster recovery via Amazon Kinesis Data Firehose replication to Amazon Simple Storage Service (Amazon S3). Multi-AZ architecture helps ensure service continuity. Supports aggressive RPO (near-zero via synchronous replication) and RTO (typically minutes for AZ failures) objectives.

Setup Process

Onboarding: Create streams through AWS Management Console, CLI, or SDK by specifying name and selecting capacity modes (on-demand or provisioned). Configure producers using Amazon Kinesis API, KPL, or Amazon Kinesis Agent, and consumers via KCL, AWS Lambda, or direct SDK calls.

Offboarding: Delete streams using DeleteStream API, AWS Management Console, or CLI after shutting down applications and ensuring ACTIVE state. Deletion permanently removes streams, shards, data, tags, and consumer registrations with automatic cleanup.

Data Management

- **Storage Options:** Configurable retention periods from 24 hours up to 365 days with automatic multi-AZ replication
- **Data Removal:** Delete streams via DeleteStream API, AWS Management Console, or CLI after shutting down applications in ACTIVE state
- **Cross-Region Replication:** Implement via producer applications writing to multiple AWS Regions or Amazon Kinesis Data Firehose to Amazon S3.
- **Encryption:** Automatic server-side encryption at rest using AWS KMS keys for comprehensive data protection.

Optimise Cost and Consumption

On-demand Advantage mode delivers up to 60% cost savings through committed usage pricing for high-throughput workloads. Warm throughput capability provides pre-configured capacity up to 10 GiB/s, preventing over-provisioning for predictable traffic spikes. Data compression before ingestion reduces throughput charges, while optimised retention periods lower storage costs. Efficient partition key strategies prevent hot spotting and reduce resharding expenses.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/streams/latest/dev/introduction.html

Getting Started: https://docs.aws.amazon.com/streams/latest/dev/examples.html

API Reference: https://docs.aws.amazon.com/kinesis/latest/APIReference/Welcome.html

Service Quotas: https://docs.aws.amazon.com/streams/latest/dev/service-sises-and-limits.html

Amazon Kinesis Video Streams

Feature of Amazon Kinesis

General Information

Amazon Kinesis Video Streams is a fully managed service that enables secure streaming of live video and time-encoded data from connected devices to AWS Cloud. It automatically provisions infrastructure for ingesting video from millions of devices, storing and indexing data for analytics, ML, and playback applications.

Main Benefits

Amazon Kinesis Video Streams helps reduce video streaming infrastructure complexity with fully managed, serverless solution that automatically scales from one to millions of video sources. Features built-in Amazon Rekognition integration for AI-powered analytics, tiered storage providing up to 50% cost savings, and comprehensive security with end-to-end encryption capabilities.

Key Use Cases

- **Smart Home Applications:** Security cameras and IoT devices for real-time monitoring and motion detection
- **Smart City Solutions:** Traffic cameras and public safety systems for urban infrastructure analytics
- **Industrial Automation:** Manufacturing equipment and drone video for quality control and predictive maintenance
- **Video Analytics:** Machine learning integration with Amazon Rekognition for face detection and content analysis.

Core Features

- **Secure Video Ingestion:** SDKs and APIs with encryption in transit using TLS and at rest using AWS Key Management Service (AWS KMS)
- **WebRTC Support:** Low-latency two-way media streaming between web browsers, mobile applications, and connected devices
- **Video Playback:** Live and on-demand video playback via HTTP Live Streaming (HLS) and MPEG-DASH protocols
- **Amazon Rekognition Integration:** Built-in integration for face detection, object recognition, and automated content analysis
- **Tiered Storage:** Hot tier for immediate access and warm tier for cost-effective long-term retention
- **Automatic Indexing:** Time-based indexing of stored media using producer timestamps and ingestion timestamps
- **Edge Recording:** Local recording capabilities with scheduled cloud streaming for IP cameras on premises
- **Automatic Scaling:** Provisions and scales infrastructure for ingesting streaming video data from millions of devices.

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/kinesis/video-streams/

FAQ: https://aws.amazon.com/kinesis/video-streams/faqs/

Technical Overview

Backup and Recovery Capabilities

Backup: Provides inherent data durability through underlying Amazon Simple Storage Service (Amazon S3) storage with automatic encryption and indexing. Custom backup solutions can be created manually via APIs or integrated with AWS services for cross-Region replication. Configurable retention periods enable cost-effective long-term data retention strategies.

Recovery: Supports low RTO through highly available architecture across multiple Availability Zones and rapid stream recreation capabilities. Multi-Region deployment enables comprehensive disaster recovery planning. Supports customer RPO objectives through configurable retention periods and durable storage, and RTO objectives through high availability and quick stream restoration.

Setup Process

Onboarding: Create Amazon Kinesis video streams through AWS Management Console, CLI, or SDKs. Configure producer devices using Java, C++, or Android SDKs. Setup includes AWS Identity and Access Management permissions, stream properties like retention periods, and storage tier selection.

Offboarding: Delete streams via AWS Management Console, CLI, or DeleteStream API permanently removing all video data. Service automatically handles cleanup of indexed metadata and associated resources upon deletion.

Data Management

- **Storage Options:** Hot tier for immediate access and warm tier for cost-effective long-term retention
- **Data Removal:** Delete streams via AWS Management Console, CLI, or DeleteStream API permanently removing all video data
- **Cross-Region Replication:** Custom solutions using APIs to copy video data to different Regions for disaster recovery
- **Encryption:** TLS encryption in transit and AWS KMS encryption at rest with automatic indexing.

Optimise Cost and Consumption

Tiered storage optimisation allows selecting hot tier for immediate access or warm tier for long-term retention with fragment-based billing. Adjusting fragment duration and configurable retention periods optimise storage costs based on access patterns. Automatic scaling prevents over-provisioning while edge recording capabilities enable scheduled cloud streaming to reduce bandwidth consumption and operational costs.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/kinesisvideostreams/latest/dg/

Getting Started: https://docs.aws.amazon.com/kinesisvideostreams/latest/dg/getting-started.html

API Reference: https://docs.aws.amazon.com/kinesisvideostreams/latest/dg/API_Reference.html

Service Quotas: https://docs.aws.amazon.com/kinesisvideostreams/latest/dg/limits.html

Amazon Lex

General Information

Amazon Lex is a fully managed AI service that enables developers to build conversational interfaces using voice and text. Powered by Alexa technology, it provides speech recognition, natural language understanding, and generative AI capabilities for sophisticated chatbots.

Main Benefits

Amazon Lex democratises AI development by helping reduce deep learning expertise requirements while providing Alexa-powered speech recognition and natural language understanding. The service supports custom vocabularies in 17 languages, offers consumption-based pricing, and includes comprehensive versioning with up to 100 versions per resource type.

Key Use Cases

- **Customer Support:** Build call centre bots and automated agents for handling inquiries and assistance.
- **E-commerce and Retail:** Create transactional bots for order processing, product recommendations, and inventory management.
- **Appointment Booking:** Develop scheduling bots for healthcare, services, and business appointments across industries.
- **IT Helpdesk:** Build automated helpdesk agents for technical support and troubleshooting assistance.

Core Features

- **Speech Recognition:** Advanced automatic speech recognition and natural language understanding capabilities powered by the same technology as Alexa
- **Generative AI Integration:** Uses Large Language Models for enhanced conversational experiences and automated FAQ responses
- **Multi-turn Dialog Management:** Maintains conversation context across multiple interactions and manages complex dialog flows
- **Custom Vocabulary Support:** Supports custom vocabularies in 17 languages to improve speech recognition accuracy
- **Visual Conversation Builder:** Drag-and-drop interface to design and visualise conversation paths and flows
- **Automated Chatbot Designer:** Design chatbots from conversation transcripts using AI-powered analysis and machine learning
- **Multi-language Support:** Supports multiple languages in a single bot with global resiliency features
- **8 kHz Telephony Audio:** Optimised for contact centre and telephony integrations with high-quality audio processing.

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/lex/

FAQ: https://aws.amazon.com/lex/faqs/

Technical Overview

Backup and Recovery Capabilities

Backup: Bot configurations, intents, slots, and dialog flows are automatically backed up through managed service architecture with comprehensive versioning supporting up to 100 versions per resource type. Backups occur automatically without customer intervention via built-in infrastructure management systems.

Recovery: Recovery enables easy rollback through version control systems, Global Resiliency features for near real-time bot replication across AWS Regions, and AWS CloudFormation template deployment. Multi-Region deployment patterns provide fault tolerance. Supports customer RPO objectives through automatic backups and RTO objectives through rapid failover scenarios using AWS Application Recovery Controller.

Setup Process

Onboarding: Access through AWS Management Console with multiple pathways: No-Code Path using pre-built templates, Developer Path for customisation, and Enterprise Path for advanced features. Visual conversation builder and automated chatbot designer enable rapid prototyping within 30-60 minutes.

Offboarding: Delete conversation logs, bot configurations, and associated data through AWS Management Console or API operations. Managed service architecture automatically handles resource cleanup, though customers should manually remove integrated AWS Lambda functions and Amazon DynamoDB tables.

Data Management

- **Storage Options:** Temporary storage for voice and text inputs with comprehensive versioning up to 100 versions per resource type
- **Data Removal:** Delete conversation logs, bot configurations, and associated data through AWS Management Console or API operations with full customer control
- **Cross-Region Replication:** Global Resiliency enables automatic bot replication, version synchronisation, and alias preservation across AWS Regions in near real-time
- **Encryption:** Built-in security and compliance features with enterprise-grade reliability through managed service architecture and AWS ecosystem integration.

Optimise Cost and Consumption

Streaming conversation pricing optimises costs for multi-turn dialogs by processing multiple interactions in single API calls. Pre-built industry templates and visual conversation builder reduce development time and associated costs. The automated chatbot designer enables rapid prototyping from existing conversation transcripts. Efficient conversation flow design minimises turn count, while hybrid text-speech approaches balance functionality with cost-effectiveness for diverse use cases.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/lexv2/latest/dg/what-is.html

Getting Started: https://docs.aws.amazon.com/lexv2/latest/dg/getting-started.html

API Reference: https://docs.aws.amazon.com/lexv2/latest/APIReference/welcome.html

Service Quotas: https://docs.aws.amazon.com/lexv2/latest/dg/quotas.html

Amazon Lightsail

General Information

Amazon Lightsail is the easiest way to get started with AWS for building websites and web applications. It provides pre-configured blueprints and bundles everything needed—instances, databases, storage, networking, and management tools—for predictable monthly pricing.

Main Benefits

Amazon Lightsail delivers simplicity with pre-configured blueprints that launch applications in minutes, helping reduce complex AWS service configuration. Features predictable bundled pricing with included data transfer allowances, automatic daily snapshots with seven-day retention, and AWS environment integration through Amazon Virtual Private Cloud peering for enterprise-grade capabilities.

Key Use Cases

- **Website Hosting:** Deploy WordPress, Drupal, or custom applications with pre-configured templates
- **Development Environments:** Create dev sandboxes and testing environments quickly launched and destroyed
- **Small Business Applications:** Run business software, blogs, e-commerce sites, and collaboration tools
- **Container Deployments:** Deploy and manage containerised applications using Docker with automatic scaling.

Core Features

- **Virtual Private Servers:** Easy-to-set-up instances with pre-configured OS and applications like WordPress, LAMP, Windows Server
- **Container Services:** Run containerised applications with automatic scaling and load balancing capabilities using Docker containers
- **Managed Databases:** MySQL, PostgreSQL, and MariaDB databases with automatic backups and high availability options
- **Load Balancers:** Distribute web traffic across instances with session persistence and health checks for reliability
- **CDN Distributions:** Content delivery network for caching and delivering content to end-users globally with improved performance
- **SSD Block Storage:** Scalable SSD-based block storage and object storage for additional data storage needs
- **DNS Management:** Domain name system management for registered domains with DNS record configuration capabilities
- **Snapshots and Backups:** Manual and automatic snapshots for instances, databases, and block storage for disaster recovery.

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/lightsail/
FAQ: https://aws.amazon.com/lightsail/faq/

Technical Overview

Backup and Recovery Capabilities

Backup: Incremental snapshots for instances, databases, and block storage can be created manually (Amazon Lightsail Console) or automated via daily scheduling with seven-day retention. Manual snapshots persist indefinitely while automatic snapshots retain latest seven copies. Cross-Region snapshot copying enables geographic redundancy.

Recovery: Resource restoration from snapshots ranges from minutes to hours depending on size, snapshots can be exported to Amazon Elastic Compute Cloud for enhanced recovery options, and system disk snapshots enable file access during instance issues. Supports customer RPO (up to 24-hour protection via daily snapshots) and RTO (minutes to hours restoration timeframes) objectives.

Setup Process

Onboarding: Launch via Amazon Lightsail Console by selecting AWS Region, choosing pre-configured blueprints or OS-only options, and selecting instance plans. Configure networking for dual-stack or IPv6-only as needed. Create snapshots before major changes for best practices.

Offboarding: Delete all resources including instances, databases, load balancers, and storage volumes systematically. Create snapshots before deletion to preserve data if needed. Remove Domain Name System (DNS) records and static IP addresses to avoid ongoing charges.

Data Management

- **Storage Options:** SSD-based block storage and scalable object storage for additional data storage needs
- **Data Removal:** Delete all resources including instances, databases, load balancers, Content Delivery Network (CDN) distributions, and storage volumes
- **Cross-Region Replication:** Snapshots can be copied across AWS Regions within same AWS account for geographic redundancy
- **Encryption:** Not explicitly detailed in source material; relies on underlying AWS infrastructure security capabilities.

Optimise Cost and Consumption

Right-size instances using burstable performance that automatically scales based on demand patterns. Use included data transfer allowances within bundled plans to avoid bandwidth overages. Stop instances during non-use periods to reduce costs while preserving configurations. Use automatic snapshot retention that keeps only the latest seven daily backups for efficient storage management. Container services provide built-in auto-scaling to optimise resource use and prevent paying for unused capacity.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/lightsail/latest/userguide/what-is-amazon-lightsail.html
Getting Started: https://docs.aws.amazon.com/lightsail/latest/userguide/getting-started.html
API Reference: https://docs.aws.amazon.com/lightsail/2016-11-28/api-reference/Welcome.html
Service Quotas: https://docs.aws.amazon.com/general/latest/gr/lightsail.html#limits_lightsail

Amazon Location Service

General Information

Amazon Location Service is a fully-managed geospatial platform providing maps, places, routes, trackers, and geofences for location-enabled applications. It delivers high-quality data from providers like Esri, HERE, and GrabMaps with seamless AWS integration and consumption-based pricing.

Main Benefits

Amazon Location Service delivers data quality from leading providers like Esri, HERE, and GrabMaps, designed for 99.95% availability SLA. Customers achieve up to 90% cost reductions while maintaining full data ownership and privacy through anonymisation, AWS integration, and consumption-based pricing model with no upfront commitments.

Key Use Cases

- **Asset Tracking and Fleet Management:** Monitor vehicles, equipment, and personnel in real-time with geofencing capabilities
- **Location-Based Marketing:** Deliver personalised recommendations, proximity-based offers, and location-aware notifications to customers
- **Delivery and Logistics Optimisation:** Calculate optimal routes, estimate travel times, and provide real-time tracking
- **Application Development:** Embed maps, geocoding, and location functionality into mobile and web applications.

Core Features

- **Maps:** Dynamic and static map visualisation with multiple pre-designed styles and high-quality base data
- **Places:** Forward and reverse geocoding, point-of-interest search, address validation, and autocomplete functionality with business information
- **Routes:** Calculates optimal travel routes, estimates travel times, supports multi-point optimisation using real-time traffic data
- **Trackers:** Stores and retrieves current and historical location data for devices with filtering capabilities
- **Geofences:** Detects when tracked devices enter or exit geographical boundaries and triggers automated actions
- **AWS Integration:** Native integration with Amazon EventBridge, AWS CloudFormation, Amazon CloudWatch, AWS Identity and Access Management (AWS IAM), and other core AWS services
- **Data Providers:** High-quality location data from leading providers including Esri, HERE, and GrabMaps
- **Developer Tools:** Comprehensive SDKs for mobile and web applications with sample code and solution guides.

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/location/

FAQ: https://aws.amazon.com/location/faqs/

Technical Overview

Backup and Recovery Capabilities

Backup: Amazon Location Service does not integrate with AWS Backup as it operates as a request-response service for location operations. The service automatically manages tracker position data with built-in 30-day retention policies. Customers must implement backup strategies using Amazon Simple Storage Service or Amazon DynamoDB for application-specific location data needs.

Recovery: Amazon Location Service is designed to deliver automatic failover and high availability through built-in multi-AZ redundancy with 99.95% availability SLA. The fully-managed architecture supports sub-minute recovery objectives for location requests. Supports customer RPO (through caching strategies and retry logic) and RTO (sub-minute recovery with automatic failover) objectives.

Setup Process

Onboarding: Setup begins with AWS account creation and authentication configuration through API Keys, Amazon Cognito, or AWS IAM. Create location resources like maps, place indexes, or trackers via AWS Management Console or APIs with comprehensive SDKs and sample applications.

Offboarding: Delete resources through AWS Management Console, APIs, or CLI with BatchDeleteDevicePositionHistory API for bulk data removal. Account deletion removes all associated resources automatically.

Data Management

- **Storage Options:** Tracker position data with 30-day automatic retention and filtering capabilities to reduce costs
- **Data Removal:** Automatic deletion of tracker data after 30 days; BatchDeleteDevicePositionHistory API for bulk removal
- **Cross-Region Replication:** Multi-Region architectures supported through Amazon Location's global endpoint availability for disaster recovery planning
- **Encryption:** Data privacy and security provided through anonymisation and encryption with AWS Key Management Service integration.

Optimise Cost and Consumption

Amazon Location Service provides tiered pricing buckets (Core, Advanced, Premium) enabling feature-appropriate selections. Tracker position filtering reduces storage costs by eliminating updates based on distance, time, or accuracy thresholds. Map tile caching and static maps optimise Maps API consumption for non-interactive scenarios. Batch APIs like BatchUpdateDevicePosition and BatchGetDevicePosition reduce per-request costs while volume discounts apply for monthly usage exceeding $5,000.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/location/latest/developerguide/index.html

Getting Started: https://docs.aws.amazon.com/location/latest/developerguide/getting-started.html

API Reference: https://docs.aws.amazon.com/location/latest/APIReference/Welcome.html

Service Quotas: https://docs.aws.amazon.com/location/latest/developerguide/manage-quotas.html

Amazon Macie

General Information

Amazon Macie is a data security service that uses ML and pattern matching to discover, monitor, and protect sensitive data in Amazon Simple Storage Service (Amazon S3). It automates discovery of Personally Identifiable Information (PII), financial data, and credentials while evaluating bucket security configurations.

Main Benefits

Amazon Macie delivers unmatched automation and scale for Amazon S3 data security, using ML to reduce manual effort and false positives. It provides comprehensive coverage with over 100 built-in sensitive data identifiers, pay-as-you-analyse pricing, and AWS integration through single API call setup.

Key Use Cases

- **Data Privacy Compliance:** Automatically discover and classify sensitive data for General Data Protection Regulation, California Consumer Privacy Act, and Health Insurance Portability and Accountability Act requirements
- **Security Posture Monitoring:** Continuously evaluate Amazon S3 bucket configurations for security and access control vulnerabilities
- **Risk Assessment and Remediation:** Identify unencrypted or publicly accessible buckets containing sensitive data for remediation
- **Multi-Account Data Governance:** Centrally manage sensitive data discovery across organisational accounts using AWS Organizations.

Core Features

- **Automated Sensitive Data Discovery:** Continuously evaluates Amazon S3 bucket inventories using ML to detect sensitive data
- **Managed Data Identifiers:** Built-in criteria for detecting over 100 types of PII, financial information, and credentials
- **Custom Data Identifiers:** User-defined regex patterns to detect organisation-specific sensitive data types and requirements
- **Policy Findings:** Detects changes to Amazon S3 bucket policies that reduce security or privacy posture
- **Sensitive Data Discovery Jobs:** On-demand and scheduled analysis of specific Amazon S3 buckets with customisable scope
- **Multi-Account Management:** Centralised management across AWS Organisations with delegated administrator capabilities for enterprise governance
- **Allow Lists:** Define text patterns and exceptions that Amazon Macie should ignore during sensitive data analysis
- **Integration Capabilities:** Publishes findings to Amazon EventBridge and AWS Security Hub for automated workflows and remediation.

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/macie/
FAQ: https://aws.amazon.com/macie/faq/

Technical Overview

Backup and Recovery Capabilities

Backup: Amazon Macie does not provide traditional backup capabilities but stores discovery results for 90 days by default in customer-specified Amazon S3 buckets for long-term retention. Findings data can be exported manually through API operations or automatically via Amazon EventBridge integration. Customers must implement backup strategies for configurations and custom data identifiers.

Recovery: Service provides automated data discovery for rapid security posture re-establishment after outages, cross-Region configuration replication using Infrastructure as Code, and multi-Region EventBridge rules for continued monitoring. Uses AWS global infrastructure for fault tolerance. Supports customer RPO objectives through continuous monitoring and RTO objectives via API-driven rapid redeployment.

Setup Process

Onboarding: Enable Amazon Macie with a single AWS Management Console selection or API call requiring minimal setup. Configure AWS Identity and Access Management permissions, select AWS Region, and optionally set up Amazon S3 repository for long-term retention. Multi-account organisations enable AWS Organizations trusted access with delegated administrator capabilities.

Offboarding: Disable service through AWS Management Console or API to delete all Macie-specific settings and resources. Discovery results automatically deleted after 90 days unless stored in customer Amazon S3 repositories. Manually remove Amazon EventBridge rules and AWS Security Hub integrations for complete data cleanup.

Data Management

- **Storage Options:** Discovery results stored 90 days by default, long-term retention in customer-managed Amazon S3 buckets
- **Data Removal:** Automatic deletion after 90 days unless stored in Amazon S3, manual cleanup of repositories required
- **Cross-Region Replication:** Supported through Amazon S3 cross-Region replication for findings data and multi-Region EventBridge rules
- **Encryption:** Integrates with AWS Key Management Service for analysing encrypted Amazon S3 objects with proper key access permissions.

Optimise Cost and Consumption

Amazon Macie's automated discovery samples representative objects rather than analysing entire datasets, significantly reducing analysis costs while maintaining comprehensive coverage. Cost estimation tools during discovery job creation enable expense forecasting before execution. Strategic bucket exclusions focus spending on critical data stores, while allowing lists reduce false positives and minimise manual review efforts. Monthly quotas control costs through pay-as-you-analyse pricing model.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/macie/latest/user/what-is-macie.html
Getting Started: https://docs.aws.amazon.com/macie/latest/user/getting-started.html
API Reference: https://docs.aws.amazon.com/macie/latest/APIReference/welcome.html
Service Quotas: https://docs.aws.amazon.com/macie/latest/user/macie-quotas.html

Amazon Managed Grafana

General Information

Amazon Managed Grafana is a fully managed data visualisation service that enables users to query, correlate, and visualise operational metrics, logs, and traces from multiple sources. It provides logically isolated workspaces for building dashboards without managing underlying infrastructure.

Main Benefits

Amazon Managed Grafana reduces infrastructure management overhead while providing enterprise-grade security and automatic scaling. It offers superior visualisation capabilities compared to Amazon CloudWatch, supports up to 10,000 provisioned users with 500 concurrent users per workspace, and provides pay-per-active-user pricing with a 90-day free trial.

Key Use Cases

- **Cloud Infrastructure Monitoring:** Monitor AWS resource use and performance metrics across multi-account environments
- **Application Performance Management:** Create dashboards for request latency, error rates, and user experience metrics
- **Security Operations:** Display security events in dedicated dashboards for rapid incident response and investigation
- **DevOps Team Collaboration:** Share interactive dashboards across teams with role-based access control and authentication.

Core Features

- **Unified Observability:** Visualise and correlate data across multiple sources with pre-built panels and dashboards
- **Multi-Account Access:** Secure access to AWS data across multiple accounts and Regions using AWS Organizations integration
- **Grafana Alerting:** Robust alerting system with centralised information and multiple alert instances from single rules
- **Team Collaboration:** Easy dashboard sharing with user authentication, authorisation, and team-based access control mechanisms
- **Plugin Management:** Install, update, and manage Grafana community and Enterprise plugins for extended data source support
- **VPC Integration:** Connect to private data sources in Virtual Private Clouds without exposing traffic to the internet
- **Enterprise Security:** Built-in security with SAML 2.0 and AWS Identity and Access Management (AWS IAM) Identity Center integration, encryption everywhere
- **Logically Isolated Workspaces:** Create separate Grafana servers for building dashboards without managing underlying infrastructure components.

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/grafana/

FAQ: https://aws.amazon.com/grafana/faqs/

Technical Overview

Backup and Recovery Capabilities

Backup: Dashboard snapshots with point-in-time backups can be created manually (Snapshot API) or automated via Grafana APIs and Terraform with custom scheduling. Snapshots store locally within workspaces or externally on external servers with configurable expiration times.

Recovery: Instant workspace provisioning and rapid dashboard restoration from snapshots, built-in high availability with automatic scaling and recovery mechanisms. Infrastructure-level resilience through managed durability. Supports customer RPO (manual backup strategies using snapshot APIs) and RTO (low objectives through automatic failover) objectives.

Setup Process

Onboarding: Create workspaces through Amazon Managed Grafana Console selecting authentication methods (AWS IAM Identity Center or SAML) and service-managed permissions. Service automatically provisions AWS IAM roles and policies with AWSGrafanaAccountAdministrator policy required. Getting started tutorials and pre-built dashboard templates accelerate configuration.

Offboarding: Navigate to AWS Management Console, select workspace, and choose Delete with workspace name confirmation. All dashboards, configurations, alerts, and snapshots permanently removed. API endpoints enable programmatic deletion with separate AWS IAM role cleanup required.

Data Management

- **Storage Options:** Dashboards and configurations stored securely in AWS managed environment with durable infrastructure
- **Data Removal:** Workspace deletion permanently removes all dashboards, configurations, alerts, and snapshots with confirmation required
- **Cross-Region Replication:** Manual workspace recreation and configuration restoration required in alternate AWS Regions for cross-Region recovery
- **Encryption:** Built-in encryption at rest and in transit with enterprise-grade security features integrated.

Optimise Cost and Consumption

Pay-per-active-user pricing charges only for users accessing workspaces within 24 hours, not provisioned users. Viewer licenses optimise costs for read-only access requirements. Free API requests reduce automated workflow expenses. Multiple workspaces enable cost segregation between development, testing, and production environments. Automatic scaling adjusts resources based on demand without manual intervention.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/grafana/latest/userguide/what-is-Amazon-Managed-Service-Grafana.html

Getting Started: https://docs.aws.amazon.com/grafana/latest/userguide/getting-started-with-AMG.html

API Reference: https://docs.aws.amazon.com/grafana/latest/APIReference/Welcome.html

Service Quotas: https://docs.aws.amazon.com/grafana/latest/userguide/AMG_quotas.html

Amazon Managed Service for Apache Flink

General Information

Amazon Managed Service for Apache Flink is a fully managed service for processing and analysing streaming data using Apache Flink. It handles infrastructure management, automatic scaling, and provides fault tolerance for real-time analytics and event-driven applications.

Main Benefits

Amazon Managed Service for Apache helps reduce infrastructure management overhead with serverless architecture and automatic scaling. Features per-second billing, over 40 pre-built connectors for AWS integration, and automatic checkpointing every 60 seconds. Delivers multi-Availability Zone (AZ) high availability with rapid recovery, reducing operational complexity while maintaining enterprise-grade performance.

Key Use Cases

- **Real-time Analytics:** Process streaming data for immediate insights and time-series analytics
- **Event-driven Applications:** Build responsive applications that react to streaming events with low latency
- **Streaming ETL:** Transform and load streaming data into data warehouses and lakes real-time
- **IoT Data Processing:** Process high-volume sensor data and telemetry from connected devices.

Core Features

- **Multiple Apache Flink APIs:** Support for SQL, Table API, DataStream API, and Process Functions with Java, Scala, Python
- **Automatic Scaling:** Auto-scales based on CPU usage with configurable parallelism and Kinesis Processing Units (KPUs)
- **Fault Tolerance:** Built-in checkpointing, snapshots, and savepoints for application state management and recovery
- **Studio Notebooks:** Interactive Apache Zeppelin notebooks for ad-hoc data exploration and real-time dashboard creation
- **High Availability:** Multi-AZ deployment with automatic failover and resilience across Availability Zones
- **Per-second Billing:** Pay-per-use pricing model with one-second increments and 10-minute minimum charge
- **Rich Connector Library:** Over 40 pre-built connectors for AWS services and external systems including Amazon MSK, Amazon DynamoDB, Amazon Simple Storage Service (Amazon S3)
- **System Rollbacks:** Automatic rollback capabilities for failed updates and scaling operations.

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/managed-service-apache-flink/

FAQ: https://aws.amazon.com/managed-service-apache-flink/faqs/

Technical Overview

Backup and Recovery Capabilities

Backup: Automatic checkpointing every 60 seconds with encrypted storage in service-managed Amazon S3 buckets can be created manually (AWS Management Console/CLI) or automated via native Apache Flink mechanisms. Supports up to 1,000 snapshots per application for long-term retention with configurable policies.

Recovery: Point-in-time restoration from checkpoints and snapshots typically within minutes, automatic failover across Availability Zones, and system rollback capabilities for failed operations. Multi-AZ deployment provides built-in fault tolerance. Supports customer RPO (configurable 60-second checkpoint intervals) and RTO (rapid restart from checkpoints) objectives.

Setup Process

Onboarding: Deploy via AWS Management Console or CLI after installing Java 11, Maven, and Git. Package applications as JAR files, upload to Amazon S3, configure runtime properties and parallelism settings. Studio notebooks enable direct console development using Apache Zeppelin.

Offboarding: Stop applications using StopApplication API, which creates final snapshots. Manually delete snapshots via DeleteApplicationSnapshot API, remove Amazon S3 application code, and clean up VPC configurations, Amazon CloudWatch logs, and AWS Identity and Access Management roles.

Data Management

- **Storage Options:** Service-managed Amazon S3 buckets for checkpoints and snapshots with configurable retention policies
- **Data Removal:** Manual deletion via DeleteApplicationSnapshot API and separate Amazon S3 bucket cleanup required
- **Cross-Region Replication:** Not supported; disaster recovery uses multi-AZ deployment within single AWS Region
- **Encryption:** All backup data automatically encrypted and stored in service-owned Amazon S3 buckets.

Optimize Cost and Consumption

Automatic scaling adjusts KPUs based on CPU usage thresholds, scaling up at 75% and down at 10% to match workload demand. Right-sizing parallelism settings optimises resource allocation while configurable checkpoint intervals balance performance with storage costs. Studio notebooks provide cost-effective development environments, and minimizing JAR file sizes reduces deployment overhead for variable streaming workloads.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/managed-flink/latest/java/what-is.html

Getting Started: https://docs.aws.amazon.com/managed-flink/latest/java/getting-started.html

API Reference: https://docs.aws.amazon.com/managed-flink/latest/apiv2/Welcome.html

Service Quotas: https://docs.aws.amazon.com/managed-flink/latest/java/limits.html

Amazon Managed Service for Prometheus

General Information

Amazon Managed Service for Prometheus is a serverless, Prometheus-compatible monitoring service for securely monitoring container environments at scale. It automatically scales ingestion, storage, and querying of operational metrics while providing native PromQL compatibility without infrastructure management overhead.

Main Benefits

Amazon Managed Service for Prometheus helps reduce operational overhead of managing Prometheus infrastructure while scaling to 1 billion active series per workspace. Features multi-Availability Zone (AZ) deployments with sub-minute RTO, native PromQL compatibility with 150+ exporters, and integration with existing Prometheus tooling for enterprise-grade container monitoring.

Key Use Cases

- **Container Monitoring:** Monitor Amazon Elastic Kubernetes Service, Amazon Elastic Container Service, AWS Fargate, and self-managed Kubernetes clusters at scale
- **Application Observability:** Collect and analyse time-series metrics from containerised applications using PromQL
- **Multi-Environment Monitoring:** Unified monitoring across AWS, on-premises, hybrid, and multi-cloud environments
- **Alerting and Notifications:** Built-in Alert Manager with Amazon Simple Notification Service (SNS) integration and PagerDuty routing.

Core Features

- **Serverless Monitoring:** Fully managed service with automatic scaling for ingestion, storage, and querying
- **Prometheus Compatibility:** Native support for PromQL query language and over 150 Prometheus exporters
- **AWS Managed Collectors:** Agentless metric collection that automatically discovers and pulls metrics from Kubernetes clusters
- **Alert Manager:** Built-in alerting with deduplication, grouping, routing to SNS topics, and PagerDuty integration
- **High Availability:** Multi-AZ deployments with data replication across three Availability Zones and automatic failover
- **Security Integration:** Encryption in transit and at rest, AWS Identity and Access Management (AWS IAM) access control, AWS Key Management Service (AWS KMS) integration
- **Query Insights:** Query logging to Amazon CloudWatch, cost control with Query Samples Processed (QSP) thresholds, and expensive query monitoring
- **Enterprise Scale:** Handles up to 1 billion active series per workspace with automatic scaling capabilities

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/prometheus/

FAQ: https://aws.amazon.com/prometheus/faqs/

Technical Overview

Backup and Recovery Capabilities

Backup: Amazon Managed Service for Prometheus provides data resilience through automatic replication across multiple Availability Zones within AWS Regions, with data durability supported by underlying AWS infrastructure. Customers can implement backup strategies by replicating metrics to secondary workspaces in different AWS Regions using Prometheus remote write configurations to maintain copies of critical metrics data.

Recovery: Cross-Region workspace replication enables comprehensive disaster recovery, with automatic infrastructure failure handling providing high availability within AWS Regions. Multi-AZ architecture helps ensure automatic failover capabilities. Supports customer RPO objectives through near-zero RPO within AWS Regions via automatic multi-AZ replication, and RTO objectives with sub-minute automatic failover scenarios.

Setup Process

Onboarding: Create workspaces through AWS Management Console, CLI, or SDKs, then configure metric ingestion using AWS managed collectors or custom Prometheus agents. Configure AWS IAM roles, policies, and VPC endpoints for private access. Integration with Amazon Managed Grafana provides immediate visualisation capabilities.

Offboarding: Delete active managed scrapers before workspace deletion to prevent continued charges. Workspace deletion permanently removes all stored metrics and metadata, with automatic data deletion after configurable retention periods (150-1095 days).

Data Management

- **Storage Options:** Serverless time-series storage with automatic scaling, supporting up to 1 billion active series per workspace
- **Data Removal:** Workspace deletion permanently removes all metrics and metadata, with configurable retention periods (150-1095 days)
- **Cross-Region Replication:** Supported through multiple workspaces in different Regions using Prometheus remote write functionality
- **Encryption:** Full encryption in transit and at rest with AWS KMS integration and key management.

Optimise Cost and Consumption

Label-based active series limits control metric volume by producer, while QSP thresholds manage expensive query costs. Query insights and control features enable expensive query monitoring and optimisation. Metric filtering limits active series consumption, and configurable collection frequencies reduce ingestion rates. Automatic scaling matches actual usage, preventing over-provisioning. Tiered pricing provides cost efficiency at scale with usage-based optimisation.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/prometheus/latest/userguide/what-is-Amazon-Managed-Service-Prometheus.html

Getting Started: https://docs.aws.amazon.com/prometheus/latest/userguide/AMP-getting-started.html

API Reference: https://docs.aws.amazon.com/prometheus/latest/APIReference/

Service Quotas: https://docs.aws.amazon.com/prometheus/latest/userguide/AMP_quotas.html

Amazon Managed Streaming for Apache Kafka

General Information

Amazon Managed Streaming for Apache Kafka (MSK) is a fully managed service that helps reduce operational overhead of managing Apache Kafka infrastructure. Amazon MSK enables real-time data streaming and processing with both provisioned and serverless deployment options for scalable streaming applications.

Main Benefits

Amazon MSK reduces management overhead by up to 80% compared to self-managed solutions through automated patching, scaling, and failure recovery. Deep AWS integrations with AWS Lambda, Amazon Simple Storage Service, and analytics services accelerate streaming application development, while flexible provisioned and serverless options optimise costs based on usage patterns.

Key Use Cases

- **Real-time Data Processing:** Capturing and processing log files, clickstreams, and IoT sensor data instantly
- **Event-driven Microservices:** Building decoupled, scalable applications where services communicate through event streams
- **Data Ingestion for Analytics:** Streaming data to lakes and warehouses for real-time analytics and ML
- **Financial Transaction Processing:** Tracking and processing financial transactions with low latency requirements.

Core Features

- **MSK Provisioned:** User-controlled broker scaling with Standard and Express options for optimised performance and cost
- **MSK Serverless:** Fully managed, auto-scaling Kafka clusters without capacity planning or infrastructure management requirements
- **Multi-AZ Deployment:** High availability with automatic failover and data replication across multiple Availability Zones (AZ)
- **MSK Connect:** Fully managed Kafka Connect service for streaming data between Kafka and external systems
- **MSK Replicator:** Cross-Region and cross-cluster data replication for disaster recovery and multi-Region architectures
- **Enterprise Security:** Encryption at rest and in transit, AWS Identity and Access Management (AWS IAM) authentication, Virtual Private Cloud isolation, and granular access control
- **Tiered Storage:** Cost-optimised storage with automatic data tiering to reduce costs while maintaining performance
- **Auto Scaling:** Automatic broker and storage scaling based on demand for provisioned and serverless clusters

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/msk/

FAQ: https://aws.amazon.com/msk/faqs/

Technical Overview

Backup and Recovery Capabilities

Backup: Amazon MSK provides data protection through automatic multi-AZ replication within clusters and configurable topic-level replication factors rather than traditional backups. MSK Replicator enables cross-Region data copying with configurable replication settings. The service focuses on continuous data availability and streaming resilience aligned with Kafka's distributed architecture.

Recovery: Fast cluster recovery typically occurs within minutes through automatic broker replacement and failover mechanisms. Amazon MSK Replicator enables quick failover to secondary AWS Regions, while serverless options provide instant scaling. Supports customer RPO objectives through sub-second synchronous replication and RTO objectives through automated recovery processes.

Setup Process

Onboarding: Setup through AWS Management Console, CLI, or APIs by creating VPC and subnets, then provisioning clusters with provisioned or serverless deployment options. Includes AWS IAM role configuration, client setup, and topic creation with step-by-step tutorials and migration tools.

Offboarding: Delete clusters through AWS Management Console, CLI, or API after stopping producers/consumers and removing topics via Kafka administrative commands. Storage volumes automatically deleted upon termination with Amazon CloudTrail verification available.

Data Management

- **Storage Options:** Tiered storage with automatic data tiering to reduce costs while maintaining performance
- **Data Removal:** Delete topics, consumer groups, and clusters through Kafka APIs or AWS Management Console
- **Cross-Region Replication:** Amazon MSK Replicator provides continuous topic replication to secondary Regions with configurable settings
- **Encryption:** Built-in encryption at rest and in transit with AWS IAM authentication and VPC isolation

Optimise Cost and Consumption

Amazon MSK optimises costs through Express brokers delivering higher throughput at reduced expense and MSK Serverless charging only for actual usage. Tiered storage automatically moves older data to cost-effective tiers while maintaining performance. Auto-scaling capabilities match capacity with demand patterns, while configurable retention policies and optimised replication factors reduce storage overhead. Granular pricing models enable fine-tuned cost control across compute, storage, and data transfer components.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/msk/latest/developerguide/what-is-msk.html

Getting Started: https://docs.aws.amazon.com/msk/latest/developerguide/getting-started.html

API Reference: https://docs.aws.amazon.com/msk/1.0/apireference/

Service Quotas: https://docs.aws.amazon.com/msk/latest/developerguide/limits.html

Amazon Managed Workflows for Apache Airflow

General Information

Amazon Managed Workflows for Apache Airflow (MWAA) is a fully managed orchestration service that simplifies setting up and operating end-to-end data pipelines using Apache Airflow. It automatically scales infrastructure while providing built-in security and native AWS service integration.

Main Benefits

Amazon MWAA help reduce operational overhead of managing Apache Airflow infrastructure while providing enterprise-grade security and automatic scaling. It integrates natively with over 20 AWS services, offers 20-30 minute environment provisioning, and supports both traditional and serverless deployment models for flexible cost optimisation.

Key Use Cases

- **ETL Pipeline Orchestration:** Coordinating extract, transform, and load operations across multiple data sources and destinations
- **Batch Data Processing:** Processing large datasets in scheduled batches with complex dependencies and orchestration requirements
- **Multi-Service Workflow Coordination:** Orchestrating tasks across AWS services like Amazon Simple Storage Service (Amazon S3), Amazon Elastic MapReduce (Amazon EMR), AWS Glue, AWS Lambda, and Amazon SageMaker
- **Machine Learning Workflow Orchestration:** Managing ML training pipelines, model deployment, and inference workflows with dependencies.

Core Features

- **Automatic Airflow Setup:** Quickly deploys Apache Airflow environments with chosen versions using open-source UI and code
- **Automatic Scaling:** Scales workers and webservers based on workload demands with configurable minimum and maximum limits
- **Built-in Security:** Runs in Amazon Virtual Private Cloud (Amazon VPC) with automatic data encryption using AWS Key Management Service (AWS KMS) and AWS Identity and Access Management integration
- **AWS Service Integration:** Native integration with over 20 AWS services including Amazon S3, Amazon EMR, AWS Glue, AWS Lambda, and Amazon SageMaker
- **Streamlined Upgrades:** Provides automatic patches and managed upgrades to new Apache Airflow versions without maintenance windows
- **Workflow Monitoring:** Integrates with Amazon CloudWatch for comprehensive logs, metrics, and monitoring without requiring third-party tools
- **Public/Private Access Modes:** Supports both public internet access and private VPC-only access to Apache Airflow webserver
- **Serverless Option:** MWAA Serverless provides pay-per-use pricing with automatic infrastructure scaling based on task execution.

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/managed-workflows-for-apache-airflow/

FAQ: https://aws.amazon.com/managed-workflows-for-apache-airflow/faqs/

Technical Overview

Backup and Recovery Capabilities

Backup: Amazon MWAA provides automatic backup through its managed metadata database with built-in snapshots and point-in-time recovery. Directed Acyclic Graph (DAG) files are stored in Amazon S3 with 11 9's durability and versioning capabilities. Customers can implement backup strategies using S3 cross-Region replication and Infrastructure as Code templates via AWS CloudFormation or Terraform.

Recovery: Environment recreation takes 20-30 minutes plus Amazon S3 replication lag, supporting backup and restore disaster recovery strategies with RTOs of 2-4 hours. The containerised Fargate architecture provides built-in fault tolerance within AWS Regions. Supports customer RPO objectives (near-zero using Amazon S3 cross-Region replication) and RTO objectives (2-4 hours for most workflow orchestration requirements).

Setup Process

Onboarding: Create environments through AWS Management Console, CLI, or CloudFormation with Amazon S3 bucket and VPC configuration. Specify Apache Airflow version, DAG storage location, and networking settings. Environment provisioning takes 20-30 minutes with quick-start AWS CloudFormation templates available.

Offboarding: Delete environments via AWS Management Console, CLI, or API automatically removing all managed infrastructure. DAG files in Amazon S3 require manual deletion if desired. Amazon CloudWatch logs retained based on configured policies.

Data Management

- **Storage Options:** Amazon S3 with versioning for DAG files and managed metadata database with automatic snapshots
- **Data Removal:** Manual deletion of Amazon S3 DAG files required; managed infrastructure automatically removed via console/CLI/API
- **Cross-Region Replication:** Supported through Amazon S3 cross-Region replication capabilities with 11 9's durability for disaster recovery
- **Encryption:** Automatic data encryption using AWS KMS with VPC isolation and role-based authentication integration.

Optimise Cost and Consumption

MWAA Serverless provides pay-per-use pricing based on actual task execution time, avoiding idle infrastructure costs. Automatic scaling of workers and webservers matches workload demands while configurable minimum/maximum worker counts prevent over-provisioning. Environment class selection from micro to large enables right-sizing for specific workload requirements. Efficient DAG designs reduce task execution time, and scheduled maintenance windows minimise operational disruption.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/mwaa/latest/userguide/what-is-mwaa.html

Getting Started: https://docs.aws.amazon.com/mwaa/latest/userguide/get-started.html

API Reference: https://docs.aws.amazon.com/mwaa/latest/API/Welcome.html

Service Quotas: https://docs.aws.amazon.com/mwaa/latest/userguide/mwaa-quotas.html

Amazon MQ

General Information

Amazon MQ is a managed message broker service for Apache ActiveMQ Classic and RabbitMQ that enables applications to communicate using various programming languages and messaging protocols. The service helps reduce operational overhead by automating setup, maintenance, and security patches while supporting industry-standard APIs.

Main Benefits

Amazon MQ helps reduce operational overhead by automating infrastructure management, patching, and maintenance while enabling zero-code migration of existing ActiveMQ and RabbitMQ applications. The service provides near-zero RPO with automatic failover, supports up to 1,000 connections per broker, and includes inter-node data transfer at no additional cost for RabbitMQ clusters.

Key Use Cases

- **Enterprise Application Modernisation:** Migrate on-premises message brokers to cloud without rewriting messaging code
- **Microservices Decoupling:** Enable loose coupling between application components using reliable message queuing patterns
- **Event-Driven Architectures:** Process real-time events and data streams through message routing workflows
- **Lambda Integration:** Automatically consume messages from brokers for serverless event-driven processing.

Core Features

- **Managed Infrastructure:** Automated provisioning, patching, maintenance, and scaling of message broker infrastructure
- **Multi-Engine Support:** Support for Apache ActiveMQ Classic and RabbitMQ with native web consoles
- **High Availability:** Active/standby deployments for ActiveMQ and cluster deployments across multiple Availability Zones (AZs)
- **Industry Standard Protocols:** Support for JMS, NMS, AMQP, STOMP, MQTT, and WebSocket protocols
- **Security and Encryption:** Virtual Private Cloud (VPC) access, TLS/SSL encryption, message encryption with AWS Key Management Service (AWS KMS) integration
- **Storage Options:** Durability-optimised with Amazon Elastic File System (Amazon EFS) and throughput-optimised with Amazon Elastic Block Store (Amazon EBS) storage
- **Network of Brokers:** ActiveMQ broker mesh topology for horizontal scaling and message routing
- **CloudWatch Integration:** Comprehensive monitoring through metrics, logs, and alarms for broker performance tracking.

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/amazon-mq/

FAQ: https://aws.amazon.com/amazon-mq/faqs/

Technical Overview

Backup and Recovery Capabilities

Backup: Amazon MQ provides data durability through managed storage solutions including Amazon EFS replication across multiple Availability Zones for durability-optimised brokers and Amazon EBS snapshots for throughput-optimised brokers. Custom backup strategies can be implemented using Amazon MQ APIs and storage-level backup mechanisms.

Recovery: Active/standby brokers provide automatic failover achieving RTO of minutes, cross-Region data replication (CRDR) enables regional disaster recovery with minimal data loss, and RabbitMQ cluster deployments offer node-level failover. Supports customer RPO (near-zero with active/standby configurations) and RTO (minutes with automatic failover) objectives.

Setup Process

Onboarding: Access Amazon MQ Console and select Apache ActiveMQ or RabbitMQ engine. Configure deployment mode (single-instance, active/standby, or cluster), broker instance types, storage options, and credentials. Set up VPC settings, security groups, and subnet assignments for network configuration.

Offboarding: Delete brokers through AWS Management Console, API, or CLI with confirmation. Deletion takes approximately five minutes and permanently removes all message data, configurations, and broker metadata irreversibly.

Data Management

- **Storage Options:** Durability-optimised with Amazon EFS and throughput-optimised with Amazon EBS storage
- **Data Removal:** Delete brokers via AWS Management Console, API, or CLI; takes five minutes and permanently removes all data
- **Cross-Region Replication:** CRDR enables real-time data replication between ActiveMQ brokers in different AWS Regions
- **Encryption:** VPC access, TLS/SSL encryption, message encryption at rest and in transit with AWS KMS.

Optimise Cost and Consumption

Amazon MQ enables rightsizing broker instances based on actual message throughput and connection requirements to optimise resource allocation. Choose between durability-optimised Amazon EFS storage for high availability or throughput-optimised Amazon EBS for performance-focused workloads. Configure message retention policies to control storage usage and monitor actual storage consumption. Single-instance brokers provide cost-effective options for development and testing environments.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/amazon-mq/latest/developer-guide/welcome.html

Getting Started: https://docs.aws.amazon.com/amazon-mq/latest/developer-guide/getting-started-activemq.html

API Reference: https://docs.aws.amazon.com/amazon-mq/latest/api-reference/resources.html

Service Quotas: https://docs.aws.amazon.com/amazon-mq/latest/developer-guide/amazon-mq-limits.html

Amazon Neptune

General Information

Amazon Neptune is a fully managed graph database service optimised for storing billions of relationships and querying highly connected datasets with millisecond latency. It supports multiple graph models and query languages including Gremlin, openCypher, and SPARQL for building recommendation engines, fraud detection systems, and knowledge graphs.

Main Benefits

Neptune delivers exceptional performance with millisecond latency queries and up to 100,000 queries per second while handling billions of relationships. The fully managed service helps reduce operational overhead and offers up to 90% cost savings through Neptune Serverless. Advanced disaster recovery provides 1-second RPO and sub-minute RTO.

Key Use Cases

- **Recommendation Engines:** Building personalised content and product recommendation systems using graph relationships
- **Fraud Detection:** Identifying fraudulent patterns and activities through complex relationship analysis in real-time
- **Knowledge Graphs:** Creating structured repositories of information with defined concepts, properties, and relationships
- **Identity Graphs:** Building unified customer views across multiple devices and identifiers for personalisation.

Core Features

- **Multi-model Graph Support:** Supports Property Graph and Resource Description Framework models with Gremlin, openCypher, and SPARQL query languages
- **High Performance:** Optimised for billions of relationships with millisecond latency queries and 100,000 queries per second
- **Neptune Serverless:** Automatically provisions and scales database capacity based on workload requirements with pay-per-use pricing
- **Global Database:** Spans multiple AWS Regions with storage-based replication and sub-second latency for disaster recovery
- **Neptune ML:** Integrates ML using graph neural networks for node classification, edge prediction, and link prediction
- **Multi-AZ Deployment:** Provides high availability with read replicas across Availability Zones (AZs) and automatic failover capabilities
- **Continuous Backup:** Automatic, continuous, incremental backups with point-in-time recovery and manual database snapshots
- **Neptune Analytics:** Memory-optimised analytics engine for analysing large graph datasets with built-in algorithms and GraphRAG capabilities.

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/neptune/

FAQ: https://aws.amazon.com/neptune/faqs/

Technical Overview

Backup and Recovery Capabilities

Backup: Automated continuous incremental backups to Amazon Simple Storage Service (Amazon S3) with 1-35 day retention periods can be created manually (AWS Management Console, CLI, APIs) or automated via AWS Backup with centralised scheduling and policy management. Cross-Region snapshot copying and Neptune streams replication provide additional backup durability.

Recovery: Point-in-time recovery to any second within retention window, Global Database cross-Region disaster recovery, and Multi-AZ automatic failover within minutes. Storage-based replication ensures fault tolerance across Availability Zones. Supports customer RPO (1 second with Global Database) and RTO (sub-minute cross-Region failover) objectives.

Setup Process

Onboarding: Launch clusters through AWS Management Console, CLI, or APIs with Amazon Neptune Serverless for immediate scaling without capacity planning. Configure Virtual Private Cloud (VPC) deployment across multiple Availability Zones with SSL connections and proper AWS Identity and Access Management roles for security.

Offboarding: Delete date base clusters to remove data, automated backups, and resources. Complete cleanup requires deleting manual snapshots, Amazon Neptune Analytics graphs, VPC endpoints, AWS Lambda functions, and cross-Region replicas separately in each Region.

Data Management

- **Storage Options:** Purpose-built graph database engine with automatic scaling and continuous backup to Amazon S3
- **Data Removal:** Delete data base clusters removes all data and automated backups; manual snapshots require separate deletion
- **Cross-Region Replication:** Global Database provides storage-based replication across AWS Regions with sub-second latency and disaster recovery
- **Encryption:** Advanced security with encryption at rest and in transit, integrated with AWS Key Management Service.

Optimised Cost and Consumption

Amazon Neptune Serverless automatically scales capacity based on workload demands with pay-per-use pricing, reducing costs up to 90% for variable workloads. Database Savings Plans provide discounts for predictable workloads with 1-year commitments. AWS Graviton4 R8g instances deliver 4.7 times better price-performance for write operations. Query optimisation through proper graph indexing reduces I/O costs, while read replicas distribute load efficiently.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/neptune/

Getting Started: https://docs.aws.amazon.com/neptune/latest/userguide/graph-get-started.html

API Reference: https://docs.aws.amazon.com/neptune/latest/apiref/

Service Quotas: https://docs.aws.amazon.com/neptune/latest/userguide/limits.html

Amazon OpenSearch Service

General Information

Amazon OpenSearch Service is a fully managed service that simplifies deployment, operation, and scaling of OpenSearch clusters for search and analytics. It provides log analytics, real-time monitoring, and business intelligence capabilities with automated infrastructure management.

Main Benefits

Amazon OpenSearch helps reduce operational overhead with fully managed infrastructure, designed for 99.99% availability SLA, and zero downtime updates. Delivers up to 44% better price/performance with Graviton2 instances, 40% cost savings through UltraWarm storage, and scaling to 3 petabytes with enterprise-grade security compliance.

Key Use Cases

- **Log Analytics:** Process machine-generated time series data from application logs and operational telemetry.
- **Real-time Application Monitoring:** Monitor performance, identify errors, and track system health in real-time.
- **Full-text Search:** Build sophisticated search applications with advanced filtering for e-commerce and content.
- **Security Analytics:** Analyse security logs, detect anomalies, and perform threat investigations with custom rules.

Core Features

- **Serverless Deployment:** On-demand auto-scaling configuration that automatically scales compute resources based on application needs.
- **Multi-tier Storage:** UltraWarm for cost-effective Amazon Simple Storage Service (Amazon S3)-based storage and cold storage for rarely accessed data.
- **Fine-grained Access Control:** Role-based security at index, document, and field level with multi-tenancy support.
- **Vector Search Engine:** GPU-accelerated vector search capabilities for AI applications supporting billion-scale vector databases.
- **Multi-Availability Zone (AZ) with Standby:** Enhanced availability configuration helping ensure resilience to single Availability Zone failures with 99.99% SLA.
- **Index State Management:** Automated lifecycle management with transitions between hot, warm, and cold storage tiers.
- **OpenSearch Dashboards:** Native visualisation tool for creating interactive dashboards and custom visualisations without additional fees.
- **Cross-cluster Operations:** Cross-cluster search and replication capabilities for disaster recovery and reduced latency.

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/opensearch-service/

FAQ: https://aws.amazon.com/opensearch-service/faqs/

Technical Overview

Backup and Recovery Capabilities

Backup: Automated hourly snapshots stored in Amazon S3 with 14-day retention at no additional cost can be created manually (AWS Management Console, CLI, API) or automated via Index State Management policies. Manual snapshots support extended retention and cross-Region storage in customer-managed Availability Zone S3 buckets.

Recovery: Complete domain restoration within minutes to hours depending on data volume, cross-Region snapshot restoration for disaster recovery, and cross-cluster replication for real-time synchronisation. Multi-AZ with Standby provides 99.99% availability SLA. Supports customer RPO (1-hour automated snapshots) and RTO (near-zero with Multi-AZ Standby) objectives.

Setup Process

Onboarding: Create domains through AWS Management Console, CLI, or SDKs specifying version, instance types, and storage options. Configure security settings including fine-grained access control, encryption, and Virtual Private Cloud (VPC) settings. Migration Assistant facilitates transfers from self-managed clusters with automated data optimisation.

Offboarding: Export data using manual snapshots to customer Amazon S3 buckets before deletion. Delete domain via AWS Management Console, CLI, or API permanently removing compute, storage, and automated snapshots. Remove associated AWS Identity and Access Management policies, security groups, and VPC configurations.

Data Management

- **Storage Options:** Hot, UltraWarm Amazon S3-based storage with 40% cost savings, and cold storage for rarely accessed data.
- **Data Removal:** Export data via manual snapshots to Amazon S3 buckets before domain deletion permanently removes all resources.
- **Cross-Region Replication:** Supported for real-time data synchronisation between domains in different AWS Regions for disaster recovery.
- **Encryption:** Built-in encryption at rest and in transit with AWS Key Management Service integration for enterprise-grade data protection.

Optimise Cost and Consumption

Index State Management policies automatically transition data between hot, UltraWarm, and cold storage tiers based on customisable age and access patterns, delivering up to 40% cost savings. AWS Graviton2 instances provide 44% better price/performance ratios. Reserved Instances offer up to 52% discounts for predictable workloads. Serverless deployment provides automatic scaling with pay-per-use pricing for variable workloads.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/opensearch-service/latest/developerguide/what-is.html

Getting Started: https://docs.aws.amazon.com/opensearch-service/latest/developerguide/gsg.html

API Reference: https://docs.aws.amazon.com/opensearch-service/latest/API Reference/

Service Quotas: https://docs.aws.amazon.com/general/latest/gr/opensearch-service.html

Amazon Personalize

General Information

Amazon Personalize is a fully managed ML service that generates personalised item recommendations and user segments based on interaction data. It enables developers to build customised experiences for video streaming, e-commerce, and marketing applications without requiring ML expertise.

Main Benefits

Amazon Personalize reduces ML expertise requirements while supporting catalogues up to five million items with lower inference latency through v2 recipes. It automatically handles model training, hyperparameter optimisation, and infrastructure scaling, reducing operational overhead and accelerating time-to-market for personalised experiences.

Key Use Cases

- **Video Streaming:** Personalising apps with Top picks, More like X, and most popular recommendations
- **E-commerce:** Adding Recommended for you and frequently bought together product recommendations to applications
- **Marketing Campaigns:** Creating targeted emails and campaigns using user segments based on item affinity
- **Search Personalisation:** Re-ranking search results based on individual user preferences and behaviour patterns.

Core Features

- **Real-time Personalisation:** Updates recommendations based on evolving user interests by recording interactions with items
- **Domain Recommenders:** Preconfigured resources optimised for VIDEO_ON_DEMAND and ECOMMERCE domains with use-case specific recipes
- **Custom Solutions:** Configurable ML solutions using customisable recipes for unique business requirements
- **Batch Recommendations:** Bulk recommendation generation for email campaigns and offline processing workflows
- **User Segmentation:** Generate segments of users based on affinity for items or metadata for targeted marketing
- **Automatic Training:** Automatically retrain solutions every seven days using latest data to maintain recommendation relevance
- **Generative AI Integration:** Content Generator adds descriptive themes to batch recommendations using generative AI
- **Rule-based Promotions:** Support for promoting specific items in recommendations based on business rules.

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/personalize/

FAQ: https://aws.amazon.com/personalize/faqs/

Technical Overview

Backup and Recovery Capabilities

Backup: Amazon Personalize does not provide native backup capabilities or AWS Backup integration. Customers must manually maintain copies of training data, schemas, and configuration parameters in external storage like Amazon Simple Storage Service (Amazon S3). Data durability relies on AWS managed infrastructure resilience and multi-Availability Zone (AZ) deployment.

Recovery: Recovery requires recreating dataset groups, reimporting data, and retraining models. Cross-Region recovery involves recreating resources in alternate AWS Regions with Amazon S3 cross-Region replication for training datasets. Multi-AZ deployment provides automatic infrastructure resilience. RPO depends on real-time event streaming frequency and customer backup strategies, while RTO is determined by model retraining duration.

Setup Process

Onboarding: Setup begins through AWS Management Console, CLI, or SDKs by choosing Domain dataset groups for VIDEO_ON_DEMAND/ECOMMERCE or Custom groups for unique cases. Prepare CSV training data, create schemas, upload to Amazon S3, and create dataset import jobs. Build recommenders or solutions then deploy campaigns.

Offboarding: Systematic resource deletion in specific order: campaigns/recommenders first, then solutions, event trackers, filters, datasets, and schemas. Data deletion jobs remove user interactions via CSV files, taking up to 24 hours to complete.

Data Management

- **Storage Options:** Uses Amazon S3 for bulk data import/export with structured CSV format and AWS managed infrastructure
- **Data Removal:** Data deletion jobs remove specific users and interactions using CSV files, taking up to 24 hours
- **Cross-Region Replication:** Customers implement S3 cross-Region replication for training datasets across Regions for disaster recovery
- **Encryption:** Relies on AWS managed infrastructure encryption with AWS Identity and Access Management access control and PrivateLink Virtual Private Cloud endpoints

Optimise Cost and Consumption

Amazon Personalize provides tiered pricing with automatic volume discounts at 72 million and 720 million request thresholds. Batch recommendations can significantly reduce costs for bulk operations compared to individual real-time requests. Customers can disable automatic training to control model retraining expenses while maintaining performance. Domain-specific recommenders minimise development overhead, while optimised data formats and import frequency scheduling reduce data processing costs.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/personalise/latest/dg/
Getting Started: https://docs.aws.amazon.com/personalise/latest/dg/getting-started.html
API Reference: https://docs.aws.amazon.com/personalise/latest/dg/API_Reference.html
Service Quotas: https://docs.aws.amazon.com/personalise/latest/dg/limits.html

Amazon Polly

General Information

Amazon Polly is a cloud-based text-to-speech service that converts text into lifelike speech using advanced ML technology. The service supports over 60 voices across multiple languages with four voice engines optimised for different quality requirements and use cases.

Main Benefits

Amazon Polly delivers high-quality, natural-sounding speech with fast response times suitable for real-time applications. The consumption-based pricing model helps reduce upfront costs, while caching and replay capabilities provide unlimited audio reuse at no additional charges, reducing costs by up to 75%.

Key Use Cases

- **Mobile Applications:** Adding voice capabilities for enhanced user experiences and accessibility features
- **eLearning Platforms:** Converting educational content to speech for auditory learners and compliance
- **Accessibility Applications:** Creating inclusive experiences through document-to-speech conversion for visually impaired users
- **Contact Centres:** Integrating with Amazon Connect for interactive voice response systems.

Core Features

- **Multiple Voice Engines:** Standard, Neural, Long-form, and Generative voices with varying quality levels and pricing
- **SSML Support:** Speech Synthesis Markup Language for controlling speech rate, pitch, emphasis, and pronunciation
- **Speech Marks:** Metadata providing timing information for synchronising speech with visual elements
- **Custom Lexicons:** Ability to modify pronunciation of specific words and acronyms for accuracy
- **Brand Voice:** Custom Neural Text-to-Speech voices designed specifically for organisations
- **Multiple Audio Formats:** Support for MP3, OGG Vorbis, and raw PCM output formats
- **Newscaster Speaking Style:** Professional news-style delivery for U.S. English, British English, and U.S. Spanish
- **Time-driven Prosody:** Ability to define maximum speech duration for localisation and dubbing applications.

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/polly/

FAQ: https://aws.amazon.com/polly/faqs/

Technical Overview

Backup and Recovery Capabilities

Backup: Amazon Polly operates as a stateless service without built-in backup capabilities, as it does not retain customer text inputs or generated audio files. Customers must implement their own backup strategies for generated audio content using Amazon Simple Storage Service (Amazon S3) with versioning and cross-Region replication.

Recovery: Recovery is supported through multi-Region availability and stateless architecture, enabling traffic redirection to alternative Regions during outages. The cloud-native design provides inherent resilience with automatic geographic distribution. Supports customer RPO of essentially zero (no data retention) and RTO of minutes through automated failover using Amazon Route 53 health checks.

Setup Process

Onboarding: Getting started requires an AWS account with automatic service access via AWS Management Console, CLI, or SDKs in multiple programming languages. Configure AWS Identity and Access Management permissions and select appropriate voice engines based on quality requirements.

Offboarding: Simply stop using the service as no long-term commitments exist. Amazon Polly automatically removes text inputs upon processing completion with no stored data requiring manual deletion.

Data Management

- **Storage Options:** Amazon Polly operates as a stateless service without persistent storage for text inputs or audio files
- **Data Removal:** Text inputs are automatically removed upon processing completion with no data retention by the service
- **Cross-Region Replication:** Customers must implement their own backup strategies using Amazon S3 with versioning and cross-Region replication
- **Encryption:** Service maintains compliance certifications for Health Insurance Portability and Accountability Act and Payment Card Industry Data Security Standard regulated workloads with secure processing capabilities.

Optimise Cost and Consumption

Amazon Polly offers audio caching capabilities that allow unlimited replay of generated speech without additional synthesis charges, significantly reducing costs for frequently accessed content. Selecting appropriate voice engines based on specific quality requirements optimises consumption, while efficient text preprocessing removes unnecessary characters to minimise billable counts. Asynchronous processing for long-form content provides more economical synthesis than real-time requests, and proper error handling with retry logic prevents unnecessary API calls.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/polly/latest/dg/what-is.html
Getting Started: https://docs.aws.amazon.com/polly/latest/dg/getting-started.html
API Reference: https://docs.aws.amazon.com/polly/latest/dg/API_Reference.html
Service Quotas: https://docs.aws.amazon.com/polly/latest/dg/limits.html

Amazon Quick

General Information

Amazon Quick is an agentic AI-powered workspace that combines research, business intelligence, and automation capabilities into a unified digital platform. It retrieves insights from business data across internet sources and documents, then transitions from analysis to taking action in applications like Jira and ServiceNow.

Main Benefits

Amazon Quick reduces application switching by consolidating research, business intelligence, and automation into one unified platform. It delivers contextualised insights in minutes, reduces operational overhead, and enables faster decision-making while maintaining secure data handling where queries are never used to train models.

Key Use Cases

- **Research and Analysis:** Conducting comprehensive research across enterprise data and external sources for actionable insights
- **Business Intelligence:** Transforming data into insights through natural language queries and interactive visualisations
- **Process Automation:** Automating repetitive tasks and complex business processes from simple to enterprise-scale workflows
- **Cross-Application Integration:** Taking actions in business applications like creating Jira tickets directly from insights.

Core Features

- **Quick Index:** Unified knowledge foundation consolidating documents, files, and application data for AI-driven insights
- **Quick Research:** Powerful agent conducting comprehensive research across enterprise data and external sources for contextualised insights
- **Quick Sight:** AI-powered business intelligence transforming data into actionable insights through natural language queries and visualisations
- **Quick Flows:** Natural language workflow automation tool for describing and automating repetitive tasks efficiently
- **Quick Automate:** Enterprise-scale process automation transforming complex business processes into multi-agent workflows for departments
- **Spaces:** Straightforward organisation system allowing users to add context and organise work within the platform
- **Cross-Application Integration:** Direct action execution in business applications like Jira and ServiceNow from insights
- **AI-Powered Teammates:** Comprehensive business productivity support handling simple tasks to complex multi-department workflows.

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/quicksuite/

FAQ: https://aws.amazon.com/quicksuite/faqs/

Technical Overview

Backup and Recovery Capabilities

Backup: Amazon Quick offers automatic backup capabilities as part of its managed service infrastructure with data persistence through AWS's underlying storage systems. The service includes automatic replication and point-in-time recovery capabilities built into platform architecture, with data replicated across multiple Availability Zones within AWS Regions for resilience.

Recovery: Amazon Quick supports disaster recovery through AWS's multi-Region infrastructure and automatic failover capabilities, designed for high availability with recovery from Regional outages through proven disaster recovery frameworks. Supports customer RPO (serverless, highly available architecture) and RTO (minimal downtime during recovery scenarios) objectives.

Setup Process

Onboarding: Amazon Quick requires AWS administrator activation in a few steps through the AWS Management Console with AWS Identity and Access Management Identity Center for authentication. The service can be quickly deployed without infrastructure setup, with guided tutorials and sample data for exploring capabilities.

Offboarding: Customers can remove their data through standard AWS data deletion practices. The service follows AWS's standard data retention and deletion policies for managed services with secure data handling.

Data Management

- **Storage Options:** AWS underlying storage systems with automatic replication and data persistence across multiple Availability Zones
- **Data Removal:** Standard AWS data deletion practices with secure data handling following AWS retention and deletion policies
- **Cross-Region Replication:** Multi-Region infrastructure supporting disaster recovery with cross-Region replication strategies for high availability
- **Encryption:** Secure data handling with encryption at rest and in transit for comprehensive data protection.

Optimise Cost and Consumption

Amazon Quick provides consumption-based monitoring for Quick Index and optional features, enabling customers to track and align usage with business requirements. The per-user subscription model allows precise user management to control costs based on actual workforce needs. Customers can optimise consumption by monitoring Quick Index usage patterns and adjusting feature usage accordingly. The serverless architecture automatically scales resources, reducing infrastructure overhead while maintaining performance efficiency.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/quicksuite/latest/userguide/

Getting Started: https://docs.aws.amazon.com/quicksuite/latest/userguide/quick-sight-getting-started.html

API Reference: https://docs.aws.amazon.com/quicksuite/latest/api/

Service Quotas: https://docs.aws.amazon.com/general/latest/gr/quicksuite.html

Amazon Redshift

General Information

Amazon Redshift is a fully managed, petabyte-scale cloud data warehouse service for analysing structured data using standard SQL and business intelligence tools. It uses massively parallel processing architecture with columnar storage to deliver fast performance at scale for complex analytical workloads.

Main Benefits

Amazon Redshift delivers up to 6x better price performance compared to other cloud data warehouses and up to 7x better performance for high concurrency workloads. It reduces traditional data warehousing complexity with automatic scaling, serverless options, and AWS ecosystem integration while scaling from gigabytes to petabytes.

Key Use Cases

- **Business Intelligence and Reporting:** Running enterprise BI workloads, generating dashboards, and creating analytical reports
- **Real-time Analytics:** Analysing streaming data from IoT devices, clickstreams, and social media for insights
- **Data Lake Analytics:** Querying Amazon Simple Storage Service (Amazon S3) data using Amazon Redshift Spectrum without loading into warehouse
- **ML and Predictive Analytics:** Creating ML models for fraud detection, risk scoring, and churn prediction.

Core Features

- **RA3 Instances with Managed Storage:** Separate compute and storage scaling with high-performance query processing and automatic storage management
- **Concurrency Scaling:** Automatically provisions additional compute capacity to handle varying workload demands without performance degradation
- **Redshift Spectrum:** Query data directly in Amazon S3 using SQL without loading it into Amazon Redshift tables
- **Zero-ETL Integrations:** Near real-time data replication from Amazon Aurora, Amazon Relational Database Service (Amazon RDS), Amazon DynamoDB, and enterprise applications without complex pipelines
- **Serverless Option:** Automatically manages compute capacity with AI-driven scaling and optimisation without infrastructure management
- **Data Sharing:** Secure sharing of live data between separate Amazon Redshift clusters without copying or moving data
- **Redshift ML:** Create, train, and apply ML models using SQL commands with SageMaker integration
- **Multi-AZ Deployment:** High availability with automatic failover and data replication across multiple Availability Zones (AZs).

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/redshift/

FAQ: https://aws.amazon.com/redshift/faqs/

Technical Overview

Backup and Recovery Capabilities

Backup: Automated and manual snapshots stored in Amazon S3 with configurable retention periods can be created manually (AWS Management Console, CLI, API) or automated via AWS Backup with policy-based scheduling every eight hours or after 5GB changes. Backup retention configurable from 1-35 days for automated snapshots, up to 3,653 days for manual snapshots.

Recovery: Point-in-time recovery through continuous snapshots, cross-Region snapshot copying for disaster recovery, and Multi-AZ deployments with automatic failover. Sub-minute recovery times with synchronous data replication. Supports customer RPO (automated snapshots every 8 hours or 5GB changes) and RTO (sub-minute automatic failover) objectives.

Setup Process

Onboarding: Create provisioned clusters or serverless workgroups through AWS Management Console, CLI, or API. Choose node types and cluster size for provisioned or base capacity settings for serverless. Automated VPC, security group, and parameter group setup with sensible defaults enables querying within minutes.

Offboarding: Delete clusters or serverless workgroups to remove compute resources. Manual snapshots require explicit deletion for permanent data removal, while automated snapshots delete with cluster termination. Complete cleanup includes removing cross-Region snapshots and associated resources.

Data Management

- **Storage Options:** RA3 instances with managed storage enabling separate compute and storage scaling with columnar storage technology
- **Data Removal:** Delete clusters or serverless workgroups, explicitly remove manual snapshots, and clean up cross-Region snapshot copies
- **Cross-Region Replication:** Snapshots can be copied across AWS Regions for disaster recovery and business continuity requirements
- **Encryption:** Built-in security features including end-to-end encryption for data protection and compliance requirements.

Optimise Cost and Consumption

Reserved Instances deliver up to 75% savings while Serverless Reservations provide up to 24% discounts on committed customer usage. RA3 instances with managed storage prevent over-provisioning by scaling storage independently from compute. Pause/resume functionality reduces development environment costs, while Amazon Redshift Spectrum enables querying cold data in Amazon S3 without loading into the warehouse. Concurrency Scaling provides one hour no-cost daily, and proper data compression with efficient table design optimises query performance and reduces compute usage.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/redshift/

Getting Started: https://docs.aws.amazon.com/redshift/latest/gsg/

API Reference: https://docs.aws.amazon.com/redshift/latest/APIReference/

Service Quotas: https://docs.aws.amazon.com/redshift/latest/mgmt/amazon-redshift-limits.html

Amazon Rekognition

General Information

Amazon Rekognition is a cloud-based computer vision service that analyses images and videos using deep learning technology. It detects objects, faces, text, celebrities, and inappropriate content without requiring ML expertise, supporting both real-time streaming and stored media analysis.

Main Benefits

Amazon Rekognition reduces the complexity of building custom computer vision solutions with pre-trained, highly accurate deep learning models accessible through simple APIs. It processes billions of images and videos daily with unmatched scalability, requires no ML expertise, consumption-based pricing, and AWS Free Tiers.

Key Use Cases

- **Identity Verification:** Verify user identities through face matching and liveness detection capabilities
- **Content Moderation:** Automatically detect and filter inappropriate, explicit, or unsafe content in media
- **Security and Surveillance:** Monitor premises through real-time detection of people, objects, and activities
- **Media Libraries:** Enable content discovery and organisation through automated tagging and metadata extraction.

Core Features

- **Object and Scene Detection:** Detects thousands of objects, scenes, concepts, and activities with confidence scores and bounding boxes
- **Facial Analysis and Recognition:** Analyses facial attributes, emotions, landmarks, and enables face comparison and search capabilities
- **Text Detection and Extraction:** Detects and extracts text from images and videos with word-level confidence scores
- **Celebrity Recognition:** Identifies thousands of celebrities with additional metadata including names and IMDb IDs
- **Content Moderation:** Detects explicit, suggestive, or violent content to enable automated content filtering
- **Custom Labels:** Creates custom models to detect specific objects, scenes, or concepts unique to business needs
- **Video Streaming Events:** Processes real-time video streams for low-latency object detection including people, pets, and packages
- **PPE Detection:** Identifies personal protective equipment including face covers, head covers, and hand covers.

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/rekognition/

FAQ: https://aws.amazon.com/rekognition/faqs/

Technical Overview

Backup and Recovery Capabilities

Backup: Face collections and user vectors are automatically replicated across multiple Availability Zones (AZs) with high durability. Customers can backup face collections by exporting face vectors and metadata manually via APIs. Analysis results stored in integrated services like Amazon DynamoDB or Amazon Simple Storage Service (Amazon S3) can use AWS Backup for automated scheduling.

Recovery: Multi-AZ architecture enables rapid recovery with built-in disaster recovery capabilities. Cross-Region face collection replication and multi-Region architectures provide comprehensive disaster recovery options. Supports low RPO objectives (near-zero RPO for face data through automatic replication) and RTO objectives (rapid recovery via API-based architecture with Regional failover capabilities).

Setup Process

Onboarding: Setup requires AWS account creation and AWS Identity and Access Management (AWS IAM) permissions configuration via web AWS Management Console or AWS CLI. Install SDKs for preferred programming languages. Amazon provides AWS CloudFormation templates, sample applications, and comprehensive developer guides with step-by-step tutorials for accelerated deployment.

Offboarding: Delete face collections using DeleteCollection and DeleteFaces operations through AWS Management Console or APIs. Remove stored analysis results and revoke AWS IAM permissions for complete termination. AWS CloudTrail enables auditing of data deletion activities.

Data Management

- **Storage Options:** Face collections stored with automatic multi-AZ replication, analysis results in Amazon DynamoDB or Amazon S3
- **Data Removal:** Delete images, videos, and face collections anytime via AWS Management Console or DeleteCollection/DeleteFaces API operations
- **Cross-Region Replication:** Customers can replicate face collections across Regions for comprehensive disaster recovery architectures
- **Encryption:** Enterprise-grade security with Health Insurance Portability and Accountability Act eligibility and encryption capabilities for compliance requirements.

Optimise Cost and Consumption

Amazon Rekognition provides tiered pricing that automatically reduces costs per image with higher processing volumes. Confidence threshold controls enable processing only relevant content, reducing unnecessary analysis calls. Custom moderation adapters minimise false positives to decrease reprocessing requirements. Batch processing capabilities improve efficiency for large datasets, while Custom Labels focus detection on specific business objects. Regional pricing variations allow strategic Region selection for optimal costs.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/rekognition/latest/dg/what-is.html
Getting Started: https://docs.aws.amazon.com/rekognition/latest/dg/getting-started.html
API Reference:
https://docs.aws.amazon.com/rekognition/latest/APIReference/Welcome.html
Service Quotas: https://docs.aws.amazon.com/rekognition/latest/dg/limits.html

Amazon Relational Database Service (Amazon RDS)

General Information

Amazon Relational Database Service (Amazon RDS) is a fully managed relational database service that automates essential administration tasks including hardware provisioning, patching, backups, and recovery. It supports multiple database engines like MySQL, PostgreSQL, and Oracle while providing cost-efficient, resizable capacity with existing application compatibility.

Main Benefits

General RDS reduces administrative overhead by up to 43% compared to self-managed databases while delivering 99.95% availability SLA with Multi-AZ deployments. The service offers up to 69% cost savings with Reserved Instances, automatic failover in 1-2 minutes, and point-in-time recovery capabilities with enterprise-grade security and performance monitoring.

Key Use Cases

- **Web and Mobile Applications:** Supporting transactional workloads with high availability and automated scaling
- **Enterprise Applications:** Running business-critical applications requiring ACID compliance and complex queries
- **E-commerce Platforms:** Managing product catalogues, customer data, and transaction processing with security
- **Content Management Systems:** Storing structured content with read replicas for improved performance.

Core Features

- **Multi-Availability Zone (AZ) Deployments:** Provides synchronous replication to standby instance for high availability and automatic failover
- **Read Replicas:** Creates read-only database copies to scale read workloads and improve application performance
- **Automated Backups:** Performs daily automated backups with point-in-time recovery capability up to 35 days
- **Security Features:** Provides encryption at rest and in transit, Virtual Private Cloud (Amazon VPC) isolation, and AWS Identity and Access Management (AWS IAM) integration
- **Performance Insights:** Offers database performance monitoring and analysis tools for query optimisation and troubleshooting
- **Blue/Green Deployments:** Enables zero-downtime database updates and deployments through managed environment switching capabilities
- **Multiple Database Engines:** Supports MySQL, PostgreSQL, MariaDB, Oracle Database, Microsoft SQL Server, and IBM Db2
- **Automated Administration:** Automates hardware provisioning, database setup, patching, and recovery tasks for operational efficiency.

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/rds/
FAQ: https://aws.amazon.com/rds/faqs/

Technical Overview

Backup and Recovery Capabilities

Backup: Automated daily backups with transaction log captures stored in Amazon Simple Storage Service can be created manually (AWS Management Console, CLI, API) or automated via AWS Backup with configurable retention up to 35 days. Cross-Region copying and backup encryption capabilities are included with indefinite manual snapshot retention.

Recovery: Point-in-time recovery to any second within retention period, Multi-AZ automatic failover in 1-2 minutes, and cross-Region disaster recovery through read replica promotion. Multi-AZ deployments provide synchronous replication for fault tolerance. Supports customer RPO (minutes via 5-minute transaction logs) and RTO (1-2 minutes for failover) objectives.

Setup Process

Onboarding: Setup through Amazon RDS Console, CLI, or API using Easy Create option with guided workflows. Configure instance class, storage type, VPC security groups, and parameter groups. Initial provisioning takes 10-20 minutes with automated deployment and Free Tier eligibility.

Offboarding: Delete date base instances via AWS Management Console, CLI, or API with optional final snapshot creation. Automated backups are removed unless retained, while manual snapshots persist until explicitly deleted with centralised deletion control.

Data Management

- **Storage Options:** Multiple storage types including gp2, gp3, and io1 with up to 64 TiB capacity
- **Data Removal:** Delete instances via AWS Management Console, CLI, or API with optional final snapshots before deletion
- **Cross-Region Replication:** Supported through cross-Region read replicas and automated backups for disaster recovery capabilities
- **Encryption:** Provides encryption at rest and in transit with AWS Key Management Service integration for key management.

Optimize Cost and Consumption

Reserved Instances deliver up to 69% savings through 1-year or 3-year commitments with instance size flexibility within families. AWS Compute Optimizer provides right-sizing recommendations based on actual usage patterns, reducing costs by 20-40%. Read replicas offload primary instance workloads enabling smaller primary configurations. Storage optimisation across gp2, gp3, and io1 types matches performance requirements. Graviton-based instances offer superior price-performance ratios while automated backup lifecycle management controls storage expenses.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/Welcome.html
Getting Started: https://docs.aws.amazon.com/AmazonRDS/latest/gettingstartedguide/what-is-rds.html
API Reference: https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/ProgrammingGuide.html
Service Quotas: https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/CHAP_Limits.html

Amazon Route 53

General Information

Amazon Route 53 is a highly available Domain Name System (DNS) web service that performs domain registration, DNS routing, health checking, and resolution. It translates domain names into IP addresses using a global network of authoritative DNS servers and provides comprehensive management for public and private domains.

Main Benefits

Route 53 delivers reliability with 100% availability SLA and low-latency DNS responses through its global anycast network. Key differentiators include advanced routing policies for intelligent traffic distribution, AWS service integration, enterprise-grade security with domain Name System Security Extensions (DNSSEC), and 60-minute RTO through Accelerated Recovery during Regional disruptions.

Key Use Cases

- **DNS Traffic Routing:** Route internet traffic using latency-based, geolocation, weighted, and failover routing policies
- **Health Monitoring and Failover:** Monitor endpoint health and automatically redirect traffic from unhealthy resources
- **Hybrid DNS Resolution:** Enable DNS resolution between on-premises networks and Amazon Virtual Private Clouds (Amazon VPCs)
- **Global Traffic Management:** Distribute traffic across multiple AWS Regions based on performance and location requirements.

Core Features

- **DNS Hosting:** Public and private hosted zones for authoritative DNS management with global anycast network
- **Advanced Routing Policies:** Latency-based, geolocation, geoproximity, weighted, failover, IP-based, and multivalue routing options
- **Health Checks:** Automated endpoint monitoring with DNS failover and Amazon CloudWatch integration for application availability
- **Route 53 Resolver:** Recursive DNS resolution between VPCs and on-premises networks with conditional forwarding
- **Domain Registration:** Register and manage domains with privacy protection and auto-renewal options
- **DNSSEC Support:** DNS Security Extensions for authentication and integrity of DNS responses
- **Traffic Flow:** Visual editor for complex routing policies and global traffic management
- **Resolver DNS Firewall:** Filter and block malicious DNS queries with customisable security rules.
- Route 53 Application Recovery Controller monitors recovery readiness and controls failover across multiple AWS environments.

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/route53/
FAQ: https://aws.amazon.com/route53/faqs/

Technical Overview

Backup and Recovery Capabilities

Backup: Amazon Route 53 offers inherent backup through globally distributed architecture with automatic DNS data replication across multiple edge locations. DNS records and configurations can be exported manually via APIs, CLI, or third-party tools for external storage, with configuration history maintained through CloudTrail logging.

Recovery: Accelerated Recovery offers 60-minute RTO during Regional disruptions with automatic control plane failover. The global anycast network helps ensure continued DNS resolution during localised outages, while health checks enable automatic endpoint failover. Supports sub-minute RPO through real-time replication and seconds-level RTO via DNS failover with 10-second health check intervals.

Setup Process

Onboarding: Setup requires AWS account with AWS Identity and Access Management permissions for Amazon Route 53 access. Access via AWS S Management Console, APIs, CLI, or SDK. Create hosted zones, configure DNS records, and register domains with step-by-step tutorials for common scenarios like S3 website routing.

Offboarding: Delete hosted zones after removing all records except Name Server and Start of Authority. Transfer domain registrations to another registrar before deletion. Remove health checks, traffic policies, and resolver configurations individually with programmatic deletion via APIs.

Data Management

- **Storage Options:** DNS records stored in globally distributed hosted zones with automatic replication across edge locations
- **Data Removal:** Delete hosted zones, records, health checks, and domain configurations via AWS Management Console or API programmatically
- **Cross-Region Replication:** Automatic DNS data replication across multiple global edge locations with built-in redundancy
- **Encryption:** DNSSEC support with customer-managed AWS Key Management Service keys and encrypted queries through Global Resolver.

Optimise Cost and Consumption

Amazon Route 53 offers alias records that provide cost-free routing to AWS resources, reducing DNS query charges. Optimising TTL values balances query responsiveness with volume costs, while appropriate health check intervals minimise unnecessary monitoring expenses. Wildcard records consolidate DNS zones efficiently under the pay-per-query model. Geographic routing directs traffic to cost-effective Regions, and Route 53 Profiles enable centralised management across multiple accounts for operational savings.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/Welcome.html
Getting Started: https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/getting-started.html
API Reference: https://docs.aws.amazon.com/Route53/latest/APIReference/Welcome.html
Service Quotas: https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/DNSLimitations.html

Amazon SageMaker and Amazon SageMaker AI

General Information

Amazon SageMaker is a fully managed ML platform providing end-to-end ML capabilities for data scientists and developers to build, train, and deploy high-quality ML models at scale. It supports both traditional machine learning workflows and generative AI development with comprehensive tools including notebooks, debuggers, pipelines, MLOps features, and automated capabilities for production-ready hosted environments.

Main Benefits

Amazon SageMaker delivers 54% lower total cost of ownership compared to self-managed solutions while providing up to 40% faster training with HyperPod infrastructure. The platform offers unmatched flexibility with support for any ML framework and the broadest selection of tools in a single platform, uniquely combining traditional ML and generative AI development with comprehensive tools from no-code to expert-level environments. Key differentiators include managed Spot training reducing costs up to 90%, automated scaling from experimentation to production, and elimination of operational overhead through fully managed infrastructure.

Key Use Cases

- **Foundation Model Development:** Building and training large-scale foundation models using Amazon SageMaker HyperPod infrastructure
- **Automated ML:** Creating ML models without coding using Canvas and Autopilot services
- **Enterprise MLOps:** Implementing end-to-end workflows with governance, monitoring, and automated deployment pipelines
- **Edge AI Deployment:** Deploy and manage optimised ML models on edge devices.
- **Predictive Analytics:** Revenue forecasting, demand planning, customer churn prediction, and business intelligence applications
- **Computer Vision:** Image classification, object detection, medical image analysis, quality control, and autonomous vehicle development
- **Natural Language Processing:** Sentiment analysis, text classification, chatbots, document processing, and language translation
- **Fraud Detection:** Credit card fraud detection, anomaly detection in financial transactions, and risk assessment.

Core Features

- **SageMaker Studio:** Fully integrated development environment with JupyterLab, Code Editor, and collaborative features
- **SageMaker Autopilot:** Automated ML service that builds, trains, and tunes models with full visibility
- **SageMaker Canvas:** Visual, no-code interface for building ML models without requiring programming skills
- **SageMaker HyperPod:** Purpose-built infrastructure for distributed training with up to 40% faster performance
- **SageMaker Feature Store:** Centralised repository for storing, sharing, and managing ML features across teams
- **SageMaker Pipelines:** Continuous Integration/Continuous Deployment service for creating, managing, and executing end-to-end ML workflows

- **SageMaker Clarify:** Bias detection and model explainability service for responsible AI development
- **SageMaker Inference (Edge):** Deploy models to edge devices via AWS IoT Greengrass" if edge capability still needed.
- **SageMaker Training:** Managed model training with distributed training, hyperparameter tuning, and spot instance support
- **SageMaker Inference:** Multiple deployment options including real-time, batch, asynchronous, and serverless inference
- **SageMaker JumpStart:** Pre-trained models, algorithms, and solution templates for common ML use cases

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/sagemaker/ai/

FAQ: https://aws.amazon.com/sagemaker/ai/faqs/

Technical Overview

Backup and Recovery Capabilities

Backup: Amazon SageMaker provides backup capabilities through integration with AWS services. Model artifacts and training data stored in Amazon S3 benefit from cross-Region replication, versioning, and lifecycle policies, with training checkpoints enabling recovery for long-running jobs. The Model Registry provides versioning and metadata management for deployed models. Automated backup can be configured via AWS CloudTrail, Amazon EventBridge, and AWS Lambda for event-driven solutions, with AWS Backup integration for centralised management and comprehensive backup strategies requiring custom solutions using AWS Lambda functions and Amazon EventBridge.

Recovery: Amazon SageMaker supports recovery through multi-Region deployments and cross-Region replication, including Amazon EFS replication for SageMaker domains and Amazon S3 cross-Region replication for model artifacts and training data. Infrastructure as Code templates enable rapid environment recreation, while multi-Availability Zone deployments provide high availability. Managed spot training includes automatic recovery from instance interruptions, and automated model deployment pipelines support rapid redeployment and scaling for failover scenarios. The service supports customer RPO objectives through cross-Region backup strategies and real-time data replication, and RTO objectives through real-time model monitoring, automated recovery workflows, pre-provisioned inference endpoints, and automated failover.

Setup Process

Onboarding: Customers can get started through quick setup for individual users with default configurations or custom setup for enterprise ML administrators. Multiple entry points include Amazon SageMaker Studio, Canvas for no-code development, and Studio Lab for learning. The onboarding process includes domain creation, user profile setup, Virtual Private Cloud configuration and security controls for enterprise deployments, and access to various development environments based on organisational requirements.

Offboarding: Data removal involves deleting Amazon SageMaker domains, user profiles, and associated resources including training data, model artifacts, and notebook files through the SageMaker Console or APIs. Customers must manually delete model artifacts stored in

Amazon S3, Feature Store entries, stop notebook instances, terminate endpoints, delete Amazon S3 buckets, and revoke AWS Identity and Access Management roles, policies, and permissions across distributed AWS service components. Shared Feature Store resources should be managed appropriately during the offboarding process.

Data Management

- Storage Options: Amazon S3 integration for scalable data lake storage and model artifacts, with Feature Store for centralised ML features and feature management
- Data Removal: Manual deletion of model artifacts, training data, Feature Store entries, and distributed AWS components through Amazon SageMaker Console or APIs with full customer control over data lifecycle management
- Cross-Region Replication: Supported via Amazon EFS replication for Amazon SageMaker domains and Amazon S3 cross-Region replication for model artifacts and training data, enabling cross-Region disaster recovery and backup
- Encryption: AWS Key Management Service integration for encryption with enterprise-grade security controls, AWS IAM integration for access control, data protection, and governance features

Optimise Cost and Consumption

Amazon SageMaker offers comprehensive cost optimization through ML Savings Plans delivering up to 64% cost reduction for consistent workloads and managed Spot training reducing costs by up to 90% for non-time-sensitive jobs with automatic fault tolerance and checkpoint recovery. Serverless inference provides usage-based pricing for variable traffic patterns and automatically scales to zero during idle periods, while multi-model endpoints enable resource sharing across models to reduce deployment costs. Built-in optimization tools include intelligent instance selection recommendations, automated resource scaling to prevent over-provisioning, notebook instance lifecycle management to prevent idle charges, and rightsizing recommendations. Asynchronous and batch inference options offer flexible, cost-effective processing alternatives for various workload patterns.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/sagemaker/latest/dg/index.html

Getting Started: https://aws.amazon.com/getting-started/hands-on/build-train-deploy-machine-learning-model-sagemaker/

API Reference: https://docs.aws.amazon.com/sagemaker/latest/dg/reference.html

Service Quotas: https://docs.aws.amazon.com/sagemaker/latest/dg/training-plan-quotas.html

Amazon Simple Email Service (Amazon SES)

General Information

Amazon Simple Email Service (Amazon SES) is a scalable email service for sending and receiving emails using your own domains. It helps reduce email infrastructure complexity while supporting marketing campaigns, transactional emails, and automated email processing workflows.

Main Benefits

Amazon SES reduces email infrastructure complexity while using Amazon's enterprise-grade email infrastructure. It offers consumption-based pricing with no upfront costs, high deliverability rates through managed IP pools, AWS service integration, and scalable infrastructure supporting volumes from small businesses to enterprise operations.

Key Use Cases

- **Marketing Campaigns:** Send promotional emails, special offers, and newsletters to customers
- **Transactional Emails:** Deliver order confirmations, password resets, and account notifications automatically
- **Email Receiving:** Process incoming emails for autoresponders, unsubscribe systems, and support tickets
- **Newsletter Distribution:** Create and manage email templates for personalised bulk messaging campaigns.

Core Features

- **Email Sending:** Send emails via SMTP interface, API calls, or AWS SDKs with formatted support
- **Email Receiving:** Process incoming emails with receipt rules, IP filtering, and AWS service integration
- **Virtual Deliverability Manager:** Suite of tools to track and improve email delivery rates and reputation
- **Dedicated IP Addresses:** Managed or standard dedicated IPs for controlling and maintaining sender reputation
- **Mail Manager:** Advanced traffic management with filtering policies, archiving, and security add-ons
- **Template Support:** Create and manage email templates for personalised bulk messaging campaigns
- **Event Publishing:** Track email events like bounces, complaints, and deliveries through Amazon CloudWatch and Amazon Simple Notification Service
- **DomainKeys Identified Mail (DKIM) Authentication:** Easy DKIM setup and email authentication features to improve email deliverability

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/ses/

FAQ: https://aws.amazon.com/ses/faqs/

Technical Overview

Backup and Recovery Capabilities

Backup: Amazon SES does not provide traditional backup capabilities as it processes emails in real-time without storing content. For email receiving, backup depends on configured Amazon Simple Storage Service (Amazon S3) storage destinations with their own durability and versioning. Service configurations can be recreated through Infrastructure as Code (IaC) practices.

Recovery: Amazon SES supports disaster recovery through multi-Region availability and configuration replication across AWS Regions using automated deployment scripts. Service configurations, identities, and templates should be version-controlled for recovery scenarios. Supports customer RPO (essentially zero for real-time processing) and RTO (minimised through multi-Region deployments) objectives.

Setup Process

Onboarding: Setup requires AWS account creation and email address or domain verification. New accounts start in sandbox environment with limited capabilities, then request production access through AWS Support Center. Configure sender identities, DKIM authentication, and optional dedicated IPs.

Offboarding: Delete configuration sets, verified identities, email templates, and receipt rules. Remove stored email content from connected Amazon S3 services. Account closure automatically removes Amazon SES configurations and stops billing.

Data Management

- **Storage Options:** Real-time email processing without persistent storage; email receiving uses Amazon S3 buckets for retention
- **Data Removal:** Emails processed in real-time without default storage; delete configurations, templates, and Amazon S3 content
- **Cross-Region Replication:** Multi-Region availability with configuration replication using IaC practices for disaster recovery
- **Encryption:** AWS Key Management Service integration for email encryption with comprehensive security through integrated AWS services.

Optimise Cost and Consumption

Amazon SES offers tiered pricing for Virtual Deliverability Manager that reduces costs for high-volume senders. Strategic Region selection minimises data transfer costs while shared IP pools provide cost-effective alternatives to dedicated IPs for lower-volume operations. Implementing efficient email practices through bounce and complaint handling reduces wasted sending costs by improving deliverability rates and sender reputation management.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/ses/latest/dg/Welcome.html

Getting Started: https://docs.aws.amazon.com/ses/latest/dg/getting-started.html

API Reference: https://docs.aws.amazon.com/ses/latest/APIReference/Welcome.html

Service Quotas: https://docs.aws.amazon.com/ses/latest/dg/quotas.html

Amazon Simple Notification Service (Amazon SNS)

General Information

Amazon Simple Notification Service (Amazon SNS) is a fully managed messaging service that delivers messages from publishers to subscribers through asynchronous communication. Amazon SNS supports both Application-to-Application and Application-to-Person messaging across multiple endpoints including Amazon Simple Queue Service (SQS), AWS Lambda, email, short message service (SMS), and mobile push notifications.

Main Benefits

Amazon SNS helps reduce operational overhead of messaging infrastructure while automatically scaling to handle millions of messages per second. The service provides immediate push-based delivery, built-in multi-Availability Zone (AZ) durability, AWS integration, and consumption-based pricing with no upfront commitments, enabling loose coupling between system components.

Key Use Cases

- **Fanout Messaging:** Replicate messages to multiple endpoints enabling decoupled microservices architectures
- **Application Monitoring:** Send threshold-triggered alerts via SMS, email for system monitoring and alerting
- **Mobile Push Notifications:** Direct message delivery to iOS, Android applications for user engagement
- **User Notifications:** E-commerce confirmations, account updates, and personalised communications to individuals or groups.

Core Features

- **Standard and 'First-In, First-Out' (FIFO) Topics:** Standard topics provide high throughput; FIFO helps ensure strict ordering and exactly-once processing
- **Message Filtering:** Subscribers receive only relevant notifications through filter policies based on message attributes or content
- **Message Archiving and Replay:** FIFO topics archive messages for 365 days with replay capabilities for recovery
- **Multi-Protocol Support:** Delivers messages to Amazon SQS, AWS Lambda, HTTP(S), email, SMS, mobile push notifications, and Amazon Data Firehose
- **Server-Side Encryption:** Protects messages using AWS Key Management Service (AWS KMS) with support for AWS managed and customer managed keys
- **Message Durability:** Helps ensure persistence through multi-server storage across data centres with delivery retry policies and dead-letter queues
- **Automatic Scaling:** Handles millions of messages per second with immediate push-based delivery and no operational overhead
- **AWS Service Integration:** Deep integration with AWS Lambda, Amazon SQS, Amazon EventBridge, AWS Step Functions, and Amazon CloudWatch for comprehensive event-driven architectures.

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/sns/

FAQ: https://aws.amazon.com/sns/faqs/

Technical Overview

Backup and Recovery Capabilities

Backup: FIFO topics provide built-in message archiving for up to 365 days with automatic multi-AZ replication across servers and data centres. Topic configurations can be backed up manually (AWS Management Console, CLI, SDKs) or automated via AWS CloudFormation Infrastructure as Code (IaC) templates.

Recovery: Message replay capabilities enable recovery from FIFO topic archives, multi-Region deployments support cross-Region failover, and built-in durability features ensure high availability. Multi-AZ redundancy provides fault tolerance. Supports customer RPO (near-zero through message replay) and RTO (sub-minute via automated failover) objectives.

Setup Process

Onboarding: Setup via Console, CLI, or SDKs after creating AWS account and AWS Identity and Access Management permissions. Create topics, configure subscriptions, set message filtering and encryption. AWS provides comprehensive tutorials and code examples.

Offboarding: Delete topics through AWS Management Console, CLI, or SDKs with automatic subscription removal. FIFO archived messages are permanently deleted. Cleanup via IaC automation is available.

Data Management

- **Storage Options:** FIFO topics archive messages for 365 days; automatic multi-AZ storage across servers and data centres
- **Data Removal:** Topic deletion removes all subscriptions automatically; archived FIFO messages permanently deleted via AWS Management Console, CLI, or SDKs
- **Cross-Region Replication:** Multi-Region deployment supported through Iac templates for disaster recovery strategies
- **Encryption:** Server-side encryption using AWS KMS with AWS managed and customer managed key support.

Optimise Cost and Consumption

Amazon SNS provides message batching capabilities that publish up to 10 messages in single requests, reducing API call costs. Attribute-based message filtering reduces unwanted deliveries at the subscription level, minimising notification charges. Message payload optimisation keeps sizes under 64KB chunks to avoid additional billing increments. Efficient retry policies prevent excessive failed delivery attempts, while strategic Regional deployments reduce cross-Region data transfer costs.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/sns/latest/dg/welcome.html

Getting Started: https://docs.aws.amazon.com/sns/latest/dg/sns-setting-up.html

API Reference: https://docs.aws.amazon.com/sns/latest/api/

Service Quotas: https://docs.aws.amazon.com/general/latest/gr/sns.html

Amazon Simple Queue Service (Amazon SQS)

General Information

Amazon Simple Queue Service (Amazon SQS) is a fully managed message queuing service that enables decoupling of distributed software systems and components. Amazon SQS provides secure, durable queues for transmitting messages between applications with two queue types: Standard and FIFO.

Main Benefits

Amazon SQS helps reduce infrastructure maintenance and scaling concerns with 99.9% availability SLA and automatic scaling. The fully managed service offers virtually unlimited throughput for standard queues, message batching reducing costs by up to 90%, and AWS environment integration, requiring no capacity planning or cluster management.

Key Use Cases

- **Decoupling Microservices:** Enable loose coupling between application components for improved fault tolerance and scalability
- **Work Queues:** Distribute tasks across multiple worker processes for parallel processing and load balancing
- **Buffer Operations:** Handle traffic spikes by queuing requests for batch processing during off-peak times
- **Event-Driven Architectures:** Enable asynchronous communication between services in serverless and container-based applications.

Core Features

- **Standard and FIFO Queues:** Standard queues offer unlimited throughput while FIFO queues provide exactly-once processing with guaranteed ordering
- **Server-Side Encryption:** Encrypt messages at rest using Amazon SQS-managed keys or AWS Key Management Service (AWS KMS) customer-managed keys for enhanced security
- **Dead-Letter Queues:** Handle failed messages by automatically moving them to separate queues after maximum receive attempts
- **Long Polling:** Reduce costs and improve performance by waiting for messages instead of frequent empty requests
- **Message Batching:** Send, receive, and delete up to 10 messages in single API call for improved throughput
- **Visibility Timeout:** Prevent multiple consumers from processing same message by making it temporarily invisible during processing
- **Message Retention:** Store messages for up to 14 days with configurable retention periods from 1 minute onwards
- **Fair Queues:** Limit per-tenant consumption rates to prevent noisy neighbour scenarios in multi-tenant applications.

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/sqs/
FAQ: https://aws.amazon.com/sqs/faqs/

Technical Overview

Backup and Recovery Capabilities

Backup: Amazon SQS does not provide traditional backup capabilities as it is designed as a transient messaging service. Messages are automatically stored redundantly across multiple servers with retention up to 14 days for operational purposes. For long-term archival, customers must implement custom solutions using Amazon Simple Storage Service or Amazon DynamoDB via application logic or AWS Lambda functions.

Recovery: Amazon SQS provides built-in disaster recovery through distributed architecture across multiple Availability Zones with automatic failover capabilities. Cross-Region recovery requires customer-implemented replication strategies using multiple queues and AWS Lambda functions. Supports near-zero RPO (synchronous replication) and seconds-level RTO (automatic failover) for regional disasters.

Setup Process

Onboarding: Create queues through AWS Management Console, CLI, or SDKs, choosing between Standard or FIFO types. Configure dead-letter queues, visibility timeouts, server-side encryption, and proper AWS Identity and Access Management (AWS IAM) policies. Proven practices include comprehensive SDKs for multiple programming languages.

Offboarding: Delete queues via AWS Management Console, CLI, or APIs with immediate, irreversible removal of all messages. Update application code to remove queue references and cleanup associated AWS IAM policies and roles for security hygiene.

Data Management

- **Storage Options:** Messages stored up to 14 days with configurable retention periods, automatically replicated across multiple servers
- **Data Removal:** Queue deletion permanently removes all messages immediately and irreversibly via AWS Management Console, CLI, or APIs
- **Cross-Region Replication:** Not natively supported; customers must implement custom solutions using AWS Lambda functions or application logic
- **Encryption:** Server-side encryption at rest using Amazon SQS-managed keys or AWS KMS customer-managed keys for enhanced security.

Optimise Cost and Consumption

Message batching reduces API request costs by up to 90% by processing up to 10 messages per call. Long polling eliminates empty receive requests and associated charges by waiting for messages to arrive. Proper dead-letter queue configuration prevents infinite retry loops that generate unnecessary costs. Fair queues limit per-tenant consumption rates in multi-tenant applications, optimising resource allocation and preventing cost spikes from noisy neighbours.

Technical Resources

Developer Guide:
https://docs.aws.amazon.com/AWSSimpleQueueService/latest/SQSDeveloperGuide/welcome.html

Getting Started: https://docs.aws.amazon.com/AWSSimpleQueueService/latest/SQSDeveloperGuide/sqs-getting-started.html

API Reference: https://docs.aws.amazon.com/AWSSimpleQueueService/latest/APIReference/Welcome.html

Service Quotas: https://docs.aws.amazon.com/AWSSimpleQueueService/latest/SQSDeveloperGuide/sqs-quotas.html

Amazon Simple Storage Service (Amazon S3)

General Information

Amazon Simple Storage Service (Amazon S3) is a scalable object storage service offering virtually unlimited capacity with individual objects up to 5 TB. It is designed to provide 99.999999999% durability and stores data as objects within buckets for web-based retrieval.

Main Benefits

S3 delivers industry-leading 99.999999999% durability and 99.99% availability with unlimited scalability from gigabytes to exabytes. Intelligent-Tiering provides up to 68% cost savings, while eight storage classes and integration with over 100 AWS services maximise flexibility and value.

Key Use Cases

- **Data Lakes and Analytics:** Central repository for structured and unstructured data with analytics integration
- **Backup and Restore:** Reliable scalable backup storage with cross-Region replication capabilities
- **Website Hosting:** Static website hosting and content distribution for web and mobile applications
- **Archive Storage:** Cost-effective long-term archival using Amazon S3 Glacier storage classes for compliance requirements.

Core Features

- **Multiple Storage Classes:** Eight storage classes including Standard, Intelligent-Tiering, and Glacier for cost optimisation
- **Amazon S3 Versioning:** Preserve, retrieve, and restore every version of objects for data protection
- **Cross-Region Replication:** Automatic replication across AWS Regions for redundancy, compliance, and disaster recovery
- **Amazon S3 Object Lock:** Write-once-read-many model preventing object deletion or modification for compliance requirements
- **Access Management:** Comprehensive control through AWS Identity and Access Management (AWS IAM), bucket policies, Access Control Lists, and Amazon S3 Block Public Access
- **Encryption:** Server-side and client-side encryption options including SSE-S3, SSE-KMS, and SSE-C
- **Amazon S3 Batch Operations:** Perform large-scale operations on billions of objects with single API request
- **Event Notifications:** Trigger workflows through AWS Lambda, Amazon Simple Queue Service (Amazon SQS), and Amazon Simple Notification Service (Amazon SNS) when objects are modified.

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/s3/

FAQ: https://aws.amazon.com/s3/faqs/

Technical Overview

Backup and Recovery Capabilities

Backup: Amazon S3 provides 99.999999999% durability with object versioning and Cross-Region Replication for data protection. Backups can be created manually through AWS Management Console, CLI, or APIs, or automated via lifecycle policies and AWS Backup with customisable scheduling. Amazon S3 integrates fully with AWS Backup for centralised management.

Recovery: Point-in-time recovery through versioning, cross-Region Replication for geographic redundancy, and multi-Region Access Points for automatic failover. Multi-Availability Zone storage helps ensure fault tolerance. Supports customer RPO objectives of 15 minutes or less via replication, and RTO objectives of minutes through automated failover capabilities.

Setup Process

Onboarding: Create buckets through AWS Management Console, CLI, or SDKs with AWS IAM user permissions. Configure unique bucket names, regions, and access settings. Upload objects via console interface, programmatic APIs, or bulk transfer tools.

Offboarding: Delete all objects within buckets before bucket deletion. Remove versioned objects, abort multi-part uploads, and revoke access policies. Use Amazon S3 Inventory and Amazon S3 Batch Operations for automated cleanup.

Data Management

- **Storage Options:** Eight storage classes including Standard, Intelligent-Tiering, Standard-IA, One Zone-IA, and Glacier classes
- **Data Removal:** Delete objects before buckets, remove versioned objects and delete markers via lifecycle policies
- **Cross-Region Replication:** Automatic replication across AWS Regions for redundancy, compliance, and disaster recovery capabilities
- **Encryption:** Server-side encryption options including SSE-S3, SSE-KMS, SSE-C plus client-side encryption for data protection

Optimise Cost and Consumption

Amazon S3 Intelligent-Tiering automatically moves data between access tiers based on usage patterns, delivering up to 68% cost savings. Lifecycle policies automate transitions to lower-cost storage classes and delete expired objects. Amazon S3 Storage Class Analysis identifies optimisation opportunities while Amazon S3 Storage Lens provides organisation-wide cost visibility. Amazon S3 Batch Operations reduce per-request charges through bulk processing, and reserved capacity pricing offers discounts for predictable workloads.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/AmazonS3/latest/userguide/Welcome.html

Getting Started: https://docs.aws.amazon.com/AmazonS3/latest/userguide/GetStartedWithS3.html

API Reference: https://docs.aws.amazon.com/AmazonS3/latest/API/

Service Quotas: https://docs.aws.amazon.com/general/latest/gr/s3.html

Amazon Simple Workflow Service (Amazon SWF)

General Information

Amazon Simple Workflow Service (Amazon SWF) is a fully managed workflow service that coordinates work across distributed components with parallel or sequential steps. It manages task state, inter-task dependencies, scheduling, and concurrency for complex application workflows.

Main Benefits

Amazon SWF provides durable state management with near-zero RPO, automatic retry mechanisms, and complete execution history tracking for resilient distributed applications. It eliminates infrastructure complexity while supporting workflows up to one year duration, with no single point of failure and flexible cloud or on-premises worker deployment.

Key Use Cases

- **Video Encoding:** Processing media files through multiple encoding stages with parallel or sequential workflows
- **Data Centre Migration:** Orchestrating complex migration processes with multiple dependent steps and rollback capabilities
- **E-commerce Order Processing:** Coordinating order verification, payment processing, inventory management, and shipping workflows
- **Financial Transactions:** Managing multi-step financial processes with strict consistency and audit requirements.

Core Features

- **Task Coordination:** Manages inter-task dependencies, scheduling, and concurrency across distributed components
- **State Management:** Maintains durable execution state allowing for resilient applications and independent component failures
- **Worker Management:** Assigns tasks to workers running on cloud or on-premises infrastructure with polling-based communication
- **Decision Tasks:** Enables workflow logic through deciders that analyse execution history and make routing decisions
- **Activity Tasks:** Supports activity workers performing business logic with automatic retry and failure handling
- **Lambda Tasks:** Direct integration with AWS Lambda functions for serverless workflow components
- **Child Workflows:** Supports nested workflows for complex process decomposition and reusability
- **Workflow History:** Maintains complete execution history for audit, debugging, and state reconstruction.

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/swf/

FAQ: https://aws.amazon.com/swf/faqs/

Technical Overview

Backup and Recovery Capabilities

Backup: Amazon SWF does not provide traditional backup capabilities as workflow state is maintained internally through managed infrastructure. Workflow definitions, activity types, and execution history are preserved by the service but cannot be directly backed up by customers. Application data requires separate backup using external services like Amazon Simple Storage Service or Amazon DynamoDB.

Recovery: Cross-Region disaster recovery requires manual replication of workflow and activity type registrations across AWS Regions. The stateless worker model enables rapid failover to alternative AWS Regions. Built-in fault tolerance with no single point of failure helps ensure workflow continuity. Supports near-zero RPO (automatic persistence within AWS Region) and RTO objectives (depends on worker restart times and failover implementation).

Setup Process

Onboarding: Access SWF through AWS Management Console, SDKs, or APIs. Register domains for logical groupings, define workflow and activity types, implement activity workers and deciders, then start workflow executions. AWS provides comprehensive documentation, tutorials, and the Flow Framework for Java development.

Offboarding: Stop starting new workflow executions and allow existing executions to complete or terminate manually. Domain deprecation prevents new registrations while preserving existing executions. Data removal occurs automatically through natural retention policies without additional cleanup processes required.

Data Management

- **Storage Options:** Internal managed infrastructure maintains workflow state, execution history, and task completion details durably
- **Data Removal:** Automatic retention management with completed workflow executions deleted after limited period automatically
- **Cross-Region Replication:** Manual replication of workflow definitions and activity types across AWS Regions for disaster recovery
- **Encryption:** Not specified in source material - encryption capabilities not detailed for Amazon SWF service.

Optimise Cost and Consumption

Efficient workflow design minimises task counts and execution duration, directly reducing per-task overhead costs. Batch processing consolidates operations to optimise task usage. Regional consolidation reduces data transfer charges between integrated services. The stateless worker model enables dynamic scaling based on actual demand without maintaining idle resources. Strategic Regional selection optimises data transfer costs, while efficient execution design becomes crucial for long-running workflows charged per 24-hour period.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/amazonswf/latest/developerguide/index.html
Getting Started: https://docs.aws.amazon.com/amazonswf/latest/developerguide/welcome.html
API Reference: https://docs.aws.amazon.com/amazonswf/latest/apireference/Welcome.html
Service Quotas: https://docs.aws.amazon.com/amazonswf/latest/developerguide/swf-dg-limits.html

Note: AWS recommends AWS Step Functions for all new workflow orchestration requirements. Step Functions offers a modern, serverless approach with visual workflows, built-in error handling, and integration with over 200 AWS services.

Amazon Textract

General Information

Amazon Textract is an ML service that automatically extracts text, handwriting, and structured data from documents and images. It goes beyond basic Optical Character Recognition (OCR) to identify forms and tables without templates. The service provides APIs for document analysis, expense processing, and custom queries using deep-learning technology.

Main Benefits

Amazon Textract reduces manual data entry by processing millions of documents quickly with scalable analysis capabilities. Unlike basic OCR, it maintains relationships between data elements like key-value pairs and table structures. The service requires no ML expertise, offers pay-per-use pricing with no upfront commitments, and processes thousands of documents per hour.

Key Use Cases

- **Intelligent Search Indexes:** Create searchable document repositories from text detected in image and PDF files
- **Form Data Automation:** Extract structured data from forms and integrate into existing business workflows
- **Document Processing:** Accelerate data capture from financial documents, research reports, and medical notes
- **NLP Text Extraction:** Provide organised text grouping for natural language processing applications.

Core Features

- **Optical Character Recognition:** Automatically detects printed and handwritten text from documents, handling various fonts and styles
- **Form Extraction:** Detects key-value pairs in document images automatically while retaining context without manual intervention
- **Table Extraction:** Preserves data composition in tables during extraction, maintaining relationships between rows and columns
- **Query-based Extraction:** Allows data specification using natural language questions without knowing the document structure beforehand
- **Custom Queries:** Enables customisation of pretrained features using business-specific documents with as few as ten samples
- **Signature Detection:** Detects signatures on any document or image with location coordinates and confidence scores
- **Analyse Expense:** Extracts vendor names, invoice numbers, item prices from invoices and receipts with normalised keys
- **Analyse Lending:** Automates extraction from mortgage loan documents with prebuilt ML models for classification and splitting.

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/textract/

FAQ: https://aws.amazon.com/textract/faqs/

Technical Overview

Backup and Recovery Capabilities

Backup: Amazon Textract operates as a stateless processing service without traditional backup capabilities since it does not store customer data long-term. Customers must implement backup strategies for source documents in Amazon Simple Storage Service (Amazon S3) and extracted data in downstream systems using AWS Backup for Amazon S3 buckets and Amazon Relational Database Service databases storing extracted information.

Recovery: Multi-Region availability enables cross-Region disaster recovery strategies with automatic failover capabilities. AWS handles underlying infrastructure disaster recovery as a managed service. Supports customer RPO objectives through document repository replication and RTO objectives with sub-hour recovery by maintaining multi-Region document repositories and cross-Region replication for critical workflows.

Setup Process

Onboarding: Setup requires AWS account creation and AWS Identity and Access Management user configuration with appropriate Textract permissions. Install AWS CLI/SDKs with access credentials. Access through AWS Management Console for immediate testing or integrate via APIs with comprehensive documentation and code samples available.

Offboarding: Documents are not stored unless explicitly requested by customers. Delete stored data via AWS Management Console or AWS CLI. Custom Queries adapters removed using DeleteAdapter and DeleteAdapterVersion API operations, automatically cleaning all associated data and amazon resource names.

Data Management

- **Storage Options:** Stateless processing service that does not store customer data long-term unless explicitly configured by customers
- **Data Removal:** Delete stored data via AWS Management Console or CLI; Custom Queries adapters removed using API operations
- **Cross-Region Replication:** Multi-Region availability enables cross-Region strategies by processing documents in different geographic locations for redundancy
- **Encryption:** Integrates with AWS services supporting encryption; customers retain full ownership of content throughout processing operations.

Optimise Cost and Consumption

Document preprocessing removes unnecessary pages to reduce processing costs through the pay-per-page model. Volume-based pricing tiers provide cost reductions after processing one million pages monthly. Asynchronous batch operations optimise costs for large-volume processing scenarios. Confidence score thresholds minimise expensive manual review requirements. Custom Queries improve extraction accuracy, reducing downstream processing costs. Selecting appropriate API combinations (Forms, Tables, Queries) matches specific use cases efficiently.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/textract/latest/dg/what-is.html
Getting Started: https://docs.aws.amazon.com/textract/latest/dg/getting-started.html
API Reference: https://docs.aws.amazon.com/textract/latest/dg/API_Reference.html
Service Quotas: https://docs.aws.amazon.com/general/latest/gr/textract.html

Amazon Timestream for InfluxDB

General Information

Amazon Timestream for InfluxDB is a fully managed time-series database engine that makes it easy for application developers and DevOps teams to run InfluxDB databases on AWS for real-time time-series applications using open-source APIs. It handles provisioning, patching, backups, and software updates, allowing organisations to focus on building applications rather than managing database infrastructure.

Main Benefits

Amazon Timestream for InfluxDB delivers single-digit millisecond query response times with automatic scaling and fully managed operations. It provides compatibility with the popular open-source InfluxDB ecosystem (2.x APIs, Telegraf, client libraries), eliminating infrastructure management overhead while offering enterprise-grade availability, durability, and security through Multi-AZ deployments and AWS KMS encryption.

Key Use Cases

- **IoT Sensor Data Analysis**: Collecting and analysing high-volume time-series data from connected devices and sensors in real time
- **Application Performance Monitoring:** Tracking application metrics, response times, and system health with sub-second granularity
- **Infrastructure Monitoring:** Ingesting and querying metrics from servers, containers, and cloud resources using Telegraf and open-source tooling
- **Industrial Telemetry:** Managing operational data from manufacturing equipment with real-time dashboards and anomaly detection

Core Features

- **Open-Source Compatibility**: Full InfluxDB 2.x API support including line protocol writes, Flux queries, and InfluxQL for familiar developer workflows
- **Telegraf Integration**: Native support for 300+ Telegraf input plugins for broad data collection from infrastructure, applications, and IoT devices
- **Managed Operations**: Automated provisioning, patching, backups, and software updates with configurable maintenance windows
- **Multi-AZ High Availability**: Automatic replication across Availability Zones for enterprise-grade durability and failover
- **Flexible Instance Types**: Choice of database instance classes (db.influx) to match workload requirements from development to production scale
- **Always Encrypted**: Data encrypted at rest and in transit with AWS Key Management Service (AWS KMS) integration
- **Automated Backups**: Configurable backup retention with point-in-time recovery capabilities
- **Advanced Metrics:** Built-in monitoring with Amazon CloudWatch integration for operational visibility

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/timestream/

FAQ: https://aws.amazon.com/timestream/faqs/

Technical Overview

Backup and Recovery Capabilities

Backup: Automated daily backups with configurable retention periods. On-demand backups available via AWS Management Console, CLI, or APIs (CreateDbBackup). Internal snapshots taken during deletes and retained for 30 days. Cross-instance backup restoration supported.

Recovery: Restore from backup to a new DB instance or replace an existing instance using REPLACE_EXISTING restore mode. Multi-AZ deployment provides automatic failover for high availability. Internal backups retained for 24 hours for availability and durability support. Supports customer RPO objectives (minutes via automated backups) and RTO objectives (minutes for Multi-AZ failover, longer for full restoration from backup).

Setup Process

Onboarding: Create a DB instance through AWS Management Console, CLI, or APIs by selecting an instance class, configuring storage, and specifying VPC/security groups. Set master user credentials and organisation name during creation. Access via InfluxDB UI, Influx API, or Influx CLI using the auto-generated DNS endpoint. Connect existing Telegraf agents and InfluxDB client libraries using standard InfluxDB 2.x protocols.

Offboarding: Delete DB instances through AWS Management Console, CLI, or APIs. Snapshots are retained for 30 days following deletion to enable restoration if needed. Export data before deletion using InfluxDB CLI or API. Billing stops when the instance is terminated.

Data Management

- **Storage Options:** Influx IOPS Included storage with three pre-configured performance tiers optimised for different workload types, up to 16 TiB per instance
- **Data Removal:** Delete DB instances to permanently remove data; snapshots retained for 30 days post-deletion. Programmatic data deletion via Influx API for selective removal
- **Replication:** Multi-AZ replication for high availability within a Region
- **Encryption:** Data encrypted at rest and in transit with AWS KMS integration for customer-managed encryption keys

Optimise Cost and Consumption

Amazon Timestream for InfluxDB enables cost optimisation through right-sizing instance classes based on workload requirements — from development (db.influx.medium) to production scale (db.influx.24xlarge). Storage costs scale with usage on the Influx IOPS Included tier. Organisations can reduce costs by using appropriate retention policies to automatically expire old data, optimising query patterns, and selecting instance sizes that match actual throughput and cardinality requirements. No upfront commitment required — pay-as-you-go pricing with per-hour billing for DB instances and per-GB-month for storage.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/timestream/latest/developerguide/timestream-for-influxdb.html

Getting Started: https://docs.aws.amazon.com/timestream/latest/developerguide/timestream-for-influxdb-getting-started.html

API Reference: https://docs.aws.amazon.com/timestream-influxdb/latest/APIReference/Welcome.html

Amazon Transcribe

General Information

Amazon Transcribe is a fully managed automatic speech recognition service that converts speech to text using advanced ML models. It supports both real-time streaming and batch processing of audio files with features like speaker identification, multi-language support, and content filtering.

Main Benefits

Amazon Transcribe delivers superior accuracy across diverse demographics and audio conditions while reducing costs by 60% compared to manual transcription services. It supports over 100 languages, identifies up to 10 speakers, and provides real-time processing with automatic scaling and pay-per-second billing.

Key Use Cases

- **Contact Centre Analytics:** Transcribe customer service calls for quality monitoring and compliance insights
- **Media Content Captioning:** Create subtitles and closed captions for video content and live streams
- **Meeting Transcription:** Convert recorded meetings and conference calls into searchable text documents
- **Voice-Enabled Applications:** Add speech-to-text capabilities for voice commands and conversational interfaces.

Core Features

- **Batch and Streaming Transcription:** Supports real-time streaming and batch processing with WebSocket and HTTP/2 protocols
- **Multi-language Support:** Supports over 100 languages with automatic language identification for multi-lingual audio streams
- **Speaker Diarisation:** Identifies and separates speech from up to 10 different speakers in audio recordings
- **Custom Vocabularies:** Create custom vocabularies and domain-specific language models for specialised terminology accuracy improvement
- **Content Redaction:** Automatically identifies and redacts personally identifiable information (PII) and filters toxic content
- **Call Analytics:** Provides sentiment analysis, call summarisation, and conversation insights for contact centre applications
- **Medical Transcription:** HIPAA-eligible service variant designed specifically for medical and healthcare transcription use cases
- **Multiple Audio Formats:** Supports FLAC, WAV, MP3, and MP4 formats with recommended sample rates.

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/transcribe/

FAQ: https://aws.amazon.com/transcribe/faqs/

Technical Overview

Backup and Recovery Capabilities

Backup: Amazon Transcribe stores transcripts in customer-controlled Amazon Simple Storage Service (Amazon S3) buckets, inheriting S3's built-in durability. Backup protection can be implemented using Amazon S3 versioning, cross-Region replication, or AWS Backup integration for comprehensive data protection of audio files and generated transcripts.

Recovery: Multi-Region architecture enables cross-Region disaster recovery strategies with automatic failover capabilities. Multi-Region deployments with AWS Lambda functions support automated failover scenarios. Supports near-zero RPO for streaming transcriptions and low RTO through cross-Region failover strategies and backup processing capabilities in secondary Regions.

Setup Process

Onboarding: Requires AWS account, AWS Identity and Access Management credentials configuration, and Amazon S3 bucket setup for batch transcriptions. Begin immediately using AWS Management Console or install AWS CLI and SDKs for programmatic access. Multiple interfaces are available including AWS Management Console, REST APIs, and WebSocket streaming.

Offboarding: Use DeleteTranscriptionJob API or AWS Management Console interface to remove transcription jobs and transcripts. Batch deletion commands available for complete data removal. Job records automatically deleted after 90 days.

Data Management

- **Storage Options:** Transcripts stored in customer-controlled Amazon S3 buckets with built-in durability and backup capabilities.
- **Data Removal:** Complete data control via DeleteTranscriptionJob API and AWS Management Console with automatic 90-day retention.
- **Cross-Region Replication:** Supported through Amazon S3 cross-Region replication for disaster recovery and multi-Region deployments.
- **Encryption:** Built-in security features with encryption capabilities and PII redaction for data protection.

Optimise Cost and Consumption

Amazon Transcribe optimises costs through pay-per-second billing and audio quality adjustment features that allow lower sample rates for voice-only content. Batch processing provides cost advantages for non-real-time requirements, while custom language models and vocabulary filtering reduce processing time and improve accuracy to minimise re-processing costs. Job queueing capabilities enable efficient resource usage during peak periods, and volume discounts are available for large workloads.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/transcribe/latest/dg/what-is.html
Getting Started: https://docs.aws.amazon.com/transcribe/latest/dg/getting-started.html
API Reference: https://docs.aws.amazon.com/transcribe/latest/APIReference/
Service Quotas: https://docs.aws.amazon.com/general/latest/gr/transcribe.html

Amazon Virtual Private Cloud (VPC)

General Information

Amazon Virtual Private Cloud (VPC) enables users to launch AWS resources in a logically isolated virtual network within the AWS Cloud. It provides complete control over networking environment including IP address ranges, subnets, routing tables, and security settings.

Main Benefits

Amazon VPC offers control and security for cloud networking with logical isolation that makes unauthorised information transfer highly improbable. Core VPC functionality is of no cost, reducing traditional network hardware costs while delivering enterprise-grade features and automatic scaling across multiple Availability Zones (AZ).

Key Use Cases

- **Multi-Tier Applications:** Create production environments with web servers in public subnets and databases in private subnets
- **Private Server Environments:** Deploy Amazon Elastic Compute Cloud (Amazon EC2) instances in private subnets with Network Address Translation (NAT) gateways for secure outbound access
- **Test and Development:** Set up isolated development environments with customisable network configurations for testing
- **Hybrid Cloud Connectivity:** Connect on-premises networks to AWS via Virtual Private Network or AWS Direct Connect for integration.

Core Features

- **Subnets and IP Addressing:** Create public and private subnets with customisable IPv4 and IPv6 CIDR blocks
- **Security Groups and Network ACLs:** Instance-level and subnet-level security controls with stateful and stateless filtering for granular traffic management
- **Internet and NAT Gateways:** Provide internet connectivity for public subnets and secure outbound access for private subnets
- **VPC Peering and Transit Gateway:** Connect multiple VPCs within same AWS Region or across AWS Regions for scalable network architectures
- **VPC Flow Logs:** Monitor and log IP traffic going to and from network interfaces for security analysis
- **AWS PrivateLink and VPC Endpoints:** Access AWS services privately without internet gateways, maintaining traffic within AWS network backbone
- **Domain Name System (DNS) Resolution:** Built-in DNS services with customisable DHCP option sets for name resolution within VPC
- **Route Tables:** Define network traffic routing with up to 200 route tables per VPC and 500 routes.
- **VPC Lattice:** Simplify service-to-service connectivity, security, and monitoring across VPCs and accounts.
- **AWS Cloud WAN:** Build and manage global networks connecting cloud and on-premises resources.

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/vpc/

FAQ: https://aws.amazon.com/vpc/faqs/

Technical Overview

Backup and Recovery Capabilities

Backup: VPC configurations are automatically maintained and replicated by AWS across multiple Availability Zones. Infrastructure as Code (IaC) templates provide version-controlled backup through AWS CloudFormation or Terraform. VPC Flow Logs can be stored in Amazon Simple Storage Service (Amazon S3) or CloudWatch Logs with versioning and lifecycle policies.

Recovery: VPC configurations can be recreated in other AWS Regions using IaC tools with cross-Region peering and Transit Gateway connections. Multi-AZ deployments provide automatic failover within minutes. Supports customer RPO (near-zero through cross-Region architectures) and RTO (sub-minute to hours depending on complexity) objectives.

Setup Process

Onboarding: AWS offers a default VPC in each AWS Region that is immediately ready for use. Custom VPCs can be created through AWS Management Console with guided wizards, AWS CLI, or AWS CloudFormation templates for IaC deployments.

Offboarding: Systematically remove resources in correct order to avoid dependency conflicts. Terminate Amazon EC2 instances, databases, and gateways first, then clean up network interfaces and subnets. AWS provides automated tools for systematic resource identification and removal.

Data Management

- **Storage Options:** VPC Flow Logs can be stored in Amazon S3 or Amazon CloudWatch Logs with versioning and lifecycle policies
- **Data Removal:** Systematically remove resources in correct order: instances, databases, gateways, then network interfaces, subnets, and VPC
- **Cross-Region Replication:** VPC configurations can be replicated across AWS Regions using IaC tools like AWS CloudFormation or Terraform
- **Encryption:** VPC provides enterprise-grade encryption controls, traffic mirroring, and advanced connectivity options for comprehensive network security.

Optimise Cost and Consumption

VPC endpoints reduce data transfer charges by accessing AWS services privately without NAT gateways. Strategic resource placement within the same Availability Zone reduces cross-AZ data transfer fees. IPv6 adoption reduces reliance on costly public IPv4 addresses. VPC peering provides cost-effective connectivity for simple multi-VPC scenarios. Rightsizing NAT gateway instances based on actual throughput requirements optimises ongoing operational expenses.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/vpc/latest/userguide/what-is-amazon-vpc.html

Getting Started: https://docs.aws.amazon.com/vpc/latest/userguide/getting-started-with-amazon-vpc-using-the-aws-cli.html

API Reference: https://docs.aws.amazon.com/AWSEC2/latest/APIReference/OperationList- query-vpc.html

Service Quotas: https://docs.aws.amazon.com/vpc/latest/userguide/amazon-vpc-limits.html

Amazon Workspaces Services – Amazon WorkSpaces Applications

Feature of Amazon Workspaces (Formerly Amazon AppStream 2.0)

General Information

Amazon WorkSpaces Applications is a fully managed application streaming service that transforms desktop applications into cloud-based SaaS solutions without code changes. Users access applications through web browsers or native clients with responsive performance indistinguishable from local installations.

Main Benefits

Amazon WorkSpaces Applications reduces infrastructure procurement and maintenance overhead while delivering 4K resolution performance across any device. Multi-session capabilities reduce costs by 50% or more, supporting multiple users per instance. Consumption-based pricing with automatic scaling ensures optimal resource usage without upfront investment.

Key Use Cases

- **SaaS Transformation:** Transform desktop applications into cloud-based services without code modifications
- **Educational Access:** Provide students and educators on-demand access to specialised software applications
- **Remote Workforce:** Deliver applications to distributed teams across various locations and devices
- **Temporary Access:** Secure application access for contractors without requiring local installations.

Core Features

- **Multi-Device Access:** Stream applications to Windows, Mac, Chromebook, iPad, and Android devices via browsers or native clients
- **Fleet Types:** Three options: Always-On for constant pools, On-Demand provisioning, and Elastic pay-per-streaming models
- **Multi-Session Support:** Multiple user sessions on single Windows fleet instances to optimise resource usage and costs
- **4K Resolution:** Native client support with multi-monitor display, USB devices, local drives, and print redirection
- **Authentication Integration:** Supports User Pool, SAML 2.0 SSO, Microsoft Active Directory, and API-based streaming URLs
- **Auto Scaling:** Schedule-based and usage-based scaling policies with configurable minimum and maximum capacity settings
- **Real-Time Audio-Video:** Local webcam and microphone input support for video conferencing within streaming sessions
- **Enterprise Integration:** Deep AWS service integration including Amazon Simple Storage Service (Amazon S3), AWS Identity and Access Management (AWS IAM), Amazon Virtual Private Cloud (VPC), Amazon CloudWatch, and AWS Directory Service connectivity.

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/workspaces-family/workspaces/
FAQ: https://aws.amazon.com/workspaces-family/workspaces/faqs/

Technical Overview

Backup and Recovery Capabilities

Backup: Image snapshots for application environments and Amazon S3 for persistent data storage like home folders and application settings can be created manually (AWS Management Console or APIs). Custom backup solutions using Amazon S3 cross-Region replication must be implemented for data persistence and maintaining image catalogues across AWS Regions.

Recovery: Manual cross-Region replication of images and configurations, disaster recovery procedures using Infrastructure as Code tools like AWS CloudFormation or CDK to recreate environments in alternate AWS Regions. Supports customer RPO (depending on image update frequency and Amazon S3 replication schedules) and RTO (hours depending on implementation complexity) objectives.

Setup Process

Onboarding: Start with 'Try it Now' feature or Getting Started tutorial via AWS Management Console. Create stacks with user access policies, choose application images, and configure fleets with instance types and scaling settings. Service automatically creates AWS IAM roles with step-by-step guidance.

Offboarding: Delete stacks, fleets, and images through AWS Management Console or APIs, automatically terminating streaming instances. Remove application data by deleting associated Amazon S3 buckets and manually remove stored snapshots from connected storage services.

Data Management

- **Storage Options:** Amazon S3 for home folder persistence and application settings, with Amazon FSx integration for enterprise file sharing
- **Data Removal:** Delete stacks, fleets, and images via AWS Management Console or APIs, plus manual removal of Amazon S3 buckets and snapshots
- **Cross-Region Replication:** Custom implementation required using Amazon S3 cross-Region replication and manual image catalogue maintenance across Regions
- **Encryption:** Data remains secure in AWS Cloud with VPC network isolation and enterprise-grade security controls throughout streaming sessions.

Optimise Cost and Consumption

Multi-session capabilities reduce per-user costs by sharing Windows fleet instances among multiple users. Elastic fleet types charge only during active streaming sessions, while schedule-based and usage-based auto-scaling policies match capacity with actual demand. The Simple Pricing Tool and Scaling Optimiser Tool help estimate costs and optimise fleet capacity configurations for different instance types.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/appstream2/latest/developerguide/what-is-appstream.html
Getting Started: https://docs.aws.amazon.com/appstream2/latest/developerguide/what-is-how-to-start.html
API Reference: https://docs.aws.amazon.com/appstream2/latest/APIReference/Welcome.html
Service Quotas: https://docs.aws.amazon.com/appstream2/latest/developerguide/limits.html

Amazon WorkSpaces

Feature of Amazon Workspaces

General Information

Amazon WorkSpaces is a fully managed virtual desktop infrastructure (VDI) service that provisions cloud-based desktops running Windows or Linux operating systems. Users can securely access their virtual desktops from multiple devices and web browsers, reducing traditional hardware procurement and complex software deployment requirements.

Note: *If you use the Amazon WorkSpaces Applications User Pool feature, email addresses may be stored and processed in AWS Regions outside your selected region for the purpose of sending service notifications. This processing is subject to AWS's Data Processing Addendum.*

Main Benefits

Amazon WorkSpaces reduces virtual desktop deployment times by up to 90% while reducing hardware procurement and maintenance overhead. The fully managed service delivers up to 72% cost savings through Reserved Instances integration, automated cost optimisation tools, and flexible billing models that charge only for actual usage.

Key Use Cases

- **Remote and Hybrid Work:** Enable secure access to corporate applications and data from anywhere
- **Developer Workstations:** On-demand access to high-performance development environments for software teams
- **Contact Centres:** Secure virtual desktops for customer service representatives and support teams
- **Educational Environments:** Virtual desktops for educational institutions and training programme.

Core Features

- **Multi-Platform Access:** Access WorkSpaces from Windows, Mac, Linux, Chromebooks, iPads, Android devices, and web browsers
- **Flexible Hardware Bundles:** Multiple compute types from Value to GraphicsPro.g4dn with varying vCPU, memory, and storage configurations
- **Multiple Operating Systems:** Support for Windows, Amazon Linux, Ubuntu Linux, Rocky Linux, and Red Hat Enterprise Linux
- **Custom Images and Bundles:** Create and deploy custom desktop images with pre-installed applications and configurations
- **Multi-Factor Authentication:** Integration with RADIUS servers and SAML 2.0 identity providers for enhanced security
- **Data Encryption:** Encryption at rest and in transit using AWS Key Management Service (AWS KMS)
- **Self-Service Management:** Users can restart, rebuild, change compute type, and manage their WorkSpaces directly
- **Multi-Region Resilience:** Cross-Region redirection and standby Amazon WorkSpaces for disaster recovery capabilities.

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/workspaces/
FAQ: https://aws.amazon.com/workspaces/faqs/

Technical Overview

Backup and Recovery Capabilities

Backup: Automatic snapshots of root and user volumes stored in Amazon Simple Storage Service (Amazon S3) for durability are created after initial creation, during regular use, and after restoration events. Integration with AWS Backup provides centralised backup management across AWS services.

Recovery: Amazon WorkSpaces can be restored to previous snapshots when in AVAILABLE, ERROR, UNHEALTHY, or STOPPED states. Multi-Region Resilience enables standby WorkSpaces in secondary AWS Regions with cross-region redirection and domain name system health check failover. Supports customer RPO (near real-time data synchronisation) and RTO (standby activation within minutes) objectives.

Setup Process

Onboarding: Launch Amazon WorkSpaces through AWS Management Console by selecting bundle types, running modes, and directory settings. Configure Virtual Private Cloud (VPC) with proper security groups, directory service, and multi-factor authentication. Deploy in private subnets with CloudWatch monitoring.

Offboarding: Terminate Amazon WorkSpaces through Console, CLI, or API, permanently deleting all data. Separately deregister directory services and delete VPC components, security groups, and custom images for complete organisational data removal.

Data Management

- **Storage Options:** Persistent storage with automatic backup to Amazon S3 and configurable storage sizes
- **Data Removal:** Termination permanently deletes all root and user volume data; process is irreversible
- **Cross-Region Replication:** One-way data replication from primary to secondary Regions for disaster recovery
- **Encryption:** Data encryption at rest and in transit using AWS KMS.

Optimise Cost and Consumption

The Cost Optimizer solution automatically analyses usage patterns and converts Amazon WorkSpaces between hourly and monthly billing models to minimise costs. AutoStop mode automatically stops inactive WorkSpaces, charging only for active hours. Amazon Elastic Compute Cloud Reserved Instances and Compute Savings Plans integration provides up to 72% cost savings for predictable workloads. Bring Your Own License options reduces per-user licensing costs significantly while optimising pool sizes reduces infrastructure overhead.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/workspaces/latest/adminguide/
Getting Started: https://docs.aws.amazon.com/workspaces/latest/adminguide/amazon-workspaces.html
API Reference: https://docs.aws.amazon.com/workspaces/latest/api/
Service Quotas: https://docs.aws.amazon.com/workspaces/latest/adminguide/workspaces-limits.html

Amazon WorkSpaces Secure Browser

Feature of Amazon Workspaces

General Information

Amazon WorkSpaces Secure Browser is a fully managed, cloud-native browser service that enables secure access to private websites and SaaS applications through disposable containers. It streams encrypted pixels to users' browsers while keeping sessions isolated in AWS, reducing data exfiltration risks.

Main Benefits

Amazon WorkSpaces Secure Browser reduces browser attack surface through AWS isolation and pixel streaming technology, preventing data exfiltration without transmitting HTML or sensitive data to local devices. Includes up to 200 streaming hours monthly, requires no specialised client software or VPN infrastructure, and supports ephemeral sessions with automatic patching.

Key Use Cases

- **Customer Support:** Secure access to internal applications and customer data without exposing sensitive information
- **Bring Your Own Device Initiatives:** Enable employees to safely access company resources from personal devices without security risks
- **Secure Analytics:** Provide controlled access to sensitive data and analytics tools while maintaining data isolation
- **High-Security Networks:** Enable secure browsing for users in high-security environments without direct internet exposure.

Core Features

- **Ephemeral Sessions:** Each session launches with a fresh, fully patched Chrome browser that is wiped clean after use
- **Pixel Streaming:** Streams encrypted pixels from remote browser sessions without transmitting sensitive data to local devices
- **Enterprise Browser Policies:** Enforce URL allow/blocking, session-level controls for clipboard, file transfer, and printer access
- **IP Access Controls:** Restrict access to trusted networks or devices through IP-based access controls
- **SAML 2.0 Integration:** Single sign-on support with existing SAML 2.0 identity providers for authentication
- **Web Content Filtering:** Control and monitor web content access with granular policies and over 25 predefined categories
- **Session Monitoring:** Monitor sessions using Amazon CloudWatch and AWS CloudTrail with detailed logging capabilities
- **Multiple Instance Types:** Support for standard.regular, standard.large, and standard.xlarge instances for different performance needs.

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/workspaces/secure-browser/

FAQ: https://aws.amazon.com/workspaces/secure-browser/faqs/

Technical Overview

Backup and Recovery Capabilities

Backup: Traditional backup capabilities are not required as browser sessions are ephemeral and automatically wiped after use. Configuration data including portal settings, browser policies, and user preferences are stored in AWS managed services like Amazon DynamoDB and Amazon Simple Storage Service (Amazon S3) with built-in durability across multiple Availability Zones.

Recovery: Recovery is supported through multi-Region availability and stateless session architecture, enabling quick restoration in alternate AWS Regions. Sessions are recreated on-demand rather than requiring instance-level recovery. Supports minimal RPO (effectively zero for session data) and RTO (rapid failover capabilities) objectives through ephemeral session provisioning.

Setup Process

Onboarding: Setup involves creating a web portal through AWS console, configuring Virtual Private Cloud, subnets, and security groups for network connectivity. Administrators configure browser policies and federate existing SAML 2.0 identity providers for authentication.

Offboarding: Deactivate user access through SAML identity provider and delete web portals via AWS Management Console. All encrypted session data, browser policies, and user preferences are automatically removed by deleting associated resources.

Data Management

- **Storage Options:** Configuration data stored in AWS managed services like Amazon DynamoDB and Amazon S3 with built-in durability
- **Data Removal:** Ephemeral sessions automatically wiped after use; no HTML, DOM, or sensitive data stored locally
- **Cross-Region Replication:** Multi-Region availability enables quick restoration in alternate AWS Regions through stateless session architecture
- **Encryption:** Data protection ensured through AWS Key Management Service integration for encryption and secure data storage

Optimize Cost and Consumption

Multiple instance types enable workload-specific right-sizing: standard.regular for static content, standard.large for improved performance, and standard.xlarge for resource-intensive applications. Consumption-based pricing includes 200 streaming hours monthly per user with competitive additional hour rates. The ephemeral session model reduces persistent storage costs and reduces infrastructure overhead through automated capacity management.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/workspaces-web/latest/adminguide/index.html

Getting Started: https://docs.aws.amazon.com/workspaces-web/latest/adminguide/setting-up.html

API Reference: https://docs.aws.amazon.com/workspaces-web/latest/APIReference/Welcome.html

Service Quotas: https://docs.aws.amazon.com/workspaces-web/latest/adminguide/request-service-quota.html

AWS Agentic AI Services

General Information

AWS Agentic AI Services provide fully managed AI agent capabilities for building, deploying, and operating intelligent agents at scale. The service encompasses Amazon Bedrock Agents for custom development, Amazon Quick for enterprise productivity, Kiro for software development automation, AWS DevOps Agent for operational incident management, and Amazon Nova Act for browser-based UI workflow automation. These services use foundation models to break down complex tasks, orchestrate interactions between data sources and applications, and autonomously complete multi-step workflows while maintaining enterprise security, compliance, and governance controls.

Main Benefits

AWS Agentic AI Services deliver 99.99% availability SLA, eliminate infrastructure management overhead, and provide seamless AWS integration. Organizations benefit from rapid agent deployment, built-in workflow orchestration, comprehensive observability, and accurate permissions-aware responses with source citations. Specialized agents offer over 90% task reliability for browser automation, autonomous operational incident triage, and best-in-class security vulnerability detection across popular programming languages. Enterprise-grade security includes full visibility and control over agent activities, secure data access, and regulatory compliance.

Key Use Cases

- **Software Development Automation:** AI-powered coding assistance with agents that generate code, perform security scanning, create documentation, write unit tests, and conduct code reviews. Kiro agents detect hard-to-find vulnerabilities, provide automated remediations, and accelerate application upgrades from legacy Java versions. Agents work across IDEs, command-line interfaces, and CI/CD pipelines.
- **Operational Incident Management:** Autonomous triage and resolution of operational incidents through AWS DevOps Agent. The agent analyzes patterns across historical incidents, strengthens observability, optimizes infrastructure, enhances deployment pipelines, and improves application resilience across AWS accounts and regions.
- **Browser-Based Workflow Automation:** Reliable UI workflow automation through Amazon Nova Act for web QA testing, data entry, data extraction, and checkout flows. Agents achieve over 90% reliability executing complex multi-step browser tasks with human-in-the-loop escalation when needed.
- **Action-Enabled Chat Agents:** Real-time conversational AI with API capabilities for answering questions, providing summaries, and executing actions based on enterprise data. Ideal for IT help desks, HR support, and customer service with proper access controls and source attribution.
- **Enterprise Knowledge Management:** Intelligent search and retrieval across enterprise content with natural language understanding. Agents connect to 40+ data sources including Amazon S3, Microsoft SharePoint, and Salesforce, providing comprehensive answers with citations while respecting user permissions.
- **Multi-Agent Collaboration:** Coordination of specialized agents working together on complex tasks with built-in intent classification. Enables sophisticated workflows where different agents handle specific domains with unified observability and automatic orchestration.

Core Features

- **Foundation Model Integration:** Support for Amazon Nova (Premier, Pro, Lite, Micro, Act), Anthropic Claude, and other leading models through Amazon Bedrock. Flexible model selection optimized for different reasoning capabilities with multimodal processing including text, images, audio, and browser-based visual understanding.
- **Orchestration and Reasoning:** Built-in ReAct orchestration breaks down complex tasks into manageable steps. Supports multiple orchestration strategies, inline agents, and Amazon Bedrock Flows for visual workflow design. Amazon Nova Act uses visual grounding to interact with web interfaces through natural language commands.
- **Developer Tools and Agents:** Kiro agent capabilities include spec-driven development, security scanning, unit test generation, documentation, and application upgrades. CLI agents provide agentic coding with step-by-step reasoning, file operations, and AWS resource queries. Security scanning detects vulnerabilities across Python, Java, JavaScript with automated remediation.
- **DevOps and Operational Agents:** AWS DevOps Agent autonomously triages operational incidents and analyzes patterns across historical data. Operates with multi-region cross-account data collection for comprehensive incident analysis. Provides recommendations for observability, infrastructure optimization, deployment pipelines, and application resilience.
- **Browser Automation Capabilities:** Amazon Nova Act SDK enables reliable browser automation with natural language workflow creation, form filling, navigation, data extraction, and human-in-the-loop escalation. Powered by custom Amazon Nova 2 Lite model with over 90% task reliability.
- **Tools and Actions:** Comprehensive tool ecosystem including action groups for API integration, code interpretation, computer use through Nova Act, and Model Context Protocol integration. Native AWS service support through Lambda, Step Functions, and direct API calls with Playwright browser integration.
- **Memory Management:** Session handling maintains conversation context while long-term memory enables persistent storage of user preferences and historical context. Agents store domain-specific knowledge and build contextual understanding over time.
- **Knowledge Base Integration:** Seamless connection to Amazon Bedrock Knowledge Bases for retrieval-augmented generation. Supports vector search and semantic retrieval. Agents automatically query knowledge bases to supplement responses.
- **Security and Guardrails:** Built-in Amazon Bedrock Guardrails for content filtering, topic restrictions, and response validation. Automated abuse detection, encryption at rest and in transit, and IAM-based access controls. Security scanning outperforms leading tools on detection across popular programming languages.
- **Observability and Evaluation:** Comprehensive logging and tracing through Amazon CloudWatch with detailed metrics for agent performance, user interactions, and system health. End-to-end evaluation frameworks for testing responses and measuring accuracy. Session replay and live viewing for browser automation workflows.
- **Enterprise Integration:** Pre-built connectors for 40+ enterprise data sources including Microsoft 365, Google Workspace, Salesforce, ServiceNow, and Slack. Browser extensions and application integrations for Microsoft Teams, Outlook, and Word. Integration with Amazon Bedrock AgentCore Browser for secure cloud-based environments.

AWS Nova Act

Feature of AWS Agentic Services

General Information

Amazon Nova Act enables organizations to build, deploy, and manage reliable AI agents for automating browser-based UI workflows at enterprise scale. Nova Act completes repetitive workflows in the browser with high reliability, executes APIs or tools, and even escalates to your attention when appropriate without breaking the agentic flow.

Main Benefits

Amazon Nova Act delivers high reliability at enterprise scale, unlike traditional automation scripts that break with UI changes. The service accelerates deployment from weeks to hours and enables 5x faster shipping velocity through reliable browser automation.

Key Use Cases

- **Web QA Testing:** Automate end-to-end UI testing to accelerate release cycles with full user-journey validation
- **Data Entry and Form Filling:** Streamline repetitive form filling across disparate systems, replacing manual copy-paste operations
- **Data Extraction:** Extract and consolidate information from web sources automatically for competitive intelligence monitoring
- **E-commerce Automation:** Automate shopping, reservations, and bookings at scale across e-commerce and travel sites

Core Features

- **High Reliability:** Achieves task reliability at scale through reinforcement learning and browser interaction training
- **Natural Language + Python:** Allows workflow definition using both natural language instructions and Python code for flexibility
- **IDE Integration:** Nova Act IDE extension enables building and testing agents within Visual Studio Code, Cursor, and Kiro
- **Enterprise Scale Deployment:** Native integration with Amazon Bedrock AgentCore, CloudWatch for monitoring, and IAM for authentication
- **Human-in-the-Loop:** Capability to escalate tasks to human supervisors when appropriate without breaking the agentic flow
- **Strands Framework Integration:** Compatibility with Strands Agents framework for building complex multi-agent workflows at enterprise scale
- **Custom Nova 2 Lite Model:** Powered by purpose-built computer use model that adapts to interface variations
- **AWS Service Integration:** Deep integration with Amazon S3, Lambda functions, and Git-based CI/CD pipelines for operations

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/nova/act/

FAQ: https://aws.amazon.com/nova/act/faq

Technical Overview

Backup & Recovery Capabilities

Backup: Workflow definitions and agent configurations are backed up using Amazon S3 with versioning and cross-region replication. Backups can be created through standard AWS practices or automated via centralized management.

Setup Process

Onboarding: Three-step process: experiment in online playground at nova.amazon.com/act, refine scripts using Nova Act IDE extension for Visual Studio Code/Cursor/Kiro, then deploy to AWS production. Requires Nova Act API key and appropriate IAM permissions.

Offboarding: Standard AWS data handling practices with customer data control. Workflow definitions and execution logs removed through standard AWS resource deletion processes. Data lifecycle managed through S3 policies and AWS data retention controls.

Data Management

- **Storage Options:** Amazon S3 integration for data storage with policy control and versioning capabilities
- **Data Removal:** Standard AWS resource deletion processes with customer data control through AWS support channels
- **Cross-Region Replication:** Supported through S3 storage with cross-region replication for workflow definitions and configurations
- **Encryption:** Nova Act stores data at rest using Amazon DynamoDB and Amazon Simple Storage Service
- (Amazon S3).

Optimize Cost and Consumption

Amazon Nova Act optimizes costs through agent-hour billing that excludes Human-in-the-Loop wait times from charges. Multiple parallel agents enable precise cost allocation per workflow. The IDE extension enables local development and testing before production deployment, minimizing billable runtime. Serverless architecture through Amazon Bedrock AgentCore provides efficient resource utilization without infrastructure provisioning costs, while AWS Enterprise agreements offer volume discounts.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/nova-act/latest/userguide/index.html

Getting Started: https://docs.aws.amazon.com/nova-act/latest/userguide/getting-started.html

API Reference: https://docs.aws.amazon.com/service-authorization/latest/reference/list_amazonnovaact.html

Service Quotas: https://docs.aws.amazon.com/nova/latest/userguide/quotas.html

AWS DevOps Agent

Feature of AWS Agentic AI Services

General Information

AWS DevOps Agent is a frontier AI agent that autonomously resolves and prevents incidents across AWS, multicloud, and hybrid environments. It investigates issues immediately when alerts occur, correlating telemetry and deployment data to reduce MTTR from hours to minutes.

Main Benefits

AWS DevOps Agent dramatically reduces mean time to resolution from hours to minutes through autonomous incident investigation that begins immediately when alerts occur. It correlates data across multiple tools simultaneously, builds comprehensive application topology maps, and provides proactive recommendations to prevent future incidents.

Key Use Cases

- **Autonomous Incident Investigation:** Automatically investigates issues when alerts occur, analyzing error patterns and root causes
- **Proactive Incident Prevention:** Analyzes historical patterns to provide observability and infrastructure optimization recommendations
- **Multi-Tool Correlation:** Integrates with CloudWatch, Datadog, Dynatrace, New Relic, Splunk, and CI/CD systems
- **Application Resource Mapping:** Builds topology graphs of resources and relationships for architecture understanding

Core Features

- **Always-On Autonomous Response:** Automated incident investigation begins immediately when alerts occur with interactive chat interface
- **Application Resource Mapping:** Builds topology graphs of application resources and relationships to understand architecture during investigations
- **Multi-Tool Integration:** Works with observability tools, code repositories, CI/CD pipelines, and supports custom MCP server connections
- **Proactive Recommendations:** Delivers targeted improvements for observability, infrastructure optimization, and deployment pipeline enhancement based on patterns
- **Agent Spaces:** Logical containers that define access boundaries with AWS account isolation, user access controls, and data isolation
- **Cross-Platform Support:** Works across AWS, multicloud, and hybrid environments for comprehensive incident response and prevention capabilities
- **Detailed Mitigation Plans:** Provides comprehensive mitigation plans with specific actions for resolving identified incidents and system issues
- **Dual-Console Architecture:** Separates administrative management through AWS Console from operational use through dedicated Agent Space web apps

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/devops-agent/

FAQ: https://aws.amazon.com/devops-agent/faqs/

Technical Overview

Backup & Recovery Capabilities

Backup: AWS DevOps Agent does not provide traditional backup capabilities and does not integrate with AWS Backup service. The service is designed for real-time incident analysis and response rather than data backup. Investigation data and incident history are maintained within Agent Spaces but focus on operational intelligence rather than data protection.

Recovery: AWS DevOps Agent does not support disaster recovery functionality and does not integrate with AWS Elastic Disaster Recovery. The service helps customers meet RTO objectives by significantly reducing MTTR from hours to minutes through autonomous investigation but does not directly address RPO objectives as it focuses on operational intelligence rather than data recovery.

Setup Process

Onboarding: Create Agent Spaces through AWS Management Console, configure IAM roles for cross-account access, integrate observability tools like CloudWatch or Datadog, and set user permissions. Setup includes defining AWS account access and external tool connections.

Offboarding: Delete Agent Spaces from AWS Management Console to remove investigation data and incident history. Clean up by removing third-party tool connections and deleting created IAM roles for straightforward resource removal.

Data Management

- **Storage Options:** Investigation data, incident history, and recommendations stored within Agent Spaces with strict data isolation boundaries
- **Data Removal:** Delete Agent Spaces from AWS Management Console to remove all investigation data, incident history, and recommendations
- **Cross-Region Replication:** Currently limited to US East (N. Virginia) Region only during preview period with no cross-region capabilities
- **Encryption:** Uses AWS security infrastructure for data separation and maintains secure cross-account access through dedicated IAM roles

Optimize Cost and Consumption

Efficient Agent Space configuration optimizes resource allocation by defining precise access boundaries and limiting agent task hours per month during preview. Targeted tool integrations reduce operational overhead by connecting only necessary observability and CI/CD systems. The service's proactive recommendations deliver infrastructure optimization suggestions that reduce incident frequency and operational costs through strengthened system resilience and deployment pipeline enhancements.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/devopsagent/latest/userguide/

Getting Started: https://docs.aws.amazon.com/devopsagent/latest/userguide/getting-started-with-aws-devops-agent.htm I

Service Quotas: https://docs.aws.amazon.com/devopsagent/latest/userguide/about-aws-devops-agent-public-preview-pricing-and-limits.html

AWS Security Agent

Feature of AWS Agentic AI Services

General Information

AWS Security Agent is an AI-powered tool that proactively secures applications throughout the development lifecycle. It conducts automated security reviews and delivers context-aware penetration testing on demand, understanding application design, code, and security requirements.

Main Benefits

AWS Security Agent eliminates the security bottleneck that forces 81% of organizations to deploy vulnerable code to meet deadlines. It transforms penetration testing from weeks-long external vendor processes into instant on-demand capabilities, enabling security validation at development speed while preventing vulnerabilities early in the lifecycle.

Key Use Cases

- **Automated Security Reviews:** Continuously validates security requirements during development against organizational standards
- **Context-Aware Penetration Testing:** Executes on-demand testing with customized attack plans before production deployment
- **Design Security Validation:** Validates architectural documents against organizational security requirements during design phase
- **Code Security Review:** Automated code reviews identifying vulnerabilities in pull requests with GitHub integration

Core Features

- **Design Security Review:** Validates architectural documents against organizational security requirements during the design phase
- **Code Security Review:** Automated code reviews identifying vulnerabilities and compliance issues in pull requests with GitHub integration
- **On-Demand Penetration Testing:** Context-aware testing that creates customized attack plans and executes attack chains to discover vulnerabilities
- **Automated Code Remediation:** Submits pull requests with suggested fixes when security issues are detected during penetration testing
- **GitHub Enterprise Integration:** Native integration with GitHub Enterprise Cloud for private repositories and development workflows
- **Organizational Security Requirements:** Enforces custom security controls and policies that applications must follow during reviews
- **CloudTrail Integration:** Comprehensive API activity logging, auditing, and compliance tracking of all security operations
- **AI-Powered Analysis:** Leverages AWS generative AI capabilities for investigation summaries and security analysis while maintaining data privacy

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/security-agent/

FAQ: https://aws.amazon.com/security-agent/faqs/

Technical Overview

Backup & Recovery Capabilities

Backup: AWS Security Agent automatically maintains security testing artifacts, findings, and configuration data within AWS secure storage systems with standard AWS availability and durability guarantees. Integration with AWS CloudTrail provides comprehensive audit trail backup for all API activities and security operations, ensuring complete traceability and compliance logging.

Recovery: Recovery capabilities are inherent to the managed service design through AWS multi-Availability Zone architecture, ensuring service continuity without customer intervention required. Supports customer RPO objectives through continuous on-demand security validation and RTO objectives via instant security response capabilities and rapid remediation through automated code fixes.

Setup Process

Onboarding: Navigate to the AWS Security Agent console and create an agent space workspace for security operations. Configure user access, set up organizational security requirements, and install the GitHub app for GitHub Enterprise Cloud integration. Define penetration testing scope and establish automated workflows.

Offboarding: Delete agent spaces, remove GitHub integrations, and terminate security testing activities through the AWS Security Agent console. Standard AWS data retention and deletion policies apply to stored artifacts, findings, and configuration data.

Data Management

- **Storage Options:** AWS secure storage systems with standard availability and durability guarantees for artifacts and findings
- **Data Removal:** Delete agent spaces and GitHub integrations through console with standard AWS data retention policies
- **Cross-Region Replication:** Currently operates in US East Virginia with plans for additional regional expansion
- **Encryption:** All customer data remains within secure AWS environments and never used to train AI models

Optimize Cost and Consumption

AWS Security Agent optimizes costs through on-demand penetration testing that activates instantly without scheduling, reducing resource waste from periodic testing cycles. Automated code remediation minimizes developer time spent on security fixes by submitting targeted pull requests with suggested solutions. Context-aware testing creates customized attack plans based on application specifics, focusing resources on relevant vulnerabilities and avoiding unnecessary broad-spectrum scanning that consumes excessive compute resources.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/securityagent/latest/userguide/

Getting Started: https://docs.aws.amazon.com/securityagent/latest/userguide/setup-agent.html

Service Quotas: https://docs.aws.amazon.com/securityagent/latest/userguide/quotas.html

AWS Amplify

General Information

AWS Amplify is a comprehensive set of tools and services for building secure, scalable fullstack web and mobile applications. It provides frontend libraries, UI components, backend building tools, and hosting capabilities with minimal configuration. The platform supports multiple frameworks including React, Angular, Vue, Next.js, Flutter, and React Native.

Main Benefits

AWS Amplify delivers comprehensive fullstack development capabilities with type-safe TypeScript backend definition and automatic AWS resource provisioning. It reduces the complexity of managing multiple services individually while providing built-in security, scalability, and performance proven practices. The serverless architecture automatically scales and offers pay-per-use optimisation with up to $200 credits in the AWS Free Tier.

Key Use Cases

- **Fullstack Web Applications:** Build complete web apps with authentication, real-time data, and hosting
- **Mobile Applications:** Develop iOS, Android, Flutter, and React Native apps with cloud backend integration
- **Real-time Collaborative Applications:** Create messaging apps, project management tools, and chat applications with real-time synchronisation
- **E-commerce Platforms:** Develop shopping apps with user accounts, product catalogues, and payment processing.

Core Features

- **Authentication:** User sign-up, sign-in, and authorisation using Amazon Cognito with SAML, OIDC, and social providers
- **Real-time Data APIs:** GraphQL and REST APIs powered by AWS AppSync with real-time subscriptions and offline synchronisation
- **Storage:** File storage capabilities using Amazon Simple Storage Service (Amazon S3) with different access levels and content management features
- **Hosting:** Fully managed hosting for static sites and server-side rendered apps with Continuous Integration and Continuous Deployment and global Content Delivery Network (CDN)
- **DataStore:** Local data store with automatic cloud synchronisation for offline-first applications
- **Functions:** Serverless functions using AWS Lambda for custom business logic and API extensions
- **Analytics:** User analytics and app usage tracking for understanding user behaviour
- **Push Notifications:** Send targeted push notifications to mobile app users.

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/amplify/

FAQ: https://aws.amazon.com/amplify/faqs/

Technical Overview

Backup and Recovery Capabilities

Backup: Amazon DynamoDB tables support point-in-time recovery and automated backups with Amazon S3 storage offering cross-Region replication and versioning capabilities. Backups can be configured manually through AWS CDK constructs or automated via AWS Backup for centralised management. The CDK-based architecture enables deletion protection and backup policies for critical resources.

Recovery: Multi-Region deployment capabilities with automatic failover and real-time data synchronisation minimise data loss. Applications use AWS AppSync for data synchronisation and Amazon CloudFront for global content delivery, providing resilience against Regional failures. Supports customer RPO objectives through real-time synchronisation and RTO objectives through rapid recovery via global CDN distribution and serverless auto-scaling.

Setup Process

Onboarding: Install AWS Amplify CLI and create projects using framework-specific quickstart guides. Configure AWS credentials, define backends using TypeScript in amplify/backend/directory. Deploy using 'npx ampx sandbox' for development or 'npx ampx pipeline-deploy' for production. Amplify Studio provides visual interface for non-technical users.

Offboarding: Delete apps from Amplify console through App settings > General settings, removing all associated resources. Delete AWS CloudFormation stacks, remove Amazon S3 buckets, and cleanup custom resources. Process completes in less than two minutes through automated console interface.

Data Management

- **Storage Options:** Amazon S3 file storage with different access levels, Amazon DynamoDB NoSQL database, and DataStore with automatic cloud synchronisation
- **Data Removal:** Delete app from Amplify console removes all associated resources including hosting, authentication, APIs, and storage within two minutes
- **Cross-Region Replication:** Amazon S3 storage offers cross-Region replication and versioning capabilities with multi-Region deployment and automatic failover support
- **Encryption:** Built-in security proven practices with AWS service integration providing automatic encryption and reliability across all backend components.

Optimise Cost and Consumption

AWS Amplify's serverless architecture automatically scales resources based on demand, preventing over-provisioning during traffic spikes while maintaining performance. Sandbox environments can be paused to reduce development costs, while tiered pricing models provide volume discounts for builds and hosting. Efficient caching strategies optimise content delivery through global CDN integration. Appropriate Amazon S3 storage class selection reduces file storage costs, and AWS cost management tools integration enables detailed consumption tracking with optimisation recommendations.

Technical Resources

Developer Guide: https://docs.amplify.aws/

Getting Started: https://docs.amplify.aws/javascript/start/

API Reference: https://aws-amplify.github.io/amplify-js/api/

Service Quotas: https://docs.aws.amazon.com/general/latest/gr/amplify.html

AWS Amplify Console

Feature of AWS Amplify

General Information

AWS Amplify Console provides a unified management interface for building, deploying, and hosting modern web applications with Git-based CI/CD workflows. It combines automatic deployments from code repositories with global hosting on Amazon CloudFront's content delivery network for scalable web application management.

Main Benefits

AWS Amplify Console delivers serverless hosting that automatically scales to millions of users while reducing operational overhead of web server management. Features atomic deployments with zero maintenance windows, sub-second global failover through Amazon CloudFront distribution, and RTOs under 15 minutes, providing consumption-based cost efficiency.

Key Use Cases

- **Continuous Deployment:** Automatically deploy applications from Git repositories with integrated CI/CD pipelines
- **Static Website Hosting:** Host SPAs and static sites with global Amazon CloudFront distribution
- **Full-Stack Applications:** Deploy complete applications with frontend and backend using serverless infrastructure
- **Development Environment Management:** Create staging environments with pull request previews for testing.

Core Features

- **Git-based CI/CD:** Automated build and deployment pipelines triggered by commits from GitHub, GitLab, Bitbucket, CodeCommit
- **Atomic Deployments:** Zero-downtime deployments where sites are ready to view immediately after deployment completion
- **Global Hosting:** Deploy applications on Amazon CloudFront's global edge network for low-latency content delivery worldwide
- **Feature Branch Deployments:** Automatic deployment of feature branches for testing and collaboration before production merging
- **Pull Request Previews:** Generate preview environments for pull requests to test changes without affecting production
- **Custom Domains & SSL:** Support for custom domains with automatic SSL certificate provisioning and wildcard subdomains
- **Real-time Monitoring:** Built-in hosting metrics monitoring with Amazon CloudWatch integration and custom alarm configuration capabilities
- **AWS WAF Integration:** One-click web application firewall protection with managed rules for common exploits.

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/amplify/

FAQ: https://aws.amazon.com/amplify/faqs/

Technical Overview

Backup and Recovery Capabilities

Backup: Build artifacts and deployment history are automatically versioned and stored, enabling rollback to previous application versions through the console interface. Git repositories serve as the primary backup mechanism for source code. Static assets are automatically replicated across Amazon CloudFront edge locations globally.

Recovery: Application rollbacks can be completed in minutes through the console interface, rapid recovery via Git-based redeployment, and global distribution provides sub-second failover capabilities. Supports customer RPO (determined by Git commit frequency) and RTO (typically under 15 minutes) objectives.

Setup Process

Onboarding: Access through unified management interface across 19 AWS Regions. Connect Git repositories (GitHub, GitLab, Bitbucket, CodeCommit), configure build settings, and deploy applications. Amplify automatically detects frameworks and suggests build configurations with quickstart templates available.

Offboarding: Delete applications through Amplify console removing build artifacts and deployment history. Manually delete Git webhooks and custom domains before deletion. Automatic Amazon Simple Storage Service (Amazon S3) and Amazon CloudFront cleanup with permanent environment variable removal.

Data Management

- **Storage Options:** Amazon S3 for static assets and build artifacts with automatic versioning capabilities
- **Data Removal:** Delete applications through console removing build artifacts, deployment history, and configuration settings automatically
- **Cross-Region Replication:** Static assets automatically replicated across global CloudFront edge locations for content delivery
- **Encryption:** Automatic SSL certificate provisioning for custom domains with built-in HTTPS encryption support.

Optimise Cost and Consumption

Optimise build artifact sizes to reduce storage costs and build times to minimise compute charges. Use feature branch deployments judiciously to avoid unnecessary builds. Use Amazon CloudFront's global caching mechanisms to minimise data transfer costs. Environment-based pricing enables cost allocation across development, staging, and production environments while automatic scaling matches resource consumption to actual usage patterns.

Technical Resources

Developer Guide: https://docs.amplify.aws/

Getting Started: https://docs.amplify.aws/react/start/

API Reference: https://docs.aws.amazon.com/amplify/latest/APIReference/

Service Quotas: https://docs.aws.amazon.com/general/latest/gr/amplify.html

AWS AppFabric

General Information

AWS AppFabric is a fully managed service that connects SaaS applications to enhance security and productivity. It aggregates and normalises audit logs using Open Cybersecurity Schema Framework (OCSF) schema for centralised security monitoring. The service leverages generative AI to automate tasks across productivity applications.

Main Benefits

AWS AppFabric reduces custom integration development costs while delivering near real-time audit log normalisation every two minutes using OCSF schema. The service supports over 20 popular SaaS applications, accelerates incident response through enriched logs, and provides unique cross-application AI insights powered by Amazon Bedrock.

Key Use Cases

- **Security Observability Enhancement:** Aggregate and normalise audit logs for centralised monitoring and alerts
- **User Access Management:** Audit employee access across SaaS applications through single interface management
- **AI-Powered Productivity Enhancement:** Automate tasks and generate insights using Amazon Bedrock generative AI
- **No-Code SaaS Integration:** Connect over 20 applications without custom development or maintenance.

Core Features

- **Normalised Audit Logs:** Automatically normalises SaaS audit logs using OCSF into JSON or Apache Parquet formats
- **Enriched Audit Logs:** Enriches log data with user emails, IP addresses, and user agent strings
- **User Access Management:** Quick search and management of employee access across connected SaaS applications
- **Generative AI Assistant:** Amazon Bedrock-powered AI provides cross-application insights and task automation for productivity
- **No-Code Integration:** Connects multiple SaaS applications without requiring custom code development or maintenance
- **Data Encryption:** Encrypts data at rest using AWS KMS and in transit with TLS 1.2
- **Real-Time Processing:** Delivers normalised audit logs every 2 minutes for near real-time security monitoring
- **Multi-Application Support:** Supports over 20 popular SaaS applications including Microsoft 365, Google Workspace, Slack.

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/appfabric/

FAQ: https://aws.amazon.com/appfabric/faqs/

Technical Overview

Backup and Recovery Capabilities

Backup: AWS AppFabric does not provide native backup capabilities as it functions as a data processing service. Audit log data flows to customer-specified destinations (Amazon Simple Storage Service (Amazon S3) or Amazon Kinesis Data Firehose), where customers implement backup strategies using AWS Backup or S3 versioning and replication features. AWS AppFabric configurations should be documented for recreation.

Recovery: AWS AppFabric operates across multiple AWS Regions with AWS-managed high availability and does not support traditional disaster recovery. Customers implement recovery strategies for audit log destinations using cross-Region replication. Supports customer RPO objectives (continuous two-minute ingestion minimises data loss) and RTO objectives (underlying AWS infrastructure availability).

Setup Process

Onboarding: Setup begins through the AWS AppFabric console in supported AWS Regions by creating app bundles, configuring encryption with AWS-owned or customer-managed AWS Key Management Service (AWS KMS) keys, authorising applications with credentials, and setting up audit log ingestions to Amazon S3 or Amazon Kinesis Data Firehose destinations.

Offboarding: Delete app bundles to permanently remove metadata and automatically retire AWS KMS key grants. Separately delete audit log data from customer-owned Amazon S3 buckets, revoke application authorisations, and remove AppFabric access from SaaS applications. Intermediate Amazon S3 buckets auto-delete after 30 days.

Data Management

- **Storage Options:** Audit logs stored in customer-owned Amazon S3 buckets or streamed through Amazon Kinesis Data Firehose
- **Data Removal:** Delete app bundles to remove metadata; separately delete audit data per retention policies
- **Cross-Reigion Replication:** Customers implement cross-Region replication strategies for Amazon S3 and Kinesis Data Firehose destinations
- **Encryption:** Data encrypted at rest using AWS KMS keys and in transit with TLS 1.2.

Optimise Cost and Consumption

AWS AppFabric offers volume discounts for organisations exceeding 2,000 users, optimising per-user costs at scale. Organisations can strategically select applications to connect based on security priorities, controlling subscription scope. Data storage costs are managed through Amazon S3 lifecycle policies and choosing storage-efficient Apache Parquet over JSON formats. The user access feature uses per-search pricing rather than per-user charges, enabling precise usage-based cost control for access auditing activities.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/appfabric/latest/adminguide/what-is-appfabric.html
Getting Started: https://docs.aws.amazon.com/appfabric/latest/adminguide/getting-started-security.html
API Reference: https://docs.aws.amazon.com/appfabric/latest/api/Welcome.html
Service Quotas: https://docs.aws.amazon.com/appfabric/latest/adminguide/limits.html

AWS Artifact

General Information

AWS Artifact provides on-demand access to AWS security and compliance documents, including International Organization for Standardization (ISO), Payment Card Industry (PCI), and System and Organization Controls (SOC) reports. The service enables centralised agreement management across AWS accounts and offers self-service access to third-party vendor compliance documentation from AWS Marketplace.

Main Benefits

AWS Artifact reduces manual compliance workflows by providing self-service access to security documentation, reducing dependency on AWS Support during audits. All documents and agreements are provided at no cost, while watermarked reports and historical versions help ensure audit trail integrity and continuity during extended compliance processes.

Key Use Cases

- **Audit and Compliance:** Providing auditors with ISO, PCI, and SOC reports for regulatory demonstrations
- **Vendor Assessments:** Self-service access to compliance reports eliminating AWS Support contact during audits
- **Organisational Compliance Management:** Centralised agreement management across all AWS Organization member accounts
- **Third-Party Vendor Reports:** Access compliance documents for AWS Marketplace independent software vendors.

Core Features

- **On-demand Compliance Reports:** Download ISO, PCI, and SOC reports with unique watermarks for audit verification
- **Agreement Management:** Review, accept, and track AWS agreement status for individual and organisational accounts
- **Third-party Vendor Reports:** Access compliance documents for independent software vendors on AWS Marketplace
- **Historical Report Versions:** Access previous versions of compliance reports directly through console for audit continuity
- **AWS Organisations Integration:** Automatically accept security compliance reports on behalf of all member accounts
- **IAM Access Control:** Fine-grained permissions and service-linked roles for secure document access management
- **CloudTrail API Logging:** Comprehensive audit trails for all AWS Artifact access and management activities.

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/artifact/

FAQ: https://aws.amazon.com/artifact/faq/

Technical Overview

Backup and Recovery Capabilities

Backup: AWS Artifact does not provide traditional backup capabilities or integrate with AWS Backup, as it serves as a document access portal for compliance reports and agreements. The service maintains report availability and agreement status across multiple AWS Regions for service reliability.

Recovery: AWS Artifact does not support disaster recovery scenarios through AWS Elastic Disaster Recovery integration. Service continuity is maintained through AWS infrastructure resilience across AWS Regions with high availability access to compliance documentation. Supports compliance-related recovery objectives by helping ensure consistent access to required audit documentation, though it does not contribute to traditional RPO/RTO objectives for data recovery.

Setup Process

Onboarding: AWS Artifact requires only an AWS account for immediate access via the console. Configure AWS Identity and Access Management (AWS IAM) permissions for specific report types and agreement management functions. For AWS Organizations integration, enable trusted access and grant service-linked role permissions.

Offboarding: Terminate organisational agreements through the console interface. Delete locally downloaded compliance documents from user systems. The service maintains minimal customer data, focusing on document access rather than data storage.

Data Management

- **Storage Options:** Document access portal with no customer data storage; downloads to user systems
- **Data Removal:** Delete locally downloaded files; terminate organisational agreements through console interface
- **Cross-Region Replication:** Report availability and agreement status maintained across multiple AWS Regions automatically
- **Encryption:** Documents secured with unique watermarks; customer responsible for downloaded file security.

Optimise Cost and Consumption

AWS Artifact's self-service access model reduces operational overhead by automating compliance document retrieval, minimising manual administrative tasks during audit processes. Centralised agreement management through AWS Organizations enables economies of scale across multiple accounts, reducing legal and administrative costs through consolidated governance. Historical report version access maintains audit continuity without requiring duplicate procurement cycles, optimising resource use for extended compliance workflows.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/artifact/latest/ug/what-is-aws-artifact.html

Getting Started: https://docs.aws.amazon.com/artifact/latest/ug/getting-started.html

API Reference: https://docs.aws.amazon.com/artifact/latest/APIReference/

Service Quotas: https://docs.aws.amazon.com/general/latest/gr/aws-service- information.html

AWS Backup

General Information

AWS Backup is a fully managed service that centralises and automates data protection across AWS services, cloud, and on-premises environments. It consolidates backup management through policy-based automation, reducing custom scripts and manual processes for consistent, compliant backup strategies.

Main Benefits

AWS Backup delivers the only fully managed, centralised backup solution across over 20 AWS services, reducing complex custom scripts and multiple backup tools. Features include sub-second RPO continuous backups, 30% cost reduction through backup tiering, automated restore testing, and enterprise-grade write once, read many (WORM) compliance for comprehensive data protection.

Key Use Cases

- **Centralised Backup Management:** Consolidate backup policies and monitoring across multiple AWS services from single console
- **Regulatory Compliance:** Meet business and regulatory requirements with automated policies, audit reporting, and WORM capabilities
- **Disaster Recovery:** Implement cross-Region backup replication and automated recovery testing for business continuity
- **Ransomware Protection:** Use logically air-gapped vaults, vault lock, and malware scanning against cyber threats.

Core Features

- **Policy-Based Backup:** Create backup plans with automated scheduling, retention, and lifecycle management across AWS services
- **Cross-Region Replication:** Automatically replicate backups across AWS Regions and accounts for enhanced security and disaster recovery
- **Backup Vault Lock:** Write-once-read-many configuration prevents deletion or modification of backups for regulatory compliance
- **Incremental Backups:** Efficient storage with full initial backups followed by incremental changes only to reduce costs
- **Item-Level Restore:** Search and restore specific files or objects from Amazon Elastic Block Store snapshots and Amazon Simple Storage Service (Amazon S3) backups precisely
- **Backup Audit Manager:** Monitor compliance posture and generate reports for governance and regulatory requirements automatically
- **Malware Protection:** Integration with Amazon GuardDuty for automated malware scanning of recovery points and data protection
- **Restore Testing:** Automated validation of backup recoverability to meet RTO/RPO targets and help ensure business continuity.

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/backup/

FAQ: https://aws.amazon.com/backup/faqs/

Technical Overview

Backup and Recovery Capabilities

Backup: Incremental backups with encrypted storage and cross-Region replication can be created manually (AWS Management Console, API) or automated via policy-based backup plans with configurable scheduling frequencies. Features include automated lifecycle management, WORM compliance, and backup vault organisation for comprehensive data protection.

Recovery: Point-in-time recovery with full resource restoration, item-level recovery, and cross-account recovery scenarios. Automated restore testing validates backup recoverability and cross-Region backup copies enable faster regional recovery. Supports customer RPO (continuous sub-second to periodic schedules) and RTO (automated testing validation) objectives.

Setup Process

Onboarding: Setup via AWS Management Console requires service opt-in configuration and IAM role setup. Create backup plans with scheduling and retention policies, then assign resources using tags or explicit selection. Default backup plans accelerate implementation.

Offboarding: Remove AWS Backup configurations through console or API while preserving underlying service data. Automated lifecycle policies delete expired backups, with AWS Backup Vault Lock preventing deletion during compliance retention periods.

Data Management

- **Storage Options:** Incremental backups with automated lifecycle management transitioning to lower-cost cold storage and backup tiering
- **Data Removal:** Automated lifecycle policies delete expired backups with manual deletion via console or API available
- **Cross-Region Replication:** Automatically replicate backups across AWS Regions and accounts for enhanced security and disaster recovery
- **Encryption:** Encrypted storage using AWS Key Management Service integration with backup vault encryption for comprehensive data protection.

Optimise Cost and Consumption

AWS Backup provides automated lifecycle management that transitions backups to lower-cost cold storage, reducing long-term storage expenses. AWS Backup tiering for Amazon S3 delivers up to 30% cost reduction for extended retention periods. Incremental backup technology stores only changed data, minimising storage consumption. Cost allocation tags enable detailed expense tracking across backup policies, while automated deletion of expired backups prevents unnecessary storage accumulation.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/aws-backup/latest/devguide/whatisbackup.html

Getting Started: https://docs.aws.amazon.com/aws-backup/latest/devguide/getting-started.html

API Reference: https://docs.aws.amazon.com/aws-backup/latest/devguide/api-reference.html

Service Quotas: https://docs.aws.amazon.com/aws-backup/latest/devguide/aws-backup-limits.html

AWS Batch

General Information

AWS Batch is a fully managed batch computing service that automatically provisions compute resources and optimises workload distribution for containerised batch ML, simulation, and analytics workloads across, Amazon Elastic Container Registry (Amazon ECR), Amazon Elastic Kubernetes Service (Amazon EKS), AWS Fargate, and Amazon Elastic Compute Cloud (Amazon EC2) instances.

Main Benefits

AWS Batch reduces infrastructure management overhead while delivering up to 90% cost savings through automatic Spot Instance integration. The fully managed service provides automatic scaling, built-in fault tolerance, and AWS service integration, enabling organisations to focus on core work rather than infrastructure management.

Key Use Cases

- **High Performance Computing:** Scientific research, engineering simulations, and complex calculations at scale
- **Machine Learning and AI:** ML model training, hyperparameter tuning with SageMaker integration
- **Financial Services:** Post-trade analytics, risk analysis, and high-frequency trading simulations
- **Life Sciences:** Drug screening, DNA sequencing analysis, and genomics research workflows.

Core Features

- **Fully Managed Service:** Automatically provisions, scales, and terminates compute resources without manual infrastructure management
- **Multi-Compute Environment Support:** Supports Amazon EC2, AWS Fargate, Amazon EKS, and Spot Instances for flexible resource allocation
- **Job Queuing and Scheduling:** Priority-based scheduling with dependency management and automatic retry capabilities for failed jobs
- **Container Orchestration:** Runs Docker containers with support for multi-container and multi-node parallel jobs
- **Resource-Aware Scheduling:** Manages consumable resources like license tokens and bandwidth limits for optimal utilisation
- **SageMaker Integration:** Provides queuing and scheduling capabilities for SageMaker Training jobs with fair-share policies
- **Auto Scaling:** Dynamically scales compute resources up and down based on job queue demand
- **Cost Optimisation:** Supports Amazon EC2 Spot Instances for up to 90% cost savings with management tools.

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/batch/
FAQ: https://aws.amazon.com/batch/faqs/

Technical Overview

Backup and Recovery Capabilities

Backup: AWS Batch integrates with AWS Backup for Amazon Elastic Block Store (Amazon EBS) volumes and Amazon Elastic File System (Amazon EFS) file systems used by batch jobs. Job definitions and compute environments are automatically maintained with high availability. Container images support Amazon ECR replication features, while Amazon CloudWatch logs can be exported to Amazon Simple Storage Service (Amazon S3) for long-term retention.

Recovery: Multi-Region deployment enables disaster recovery through Infrastructure as Code (IaC) tools like AWS CloudFormation. AWS Batch integrates with AWS Elastic Disaster Recovery for comprehensive planning and supports automatic failover using Amazon Route 53 and Application Recovery Controller. Supports customer RPO (Amazon S3 versioning and cross-Region replication) and RTO (rapid resource provisioning and automated job restart) objectives.

Setup Process

Onboarding: Setup involves creating AWS Identity and Access Management roles, compute environments, job queues, and job definitions sequentially through wiZard-based or CLI tutorials. Configure Amazon Virtual Private Cloud networking, security groups, and container repositories. AWS provides comprehensive documentation and best practices guides.

Offboarding: Delete resources in reverse order: disable job queues, delete compute environments, deregister job definitions. Manually remove data from Amazon S3, Amazon EBS, Amazon EFS, and Amazon ECR containers. Amazon CloudWatch logs retained per policies.

Data Management

- **Storage Options:** Amazon S3 for data input/output, Amazon EFS for shared file systems, Amazon EBS for persistent storage
- **Data Removal:** Manual deletion required for Amazon S3 buckets, Amazon EBS volumes, Amazon EFS systems, and Amazon ECR container images
- **Cross-Region Replication:** Supported through Amazon S3 cross-Region replication, Amazon ECR replication, and IaC tools like AWS CloudFormation
- **Encryption:** Uses underlying AWS service encryption capabilities including Amazon S3, Amazon EBS, and Amazon EFS encryption features.

Optimise Cost and Consumption

AWS Batch optimises costs through automatic Amazon EC2 Spot Instance integration delivering up to 90% savings with automated bidding and interruption handling. Dynamic scaling helps ensure payment only when jobs run, preventing idle capacity costs. Fair-share scheduling and resource-aware scheduling optimise use while mixed instance allocation strategies balance performance and cost. Integration with Savings Plans and Reserved Instances supports predictable workloads.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/batch/latest/userguide/what-is-batch.html
Getting Started: https://docs.aws.amazon.com/batch/latest/userguide/Batch_GetStarted.html
API Reference: https://docs.aws.amazon.com/batch/latest/APIReference/
Service Quotas: https://docs.aws.amazon.com/batch/latest/userguide/service_limits.html

AWS Budgets

General Information

AWS Budgets is a cloud financial management service that enables users to set custom budgets for AWS costs and usage while tracking Reserved Instances and Savings Plans metrics. The service provides automated notifications and actions when spending approaches defined thresholds, offering proactive cost control and governance.

Main Benefits

AWS Budgets offers native integration with all AWS services, providing real-time cost and usage data updated up to three times daily. The service delivers proactive financial governance through forecasting alerts and automated actions, preventing cost overruns before they occur while optimising Reserved Instance and Savings Plans usage.

Key Use Cases

- **Cost Control and Monitoring:** Set spending limits and receive alerts to prevent unexpected charges
- **Usage Tracking and Governance:** Establish usage limits for services and get notified when approaching thresholds
- **Reserved Instances Optimisation:** Monitor Reserved Instance use and coverage to identify unused or under-utilised commitments
- **Savings Plans Management:** Track Savings Plans usage and coverage metrics to maximise cost savings.

Core Features

- **Cost Budgets:** Set spending limits for services and receive alerts when costs approach thresholds
- **Usage Budgets:** Establish usage limits for services and get notified when usage approaches set thresholds
- **Reserved Instance Budgets:** Define usage and coverage thresholds for Reserved Instances with automated alerts
- **Savings Plans Budgets:** Monitor Savings Plans usage and coverage metrics with threshold-based notifications
- **Budget Actions:** Configure automated actions like applying AWS Identity and Access Management (AWS IAM) policies or stopping Amazon Elastic Compute Cloud (Amazon EC2) instances when thresholds
- **Forecasting Alerts:** Receive notifications based on forecasted spending before actual costs are incurred
- **Custom Period Budgets:** Create budgets aligned with fiscal years, project durations, or grant periods
- **Multi-dimensional Filtering:** Apply filters based on services, AWS Regions, accounts, tags for granular budget tracking.

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/aws-cost-management/aws-budgets/
FAQ: https://aws.amazon.com/aws-cost-management/aws-budgets/faqs/

Technical Overview

Backup and Recovery Capabilities

Backup: AWS Budgets maintains budget configurations and historical data across multiple Availability Zones for high availability. AWS Budget history data can be exported manually as CSV files through the console or automated via AWS CloudFormation templates and API calls for programmatic backup of configurations.

Recovery: AWS Budget configurations can be quickly recreated using Infrastructure as Code templates or API automation, with automatic service restoration during regional issues. The service is inherently resilient across Availability Zones. Supports customer RPO through downloadable CSV exports and RTO through rapid configuration recreation via templates.

Setup Process

Onboarding: Access AWS Billing and Cost Management console and navigate to AWS Budgets page. Choose simplified templates or advanced customisation methods, configure notifications through email and Amazon Simple Notification Service (Amazon SNS) topics. Proven practices recommend starting with basic cost budgets.

Offboarding: Delete individual budgets through console or API, permanently removing budget history and notifications. Remove associated SNS topics and AWS IAM policies, disable budget actions to prevent automated resource management.

Data Management

- **Storage Options:** Configuration data and historical performance data maintained across multiple Availability Zones
- **Data Removal:** Delete budgets through console or API, permanently removing history, notifications, and actions
- **Cross-Region Replication:** Inherently resilient with configurations managed across multiple AWS Availability Zones automatically
- **Encryption:** Standard AWS encryption applies to budget configurations and financial tracking data.

Optimise Cost and Consumption

AWS Budgets enables proactive cost optimisation through automated budget actions that apply restrictive AWS IAM policies or stop Amazon EC2 instances when thresholds are exceeded. Reserved Instance and Savings Plans usage monitoring identifies underutilised commitments for reallocation, maximising discount benefits. Forecasting alerts prevent cost overruns by enabling spending pattern adjustments before limits are reached, while multi-dimensional filtering isolates specific cost drivers across services and teams for targeted optimisation efforts.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/cost-management/latest/userguide/budgets-managing-costs.html

Getting Started: https://docs.aws.amazon.com/cost-management/latest/userguide/budgets-managing-costs.html

API Reference: https://docs.aws.amazon.com/aws-cost-management/latest/APIReference/API_Operations_AWS_Budgets.html

Service Quotas: https://docs.aws.amazon.com/cost-management/latest/userguide/management-limits.html

AWS Certificate Manager

General Information

AWS Certificate Manager (ACM) simplifies creating, managing, and renewing SSL/TLS certificates for secure communications. The service provides free public certificates for integrated AWS services and automates certificate deployment, renewal, and lifecycle management across your infrastructure.

Main Benefits

ACM reduces certificate authority fees saving hundreds of dollars annually per certificate, while providing free public certificates for integrated AWS services. Automatic renewal after 11 months prevents service disruptions, and enterprise-grade security with AWS Key Management Service (AWS KMS) encryption ensures protected private key storage.

Key Use Cases

- **Website Security:** Deploy SSL/TLS certificates on load balancers, Amazon CloudFront, and API Gateway for HTTPS
- **Enterprise PKI:** Use private certificates for securing internal applications and services within corporate networks
- **Container Workloads:** Secure Amazon Elastic Kubernetes Service pods and terminate TLS at Kubernetes Ingress or load balancers
- **Nitro Enclaves:** Install certificates on Amazon Elastic Compute Cloud (Amazon EC2) instances for enhanced security applications requiring isolation

Core Features

- **Automatic Certificate Renewal:** ACM automatically renews certificates after 11 months of 13-month validity period
- **Domain Validation:** Supports Domain Name System (DNS), email, and HTTP validation methods to verify domain ownership before issuance
- **Exportable Public Certificates:** Allows exporting certificates for use on Amazon EC2 instances, containers, or on-premises hosts
- **AWS Service Integration:** Native integration with AWS Elastic Load Balancing, CloudFront, Amazon API Gateway, AWS Elastic Beanstalk, and AWS App Runner
- **Multiple Certificate Types:** Supports single domain, multi-domain, and wildcard certificates with RSA and ECDSA encryption
- **Private Certificate Authority:** Integration with AWS Private CA for issuing private certificates for internal applications
- **AWS KMS Encryption:** Enterprise-grade security with encrypted private key storage using AWS KMS
- **IPv4 and IPv6 Support:** Supports both IPv4 and IPv6 on public endpoints for comprehensive connectivity.

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/certificate-manager/

FAQ: https://aws.amazon.com/certificate-manager/faqs/

Technical Overview

Backup and Recovery Capabilities

Backup: ACM does not provide traditional backup capabilities as certificates are managed resources that can be re-provisioned. Certificate metadata is maintained by AWS and certificates can be re-issued upon request. For exportable certificates, customers can export and securely store certificates externally via AWS Management Console, CLI, or API.

Recovery: ACM supports rapid certificate provisioning with public certificates issued within seconds after domain validation. Certificates can be quickly re-issued or imported in alternate AWS Regions for disaster recovery. Built-in resilience through multiple Availability Zones provides fault tolerance. Supports minimal RPO (certificate metadata maintained by AWS) and fast RTO (seconds for certificate issuance) objectives.

Setup Process

Onboarding: Access ACM through AWS Management Console, CLI, or API. Request certificates by specifying domain names and choosing DNS, email, or HTTP validation methods. Amazon Route 53 enables automatic CNAME (Canonical Name) record creation for DNS validation. Supports AWS CloudFormation templates for Infrastructure as Code deployments.

Offboarding: Delete certificates via AWS Management Console, CLI, or API after removing them from integrated services first. Private keys are securely deleted using AWS KMS management. Certificate removal is immediate and irreversible from ACM services.

Data Management

- **Storage Options:** Managed certificate storage with built-in resilience across multiple Availability Zones within AWS Regions
- **Data Removal:** Certificates deleted immediately via AWS Management Console, CLI, or API with secure AWS KMS-managed private key deletion
- **Cross-Region Replication:** Not supported; certificates are Regional resources requiring separate provisioning in each Region
- **Encryption:** Enterprise-grade security with encrypted private key storage using AWS KMS protection.

Optimise Cost and Consumption

Wildcard certificates optimise costs by securing multiple subdomains under a single certificate, reducing total certificate requirements. Automatic renewal after 11 months prevents expensive service disruptions and manual certificate management overhead. The service includes 10,000 monthly certificate export API calls at no charge, enabling efficient programmatic certificate retrieval and deployment automation across workloads.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/acm/latest/userguide/acm-overview.html

Getting Started: https://docs.aws.amazon.com/acm/latest/userguide/gs.html

API Reference: https://docs.aws.amazon.com/acm/latest/APIReference/Welcome.html

Service Quotas: https://docs.aws.amazon.com/acm/latest/userguide/acm-limits.html

AWS Clean Rooms

General Information

AWS Clean Rooms is a secure collaboration workspace enabling organisations to analyse collective datasets with partners without revealing underlying data. Users can run SQL queries, PySpark jobs, and ML models while maintaining privacy through built-in controls and cryptographic computing.

Main Benefits

AWS Clean Rooms delivers IDC MarketScape-recognised leadership in data clean room technology, reducing data movement costs while enabling secure multi-party analytics. Organisations benefit from pay-per-query cost optimisation, shared compute expenses between partners, and comprehensive privacy controls including differential privacy and cryptographic computing.

Key Use Cases

- **Advertising and Marketing:** Improve advertiser reach and publisher monetisation through secure campaign data collaboration
- **Machine Learning and Analytics:** Train custom ML models for fraud detection, campaign optimisation, and medical research
- **Entity Resolution:** Match customer records using identity resolution partners like LiveRamp and Neustar for common IDs
- **Cross-industry Data Collaboration:** Enable secure multi-party analytics across healthcare, finance, and regulated industries.

Core Features

- **Multi-party Collaboration:** Create collaborations with multiple organisations and configure member roles for secure data sharing
- **SQL and PySpark Analytics:** Run SQL queries and PySpark jobs on collective datasets with configurable compute sizes
- **Privacy-Enhancing Controls:** Implement fine-grained analysis rules, differential privacy, and cryptographic computing to protect sensitive data
- **Cross-Cloud Support:** Collaborate with datasets from multiple clouds including AWS, Snowflake and Amazon Athena
- **ML Capabilities:** AWS Clean Rooms ML for custom modelling, lookalike modelling, and synthetic dataset generation
- **Entity Resolution:** AWS Entity Resolution integration for matching customer records across different data sources
- **Cross-Region Collaboration:** Support for collaboration with data sources stored in different AWS and Snowflake AWSRegions
- **Analysis Templates:** Pre-approved SQL and Python queries that can be reused across collaborations.

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/clean-rooms/

FAQ: https://aws.amazon.com/clean-rooms/faqs/

Technical Overview

Backup and Recovery Capabilities

Backup: AWS Clean Rooms uses underlying AWS service backup capabilities since data remains in original locations like Amazon Simple Storage Service (Amazon S3) and AWS Glue. Organisations can use AWS Backup for centralised management, Amazon S3 versioning and replication for data protection, and AWS CloudFormation templates to backup collaboration configurations and analysis templates.

Recovery: Recovery strategies use AWS CloudFormation for infrastructure recreation, Amazon S3 cross-Region replication for data availability, and stateless query execution for rapid recovery. Cross-Region collaboration support enables automated failover to alternate AWS Regions. Supports customer RPO (sub-second through cross-Region replication) and RTO (quick collaboration recreation without data movement) objectives.

Setup Process

Onboarding: Setup service roles using AWS Management Console, CLI, or AWS CloudFormation. Create collaborations, configure data sources by linking AWS Glue tables, Amazon S3, or Snowflake sources, and establish analysis rules with privacy controls. Guided workflows through AWS Clean Rooms Solutions simplify preparation.

Offboarding: Delete collaborations, memberships, and resources through AWS Management Console or API. Remove configured tables, analysis templates, and query results from Amazon S3. Original data sources remain in customer accounts helping ensure complete data sovereignty.

Data Management

- **Storage Options:** Supports AWS Glue tables, Amazon S3 data, Snowflake sources, and Amazon Athena without data movement
- **Data Removal:** Delete collaborations, memberships, configured tables, analysis templates, and query results through console or API
- **Cross-Region Replication:** Supports cross-Region collaboration with data sources stored in different AWS and Snowflake Regions
- **Encryption:** Integrates with AWS Key Management Service for encryption key management and cryptographic computing for data protection.

Optimise Cost and Consumption

AWS Clean Rooms offers table caching for complex queries to reduce processing costs and configurable compute sizes for SQL query right-sizing. Customisable Spark properties enable performance tuning, while advanced SQL configurations optimise execution time. The shared cost model allows collaboration creators to designate cost responsibility between partners and pay-per-query pricing with Clean Rooms Processing Units helps ensure consumption scales with actual usage.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/clean-rooms/latest/userguide/what-is.html

Getting Started: https://docs.aws.amazon.com/clean-rooms/latest/userguide/setting-up.html

API Reference: https://docs.aws.amazon.com/clean- rooms/latest/apireference/Welcome.html

Service Quotas: https://docs.aws.amazon.com/clean-rooms/latest/userguide/quotas.html

AWS Cloud Map

General Information

AWS Cloud Map is a fully managed service discovery solution that enables applications to automatically discover and connect to backend services using logical names. It serves as a centralised service registry for dynamic cloud resources like Amazon Elastic Compute Cloud (Amazon EC2) instances, Amazon Elastic Container Service (Amazon ECS) tasks, and other AWS services.

Main Benefits

AWS Cloud Map reduces operational overhead of maintaining custom service registries while providing fully managed, highly available service discovery. It supports 1,000 steady-state API calls per second, automatic health-aware routing, and dual discovery methods (Domain Name System (DNS) and API) with deep container integration for reduced manual effort.

Key Use Cases

- **Microservices Discovery:** Enable automatic discovery and connection in containerised environments using Amazon ECS and Amazon Elastic Kubernetes Service (Amazon EKS)
- **Dynamic Resource Mapping:** Map logical names to scaling cloud resources like load balancers and databases
- **Health-Aware Routing:** Automatically route traffic only to healthy service instances using integrated health checks
- **Cross-Account Discovery:** Share service registries across multiple AWS accounts for centralised multi-account management.

Core Features

- **Namespace Management:** Create HTTP, private DNS, and public DNS namespaces to organise services with different discovery methods
- **Service Registration:** Register and manage service instances with custom attributes using RegisterInstance API for filtering and discovery
- **Health Checking:** Automatic health monitoring using c Route 53 health checks or custom checkers to serve only healthy resources
- **Discovery Methods:** Support both DNS-based discovery and API-based discovery using DiscoverInstances for flexible service location options
- **Custom Attributes:** Assign custom metadata attributes to services and instances for advanced filtering and routing decisions
- **Cross-Account Sharing:** Share Cloud Map namespaces across AWS accounts using Resource Access Manager for centralised service registries
- **Container Integration:** Deep integration with Amazon ECS and Amazon EKS for automatic service registration and deregistration lifecycle management
- **Service-Level Attributes:** Store and access service information like traffic weights and configuration directly at the service level.

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/cloud-map/

FAQ: https://aws.amazon.com/cloud-map/faqs/

Technical Overview

Backup and Recovery Capabilities

Backup: AWS Cloud Map does not provide traditional backup capabilities as it serves as a dynamic service registry with ephemeral data. Service configurations can be recreated using Infrastructure as Code (IaC) via AWS CloudFormation or custom scripts, with service instances designed for automatic re-registration by applications during startup.

Recovery: Recovery involves automated service re-registration in recovery AWS Regions, with applications implementing logic to re-register services automatically. The service operates across multiple Availability Zones for high availability. Supports customer RPO objectives with effectively zero data loss and RTO objectives with recovery typically within minutes through fast DNS propagation.

Setup Process

Onboarding: Setup involves creating namespaces through AWS Management Console, CLI, SDKs, or REST API. Choose namespace type (HTTP, private DNS, or public DNS), create services with health check configurations, then register instances using RegisterInstance API.

Offboarding: Deregister service instances using DeregisterInstance API, delete services with DeleteService, then delete namespaces using DeleteNamespace. Process must follow reverse creation order to avoid dependencies. CLI commands automate bulk cleanup operations.

Data Management

- **Storage Options:** Dynamic service registry with ephemeral data designed for transient, automatically managed application instances
- **Data Removal:** Use DeregisterInstance API, DeleteService, then DeleteNamespace in reverse creation order for complete cleanup
- **Cross-Region Replication:** Not natively supported; requires separate deployments in each Region with IaC templates
- **Encryption:** Standard AWS encryption capabilities apply; integrates with AWS Identity and Access Management for fine-grained access control

Optimise Cost and Consumption

Choose HTTP namespaces for API-only discovery to avoid DNS-related charges when DNS functionality is not required. Optimise DNS TTL values to balance discovery freshness with query frequency costs. Implement connection pooling and caching mechanisms to reduce API call frequency. Deregister unused instances immediately to minimise registered resource charges. Use cross-account namespace sharing to consolidate registries and reduce duplicate costs across multi-account architectures. Service-level attributes provide metadata in single responses, reducing additional discovery calls.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/cloud-map/latest/dg/what-is-cloud-map.html

Getting Started: https://docs.aws.amazon.com/cloud-map/latest/dg/tutorials.html

API Reference: https://docs.aws.amazon.com/cloud-map/latest/api/Welcome.html

Service Quotas: https://docs.aws.amazon.com/cloud-map/latest/dg/cloud-map-limits.html

AWS CloudFormation

General Information

AWS CloudFormation helps model and set up AWS resources using Infrastructure as Code (IaC) through JSON or YAML templates. It provisions and configures resources automatically as managed collections called stacks, handling dependencies and providing change sets for deployment previews.

Main Benefits

AWS CloudFormation offers no additional service costs with native integration across all AWS services and comprehensive resource dependency management. It provides enterprise-grade capabilities including automatic rollback, drift detection for compliance, and StackSets for multi-account deployments, delivering rapid infrastructure recreation with stack deployment times typically measured in minutes.

Key Use Cases

- **IaC:** Define and version control AWS infrastructure using JSON or YAML templates
- **Multi-Region Deployment:** Replicate application infrastructure across multiple AWS Regions for disaster recovery
- **Environment Management:** Create and manage development, staging, and production environments with consistent configurations
- **DevOps Automation:** Integrate with Continuous Integration and Continuous Deployment pipelines to automate infrastructure provisioning and deployment workflows

Core Features

- **Template-based Infrastructure:** Define infrastructure using JSON or YAML templates with parameters, mappings, conditions, and outputs
- **Stack Management:** Deploy and manage resources as cohesive units called stacks with full lifecycle management capabilities
- **Change Sets:** Preview proposed changes before deployment to understand the impact of template modifications
- **Automatic Rollback:** Automatically rollback to previous working state if stack creation or updates fail
- **Drift Detection:** Detect when resources have been modified outside of AWS CloudFormation and remediate configuration drift
- **StackSets:** Deploy stacks across multiple AWS accounts and AWS Regions with centralised management and governance
- **Custom Resources:** Extend AWS CloudFormation capabilities with custom provisioning logic for third-party or custom AWS resources
- **Import Existing Resources:** Bring existing AWS resources under AWS CloudFormation management without recreating them.

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/cloudformation/
FAQ: https://aws.amazon.com/cloudformation/faqs/

Technical Overview

Backup and Recovery Capabilities

Backup: AWS CloudFormation integrates with AWS Backup to create composite recovery points for entire stacks, including templates and associated resources stored in Amazon Simple Storage Service (Amazon S3) with cross-Region replication. Templates can be versioned and stored in version control systems, with stack event history maintained for operational insights.

Recovery: Infrastructure recreation through automated template deployment with stack creation times typically measured in minutes, cross-Region StackSets for multi-Region deployment, and automatic rollback capabilities. Supports customer RPO objectives through point-in-time infrastructure snapshots and RTO objectives via rapid template-based infrastructure recreation.

Setup Process

Onboarding: Access AWS CloudFormation through AWS Management Console, CLI, or SDKs with sample templates and guided learning resources. Infrastructure Composer provides drag-and-drop interface for visual template creation. Begin with simple templates and progressively adopt advanced features like parameters and nested stacks.

Offboarding: Stack deletion automatically removes all resources in correct dependency order. Delete all stacks, remove templates from Amazon S3 buckets, and manually delete retained resources for complete cleanup.

Data Management

- **Storage Options:** Templates stored in Amazon S3 with versioning support and stack event history maintained for operational insights
- **Data Removal:** Stack deletion automatically removes resources in correct dependency order while respecting configured DeletionPolicy settings
- **Cross-Region Replication:** Templates support cross-Region replication in Amazon S3 for disaster recovery and multi-Region deployment scenarios
- **Encryption:** Integrates with AWS Key Management Service for encryption of underlying resources provisioned through AWS CloudFormation templates.

Optimise Cost and Consumption

AWS CloudFormation enables parameterised templates for environment-specific resource sizing, allowing development environments to use smaller, cost-effective configurations. Stack-level resource management facilitates automated cleanup preventing forgotten resource costs. Drift detection identifies unauthorised modifications that impact expenses. Template reuse across environments reduces development time while ensuring consistent cost-optimised configurations and supporting cost governance policy implementation.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/default.html
Getting Started: https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/GettingStarted.html
API Reference: https://docs.aws.amazon.com/AWSCloudFormation/latest/APIReference/
Service Quotas: https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/cloudformation-limits.html

AWS CloudHSM

General Information

AWS CloudHSM provides dedicated, FIPS 140-2 Level 3 validated hardware security modules for secure cryptographic operations and key management. The service offers single-tenant hardware security module (HSM) with complete customer control over encryption keys while delivering high availability clustering across multiple Availability Zones (AZs).

Main Benefits

AWS CloudHSM delivers single-tenant, FIPS 140-2 Level 3 validated HSMs with complete customer control over encryption keys. The service is designed to provide a 99.95% uptime SLA, up to 10,000 cryptographic operations per second, and reduce hardware procurement costs while helping ensure regulatory compliance and application integration.

Key Use Cases

- **Regulatory Compliance:** Meet Payment Card Industry Data Security Standard (PCI DSS), PCI Personal Identification Numbers (PINs), and PCI Security Standards Council (SSC) standards using dedicated HSMs
- **Database Encryption:** Enable Transparent Data Encryption by storing Transparent Data Encryption (TDE) master keys in HSMs
- **SSL/TLS Processing:** Offload SSL/TLS processing for web servers to reduce computational burden
- **Certificate Authority Operations:** Store and manage private keys for Public Key Infrastructure and signing.

Core Features

- **FIPS Validated HSMs:** FIPS 140-2 Level 3 and FIPS 140-3 Level 3 validated hardware security modules
- **Single-Tenant Architecture:** Dedicated HSM instances with complete isolation and full customer control over keys
- **High Availability Clustering:** Multi-AZ deployment with automatic synchronisation and 99.95% uptime SLA across multiple HSMs
- **Industry-Standard APIs:** Support for PKCS#11, JCE, Microsoft CryptoNG, and KSP APIs for integration
- **Scalable Performance:** Up to 10,000 cryptographic operations per second with support for 28 HSMs per cluster
- **Automatic Backup Management:** Daily encrypted backups with configurable retention policies and cross-Region backup copying
- **Client SDK Integration:** Comprehensive SDKs for Linux and Windows platforms with AWS CloudHSM CLI for management
- **Third-Party Integrations:** Native support for Oracle TDE, SSL offload, Windows CA, and HashiCorp Vault.

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/cloudhsm/

FAQ: https://aws.amazon.com/cloudhsm/faqs/

Technical Overview

Backup and Recovery Capabilities

Backup: Encrypted backups with secure AWS storage and durability can be created manually (AWS Management Console, CLI, API) or automated via daily scheduling with configurable retention policies. Cross-Region backup copying and on-demand deletion capabilities provide comprehensive backup management.

Recovery: Full cluster restoration from backups, cross-Region disaster recovery capabilities, and multi-AZ automatic failover between HSMs in different Availability Zones. Multi-AZ clustering provides fault tolerance with 99.95% uptime SLA. Supports customer RPO (daily automated backups with forced backup options) and RTO (near-zero within AWS Region through cluster redundancy) objectives.

Setup Process

Onboarding: Create Virtual Private Cloud and subnets across multiple Availability Zones, then create AWS CloudHSM cluster through AWS Management Console, CLI, or API. Configure security groups, install client software on Amazon Elastic Compute Cloud instances, and initialise cluster with self-signed certificate. Deploy multi-AZ for high availability.

Offboarding: Export required keys using supported APIs before cluster deletion. Stop client applications, remove HSMs from clusters via console or CLI. Customer data automatically destroyed upon HSM deletion with tamper-resistant hardware design.

Data Management

- **Storage Options:** Dedicated HSM hardware with up to 16,666 total keys and 3,333 asymmetric keys per instance
- **Data Removal:** Customer data automatically destroyed upon HSM deletion due to tamper-resistant hardware design
- **Cross-Region Replication:** Encrypted backups can be copied across AWS Regions for disaster recovery purposes
- **Encryption:** FIPS 140-2 Level 3 validated HSMs provide secure cryptographic operations and key storage.

Optimise Cost and Consumption

Dynamic HSM scaling enables adding or removing HSMs based on workload demands, preventing over-provisioning and optimising capacity use. Configurable backup retention policies automatically purge expired backups, reducing storage costs. On-demand backup deletion through AWS SDK and CLI provides granular backup cost control. Regional deployment flexibility allows selecting cost-effective AWS Regions while maintaining compliance requirements for optimised resource placement.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/cloudhsm/latest/userguide/introduction.html

Getting Started: https://docs.aws.amazon.com/cloudhsm/latest/userguide/getting-started.html

API Reference: https://docs.aws.amazon.com/cloudhsm/latest/APIReference/Welcome.html

Service Quotas: https://docs.aws.amazon.com/cloudhsm/latest/userguide/limits.html

AWS CloudTrail

General Information

AWS CloudTrail enables operational auditing, governance, and compliance by recording user, role, and service actions as events across AWS accounts. It captures API calls from the AWS Management Console, CLI, and SDKs, delivering comprehensive audit trails to Amazon Simple Storage Service (Amazon S3), Amazon CloudWatch Logs, or Amazon EventBridge.

Main Benefits

AWS CloudTrail delivers comprehensive native AWS audit logging without infrastructure management, offering 90 days free event history and up to 10-year retention. It provides automatic scaling, immutable records, real-time monitoring within 15 minutes, and integration with security services, requiring no agent installation.

Key Use Cases

- **Security Analysis:** Monitor suspicious activities, unauthorised access, and security breaches across AWS resources
- **Compliance Auditing:** Meet regulatory requirements by maintaining detailed logs of all account activities
- **Operational Troubleshooting:** Debug issues by analysing API call patterns and resource interactions
- **Anomaly Detection:** Identify unusual API call patterns and error rates using ML.

Core Features

- **Event History:** Always-on recording of 90 days of management events with search and download capabilities
- **AWS CloudTrail Lake:** Managed data lake with up to 10-year retention and SQL-based querying capabilities
- **Trails:** Configurable event delivery to Amazon S3 supporting management, data, and network activity events
- **AWS CloudTrail Insights:** ML-powered anomaly detection identifying unusual API call patterns and error rates
- **Multi-Region Support:** Centralised logging across all AWS Regions with consolidated event delivery and management
- **Multi-account Support:** Organisation-wide logging across multiple AWS accounts with centralised audit trail management
- **Event Delivery Options:** Multiple delivery destinations including Amazon S3 buckets, Amazon CloudWatch Logs, and Amazon EventBridge integration
- **Log File Integrity:** Digest file validation ensuring immutable audit records and tamper-proof event logging.

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/cloudtrail/

FAQ: https://aws.amazon.com/cloudtrail/faqs/

Technical Overview

Backup and Recovery Capabilities

Backup: AWS CloudTrail automatically stores immutable audit logs in Amazon S3 with high durability and replication across multiple Availability Zones. Event data stores provide automatic backup with configurable retention periods. Cross-Region replication can be configured manually through AWS Management Console, CLI, or APIs for additional protection. Log file integrity validation ensures audit trail authenticity.

Recovery: Multi-Region trail configuration and Amazon S3 cross-Region replication enable disaster recovery with automatic failover capabilities. Event data stores provide rapid access to audit logs during regional outages. Supports customer RPO objectives through real-time event logging within 15 minutes and RTO objectives via multi-Region access during disasters.

Setup Process

Onboarding: AWS CloudTrail automatically enables Event History for all AWS accounts at no cost. Advanced features require creating trails via AWS Management Console, CLI, or APIs with Amazon S3 bucket specification or AWS CloudTrail Lake event data stores with chosen retention periods.

Offboarding: Delete trails, event data stores, and channels through console or APIs. Event data stores enter 7-day PENDING_DELETION state. Manual deletion of Amazon S3 log files and Amazon CloudWatch Logs required for complete data removal.

Data Management

- **Storage Options:** Amazon S3 buckets, CloudWatch Logs, EventBridge, and CloudTrail Lake with 90-day to 10-year retention
- **Data Removal:** Delete trails, event data stores via console/APIs; manual Amazon S3 and Amazon CloudWatch Logs deletion required
- **Cross-Region Replication:** Supported through Amazon S3 cross-Region replication and multi-Region trail configuration with automatic failover
- **Encryption:** Automatic encryption with AWS Key Management Service for log files and event data stores integration.

Optimise Cost and Consumption

Advanced event selectors filter unnecessary events to reduce logging volume and charges. AWS CloudTrail Lake offers one-year extendable retention for variable workloads with flexible pricing tiers based on ingestion patterns. Organisation trails centralise multi-account logging while Amazon S3 lifecycle policies optimise long-term storage costs. Data event aggregation minimises high-volume environment expenses through selective resource logging and appropriate retention periods.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/awscloudtrail/latest/userguide/cloudtrail-user-guide.html

Getting Started: https://docs.aws.amazon.com/awscloudtrail/latest/userguide/cloudtrail-user-guide.html

API Reference:
https://docs.aws.amazon.com/awscloudtrail/latest/APIReference/Welcome.html

Service Quotas:
https://docs.aws.amazon.com/awscloudtrail/latest/userguide/WhatIsCloudTrail-Limits.html

AWS CodeBuild

General Information

AWS CodeBuild is a fully managed build service that compiles source code, runs tests, and produces deployment-ready artifacts. It reduces build server management while providing prepackaged environments for popular programming languages and automatic scaling based on demand.

Main Benefits

AWS CodeBuild reduces build server management overhead. Pay-per-use pricing avoids idle infrastructure costs, while parallel test execution and persistent caching dramatically reduce build times and enhance performance.

Key Use Cases

- **Continuous Integration:** Automatically compile source code, run tests, and validate code quality in Continuous Integration and Continuous Deployment pipelines
- **Automated Testing:** Execute unit, functional, and integration tests with visual reporting of test case results
- **Multi-Environment Builds:** Build applications across different programming languages using preconfigured or custom Docker environments
- **Container Image Building:** Build and publish Docker images to Amazon Elastic Container Registry (Amazon ECR) with persistent caching capabilities.

Core Features

- **Fully Managed Build Service:** Reduces need to set up, patch, update, and manage build servers with automatic scaling
- **Multiple Source Providers:** Supports GitHub, GitHub Enterprise, Bitbucket, AWS CodeCommit, and Amazon Simple Storage Service (Amazon S3) as source providers
- **Preconfigured Build Environments:** Provides ready-to-use environments for popular programming languages and build tools like Maven and Gradle
- **Custom Docker Environments:** Allows bringing customised build environments as Docker containers for specialised build requirements
- **Parallel Test Execution:** Enables running test suites concurrently to reduce build times significantly and improve performance
- **Test Reporting:** Generates comprehensive reports for unit, functional, and integration tests with visual results and analytics
- **VPC Integration:** Can operate within Amazon Virtual Private Cloud (Amazon VPC) for secure integration with private services and databases
- **Docker Server Capability:** Provides dedicated Docker server with persistent cache for faster container builds and deployments

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/codebuild/

FAQ: https://aws.amazon.com/codebuild/faqs/

Technical Overview

Backup and Recovery Capabilities

Backup: Build project configurations are automatically replicated across AWS infrastructure through standard resilience. Source code repositories, build artifacts in Amazon S3, and custom Docker images in Amazon ECR require separate backup through respective services. Build specifications should be version-controlled using Infrastructure as Code (IaC) practices with AWS CloudFormation or CDK for automated project recreation.

Recovery: Build projects can be instantly recreated from IaC templates, cross-Region replication via AWS CloudFormation templates, and geographically distributed Amazon S3 storage for artifacts. Multi-AZ resilience provides automatic infrastructure failure handling. Supports customer RPO objectives through multi-Region repositories and Amazon S3 cross-Region replication and RTO objectives with near-zero project recreation timeframes.

Setup Process

Onboarding: Create build projects through AWS Management Console, CLI, or SDKs defining source code location, build environment, and output settings. Configure buildspec YAML files for build processes with environment settings, artifact outputs, and logging options. Step-by-step tutorials available for common scenarios.

Offboarding: Delete build projects, builds, and associated resources using AWS Management Console, CLI, or SDKs with batch-delete-builds API for multiple builds. Separately manage Amazon S3 artifacts, Amazon CloudWatch logs, test reports, AWS Identity and Access Management roles, and Amazon CloudWatch log groups for complete data removal.

Data Management

- **Storage Options:** Build artifacts stored in Amazon S3 buckets, Amazon CloudWatch logs for build history and monitoring
- **Data Removal:** Delete projects via console, CLI, or SDKs; separately manage Amazon S3 artifacts and Amazon CloudWatch logs
- **Cross-Region Replication:** Supported through Amazon S3 cross-Region replication for critical build artifacts and geographically distributed storage
- **Encryption:** Integration with AWS Secrets Manager and Systems Manager Parameter Store for secure environment variable storage.

Optimise Cost and Consumption

Reserved capacity instances provide cost savings for predictable workloads with minimum usage commitments. Parallel test execution reduces total build time, directly decreasing consumed build minutes. Build dependency caching improves performance across builds while reducing processing time. Docker server capability with persistent caching significantly reduces build times for container-based workflows. Strategic spot pricing availability for non-critical builds and proper compute type selection optimise resource right-sizing expenses.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/codebuild/latest/userguide/welcome.html
Getting Started: https://docs.aws.amazon.com/codebuild/latest/userguide/getting-started-overview.html
API Reference: https://docs.aws.amazon.com/codebuild/latest/APIReference/Welcome.html
Service Quotas: https://docs.aws.amazon.com/codebuild/latest/userguide/limits.html

AWS CodeCommit

General Information

AWS CodeCommit is a fully managed source control service that provides secure, highly scalable private Git repositories in the cloud. It eliminates infrastructure management while offering standard Git functionality and seamless integration with existing Git-based tools for team collaboration.

Main Benefits

CodeCommit eliminates infrastructure management overhead while providing unlimited private repositories, enterprise-grade security with AWS KMS encryption and FIPS 140-2 compliance, and seamless integration with AWS DevOps services. The pay-per-active-user model with permanent free tier for five users ensures cost-effective scaling.

Key Use Cases

- **Source Code Management:** Store and version control application source code with full Git functionality
- **CI/CD Pipeline Integration:** Serve as source repository for automated build, test, and deployment workflows
- **Collaborative Development:** Enable team collaboration through pull requests, code reviews, and branching capabilities
- **Enterprise Compliance:** Meet regulatory requirements with encryption, audit trails, and IAM-based access controls

Core Features

- **Private Git Repositories:** Unlimited private repositories with full Git functionality including branching, merging, and tagging
- **Pull Requests and Code Reviews:** Built-in workflows with approval rules and team collaboration capabilities
- **Encryption and Security:** Data encrypted in transit and at rest using AWS KMS with IAM access controls
- **High Availability and Durability:** Repositories stored across multiple Availability Zones using Amazon S3 and DynamoDB
- **Event Notifications and Triggers:** Repository notifications via Amazon SNS and configurable triggers for AWS Lambda functions
- **Cross-Account Access:** Support for cross-account repository sharing and federated access using IAM roles and policies
- **Multiple Connection Methods:** Access via HTTPS credentials, SSH keys, git-remote-codecommit, and AWS CLI credential helper
- **AWS Service Integration:** Native integration with CodePipeline, CodeBuild, and other AWS development services

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/codecommit/

FAQ: https://aws.amazon.com/codecommit/faqs/

Technical Overview

Backup & Recovery Capabilities

Backup: Repository data stored redundantly in Amazon S3 and DynamoDB across multiple Availability Zones can be created manually (Git clone operations) or automated via AWS Backup service with centralized snapshot management. Git's distributed nature provides inherent backup capabilities through multiple repository copies.

Recovery: Repository data can be quickly restored from local clones or backup snapshots, disaster recovery through multi-AZ architecture, and point-in-time recovery capabilities. Automatic replication provides fault tolerance. Supports customer RPO (near-zero through multi-AZ replication) and RTO (rapid recovery from backup sources) objectives.

Setup Process

Onboarding: Setup requires AWS credentials and IAM permissions configuration via Management Console, AWS CLI, or APIs. Multiple connection methods include HTTPS credentials, SSH keys, git-remote-codecommit, or CLI credential helper. Comprehensive tutorials guide repository creation and team access configuration.

Offboarding: Repository deletion through Management Console, CLI, or API operations. Backup data using Git clone commands before deletion permanently removes all commit history and metadata. Supports programmatic deletion for bulk migration scenarios.

Data Management

- **Storage Options:** Repository data stored redundantly in Amazon S3 and DynamoDB across multiple Availability Zones
- **Data Removal:** Repository deletion permanently removes all commit history, branches, tags, and metadata from CodeCommit storage
- **Cross-Region Replication:** Manual implementation using Git's distributed nature by maintaining repository mirrors across regions
- **Encryption:** Data encrypted in transit via HTTPS/SSH and at rest using AWS KMS with FIPS 140-2 compliance

Optimize Cost and Consumption

CodeCommit's pay-per-active-user model charges only for users performing Git operations monthly, enabling cost control through service account consolidation to reduce active user counts. Organizations can leverage unlimited private repositories to avoid per-repository fees while timing development activities using batch operations to minimize active user months and maximize storage and request quota allowances.

Technical Resources

Developer Guide:
https://docs.aws.amazon.com/codecommit/latest/userguide/welcome.html

Getting Started: https://docs.aws.amazon.com/codecommit/latest/userguide/getting-started.html

API Reference:
https://docs.aws.amazon.com/codecommit/latest/APIReference/Welcome.html

Service Quotas: https://docs.aws.amazon.com/codecommit/latest/userguide/limits.html

AWS CodePipeline

General Information

AWS CodePipeline is a fully managed continuous delivery service that automates software release processes from source code changes to production deployment. It orchestrates Continuous Integration and Continuous Deployment (CI/CD) workflows through configurable stages and actions, integrating natively with AWS services and third-party tools for reliable application delivery.

Main Benefits

AWS CodePipeline reduces custom deployment infrastructure maintenance while providing enterprise-grade features like automatic rollbacks and cross-Region deployment. It reduces time-to-market through native AWS integration, supports up to 1,000 pipelines per AWS Region, and offers cost-effective pay-per-use pricing that reduces operational overhead and third-party licensing costs.

Key Use Cases

- **CI/CD Automation:** Automate build, test, and deployment workflows across development, staging, and production environments
- **Multi-stage Application Deployment:** Deploy applications to Amazon Elastic Compute Cloud, AWS Elastic Beanstalk, Amazon Elastic Container Service (Amazon ECS) containers, and AWS Lambda functions automatically
- **Infrastructure as Code:** Automate AWS resource deployment using AWS CloudFormation templates with version control capabilities
- **Cross-Region Deployment:** Deploy applications across multiple AWS Regions with automated artifact replication and progression.

Core Features

- **Visual Pipeline Modelling:** Model release workflows using visual pipelines with multiple configurable stages and actions
- **AWS Service Integrations:** Native integration with AWS CodeCommit, AWS CodeBuild, AWS Elastic Beanstalk, Amazon ECS, AWS Lambda, and AWS CloudFormation
- **Stage-level Conditions:** Configure entry, success, and failure conditions with automatic rollback capabilities for enhanced release control
- **Pipeline Variables:** Use pipeline-level and action-level variables for dynamic configuration and parameterised deployments
- **Parallel and Sequential Execution:** Execute actions in parallel within stages or sequentially across stages with flexible execution modes
- **Cross-Region Support:** Deploy applications across multiple AWS Regions using cross-Region actions and artifact replication
- **Third-party Integrations:** Integration with GitHub, Jenkins, and other third-party tools through prebuilt and custom plugins
- **Manual Approval Gates:** Include manual approval actions to control progression between pipeline stages for controlled releases.

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/codepipeline/

FAQ: https://aws.amazon.com/codepipeline/faqs/

Technical Overview

Backup and Recovery Capabilities

Backup: AWS CodePipeline lacks traditional backup capabilities as it orchestrates workflows rather than storing persistent data. Pipeline configurations can be exported and recreated using AWS CloudFormation templates or AWS CLI for disaster recovery. The service automatically maintains pipeline execution history for 12 months.

Recovery: Pipeline configurations can be restored via AWS CloudFormation templates or cross-Region replication, automated rollback through stage-level conditions, and cross-Region deployment capabilities for geographic distribution. Supports faster RTO through automated deployment processes and reduced RPO via quick deployment rollbacks.

Setup Process

Onboarding: Create AWS account with AWS CodePipeline permissions and use AWS Management Console with pre-built pipeline templates for Build, Automation, and Deployment use cases. Select source repository, configure build and deployment stages using simplified getting started experience.

Offboarding: Delete all pipelines to remove execution history and metadata. Manually delete Amazon Simple Storage Service (Amazon S3) artifacts, remove custom actions and webhooks via console or API, and clean up AWS Identity and Access Management (AWS IAM) roles and policies.

Data Management

- **Storage Options:** Pipeline artifacts stored in Amazon S3 buckets with 7GB limit for S3 sources, 1GB for CodeCommit/GitHub
- **Data Removal:** Delete pipelines to remove execution history automatically; manually delete Amazon S3 artifacts and custom actions via console/API
- **Cross-Region Replication:** Supports cross-Region actions and artifact replication for deploying applications across multiple AWS Regions
- **Encryption:** Integrates natively with AWS security features and AWS IAM for access control and compliance capabilities

Optimise Cost and Consumption

AWS CodePipeline offers dual pricing models with V1 flat-rate pipelines for consistent workloads and V2 usage-based pipelines for variable execution patterns. Efficient triggering mechanisms minimise unnecessary pipeline executions while consolidating similar workflows reduces management overhead. Optimising build times and implementing efficient Amazon S3 artifact storage practices further control costs. Pipeline variables enable dynamic configuration reducing duplicate pipeline creation for similar deployment scenarios.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/codepipeline/latest/userguide/welcome.html

Getting Started: https://docs.aws.amazon.com/codepipeline/latest/userguide/getting-started-codepipeline.html

API Reference: https://docs.aws.amazon.com/codepipeline/latest/APIReference/Welcome.html

Service Quotas: https://docs.aws.amazon.com/general/latest/gr/codepipeline.html

AWS Compute Optimizer

General Information

AWS Compute Optimizer is a ML-powered service that analyses AWS resource configuration and usage metrics to provide optimisation recommendations. It evaluates Amazon Elastic Compute Cloud (Amazon EC2) instances, Auto Scaling groups, Amazon Elastic Block Store (Amazon EBS) volumes, AWS Lambda functions, and databases to identify rightsizing opportunities and reduce costs while improving performance.

Main Benefits

AWS Compute Optimizer delivers measurable cost savings and performance improvements through ML-powered recommendations based on actual usage patterns rather than theoretical capacity. The service offers 93-day lookback analysis for seasonal patterns, automated rightsizing with performance risk assessment, and reversible actions for operational safety across organisations.

Key Use Cases

- **Cost Optimisation:** Identify over-provisioned and idle resources to reduce infrastructure costs through rightsizing
- **Performance Optimisation:** Recommend instance types and configurations that better match workload requirements
- **Resource Rightsizing:** Analyse usage patterns across Amazon EC2, AWS Lambda, Amazon EBS, and databases for optimal sizing
- **License Optimisation:** Optimise commercial software licenses like Microsoft SQL Server to reduce licensing costs.

Core Features

- **Resource Analysis:** Analyses Amazon EC2, Auto Scaling groups, Amazon EBS, AWS Lambda, Amazon Elastic Container Service, and Amazon Relational Database Service/Amazon Aurora using ML
- **Automation Rules:** Creates automated optimisation implementation based on AWS Region and resource tags on schedules
- **Enhanced Infrastructure Metrics:** Extends analysis lookback period from 14 days to 93 days for quarterly patterns
- **Idle Resource Detection:** Identifies unused or underused resources like unattached Amazon EBS volumes for cleanup
- **Performance Risk Assessment:** Evaluates performance risk of changes with ratings from very low to very high
- **Savings Estimation:** Calculates potential cost savings and considers pricing discounts when generating optimisation recommendations
- **External Metrics Integration:** Ingests metrics from external observability products to enhance recommendation accuracy and precision
- **Organisation Integration:** Integrates with AWS Organizations for centralised management and delegated administration across multiple accounts.

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/compute-optimizer/

FAQ: https://aws.amazon.com/compute-optimizer/faqs/

Technical Overview

Backup and Recovery Capabilities

Backup: AWS Compute Optimizer does not provide traditional backup capabilities as it is an optimisation service rather than a data storage service. However, automation features include safety mechanisms like creating Amazon EBS snapshots before volume modifications to enable data recovery if needed.

Recovery: AWS Compute Optimizer does not directly support disaster recovery scenarios or integrate with AWS Elastic Disaster Recovery. The service itself is regionally distributed for high availability but focuses on cost and performance recommendations rather than business continuity. Does not directly help meet RPO and RTO objectives.

Setup Process

Onboarding: Opt into AWS Compute Optimizer through AWS Management Console, CLI, or API, which automatically creates service-linked roles. Configure enhanced infrastructure metrics for 93-day analysis and set up automation rules for organisation-wide optimisation.

Offboarding: Opt out through AWS Management Console, CLI, or API to stop analysis and recommendations. Contact AWS Support to delete historical data and remove service-linked roles if no longer needed.

Data Management

- **Storage Options:** Stores Amazon CloudWatch metrics and configuration data for analysis, with 14-93 day retention periods
- **Data Removal:** Contact AWS Support to delete historical metrics and configuration data after opting out
- **Cross-Region Replication:** Service is regionally distributed for high availability but does not replicate customer data
- **Encryption:** Standard AWS encryption applies to stored recommendation data and Amazon CloudWatch metrics integration.

Optimize Cost and Consumption

AWS Compute Optimizer provides ML-driven rightsizing recommendations based on actual usage patterns to reduce over-provisioned resources. Automated idle resource detection identifies unattached Amazon EBS volumes and unused Amazon EC2 instances for cleanup. Enhanced 93-day metrics analysis captures seasonal patterns for accurate quarterly optimisation. Commercial software license optimisation reduces Microsoft SQL Server costs, while automation rules implement recommendations at scale with performance risk assessment ensuring operational safety.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/compute-optimizer/latest/ug/what-is-compute-optimizer.html

Getting Started: https://docs.aws.amazon.com/compute-optimizer/latest/ug/getting- started.html

API Reference: https://docs.aws.amazon.com/compute- optimizer/latest/APIReference/Welcome.html

Service Quotas: https://docs.aws.amazon.com/general/latest/gr/compute-optimizer.html

AWS Config

General Information

AWS Config is a fully managed service that continuously monitors and records AWS resource configurations, providing detailed visibility into configuration changes and resource relationships. It enables organisations to assess compliance, audit configurations, and evaluate resources against desired states using automated rules and conformance packs.

Main Benefits

AWS Config delivers native integration with over 150 AWS services, providing sub-second configuration change detection without performance impact. It offers comprehensive managed rules for industry standards like Health Insurance Portability and Accountability Act and Payment Card Industry Data Security Standard, scales automatically across 10,000 accounts per aggregator, and reduces compliance overhead through automated remediation capabilities.

Key Use Cases

- **Compliance Monitoring:** Continuously monitor configurations against industry standards using rules and conformance packs
- **Configuration Management:** Track and audit configuration changes to maintain visibility and troubleshoot issues
- **Security Governance:** Assess security postures and enable automated remediation of non-compliant resources
- **Multi-Account Governance:** Centralise configuration data across multiple accounts, AWS Regions, and organisations using aggregators.

Core Features

- **Configuration Recording:** Records configuration changes for AWS resources as configuration items with point-in-time snapshots
- **AWS Config Rules:** Evaluates resource compliance against desired configurations using managed or custom rules automatically
- **Conformance Packs:** Bundles multiple rules and remediation actions for compliance against industry standards
- **Multi-Account Aggregation:** Aggregates configuration and compliance data across multiple AWS accounts, AWS Regions, and organisations
- **Advanced Query:** Provides SQL-like queries to search and analyse current configuration state across resources
- **Automated Remediation:** Automatically fixes non-compliant resources using AWS Systems Manager documents or AWS Lambda functions
- **Resource Relationships:** Maps and tracks relationships between AWS resources to understand dependencies and impact
- **Configuration History:** Maintains detailed configuration history for up to seven years with customisable retention periods.

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/config/

FAQ: https://aws.amazon.com/config/faqs/

Technical Overview

Backup and Recovery Capabilities

Backup: AWS Config maintains historical configuration data in customer-owned Amazon Simple Storage Service (Amazon S3) buckets with built-in redundancy and customisable retention periods up to seven years. Configuration snapshots can be generated on-demand via AWS Management Console and stored automatically through continuous or daily recording strategies. Cross-Region replication is supported through Amazon S3 bucket replication.

Recovery: Point-in-time configuration recovery enabled through Advanced Query API and aggregators for rapid assessment during incidents. Cross-Region Config aggregator and rule replication available for disaster recovery scenarios. Supports customer RPO objectives (detailed seven-year configuration history) and RTO objectives (sub-second configuration change detection).

Setup Process

Onboarding: Start with 1-click setup in AWS Management Console, automatically configuring resource recording, delivery channels, and AWS Identity and Access Management roles. Configure recording strategies (continuous or daily), select resource types, and set up Amazon S3 buckets and Amazon Simple Notification Service (SNS) notifications. Multi-account deployment is available via AWS CloudFormation, Terraform, or Systems Manager Quick Setup.

Offboarding: Stop and delete configuration recorders to cease data collection. Delete Config rules, conformance packs, aggregators, and delivery channels. Manually delete historical data from customer-owned S3 buckets and configure retention periods between 30 days to 7 years.

Data Management

- **Storage Options:** Configuration data stored in customer-owned Amazon S3 buckets with built-in redundancy and durability
- **Data Removal:** Configurable retention periods from 30 days to 7 years; historical Amazon S3 data requires manual deletion
- **Cross-Region Replication:** Supported through Amazon S3 bucket replication for configuration data and Config aggregator cross-Region deployment
- **Encryption:** Inherits Amazon S3 encryption capabilities for configuration items stored in customer-owned Amazon S3 buckets.

Optimise Cost and Consumption

AWS Config offers selective resource recording to monitor only critical resources, reducing evaluation costs. Periodic recording modes enable daily monitoring for static resources while reserving continuous recording for frequently-changing assets. Configurable retention periods from 30 days to 7 years balance compliance requirements with storage expenses. Multi-account aggregation achieves economies of scale by reducing redundant rule evaluations across accounts, while managed rules reduce AWS Lambda execution costs compared to custom rule implementations.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/config/latest/developerguide/WhatIsConfig.html
Getting Started: https://docs.aws.amazon.com/config/latest/developerguide/getting-started.html
API Reference: https://docs.aws.amazon.com/config/latest/APIReference/Welcome.html
Service Quotas: https://docs.aws.amazon.com/config/latest/developerguide/configlimits.html

AWS Control Tower

General Information

AWS Control Tower provides a straightforward way to set up and govern secure, compliant multi-account AWS environments following prescriptive best practices. It orchestrates AWS Organizations, AWS Service Catalog, and AWS Identity and Access Management (AWS IAM) Identity Center to deploy landing zones rapidly. The service implements preventive, detective, and proactive controls with automated Account Factory provisioning capabilities.

Main Benefits

AWS Control Tower delivers secure, compliant multi-account environments in under 30 minutes with over 750 managed controls and automated governance at scale. The no cost orchestration service reduces manual configuration overhead, reduces operational complexity, and provides centralised management while supporting up to 10,000 accounts per landing zone with enterprise-grade compliance capabilities.

Key Use Cases

- **Multi-Account Environment Setup:** Establish secure, compliant landing zones with standardised governance controls
- **Account Factory Provisioning:** Enable teams to rapidly create AWS accounts using configurable templates
- **Regulatory Compliance Management:** Automate security guardrails and policy enforcement across organisational units
- **Enterprise Governance at Scale:** Centralised oversight with preventive, detective, and proactive controls.

Core Features

- **Landing Zone:** Well-architected multi-account environment based on security and compliance proven practices that scales to enterprise needs
- **Controls (Guardrails):** High-level rules providing ongoing governance with preventive, detective, and proactive control types for compliance management
- **Account Factory:** Configurable account template that standardises provisioning of new accounts with pre-approved configurations and automated workflows
- **Dashboard:** Provides continuous oversight allowing administrators to see provisioned accounts, enabled controls, and noncompliant resources organised
- **Control Catalog:** Over 750 managed controls with support for various compliance frameworks and configurable exemptions
- **Automated Account Enrolment:** Automatically enrols accounts and applies appropriate baselines and controls when accounts are moved between OUs
- **Service Integration:** Deep integration with AWS Organisations, AWS Service Catalog, AWS IAM Identity Center, AWS Config, Amazon CloudTrail, and AWS Security Hub
- **Centralised Governance:** Reduces manual configuration overhead while helping ensure consistent security posture across multi-account environments at scale.

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/controltower/
FAQ: https://aws.amazon.com/controltower/faqs/

Technical Overview

Backup and Recovery Capabilities

Backup: AWS Control Tower integrates with AWS Backup to provide automated, centralised backup capabilities with prescriptive backup plans at the landing zone level. Features include four types of backup plans with configurable retention periods, automated resource tagging for backup inclusion, and centralised backup vault creation with multi-Region AWS Key Management Service (AWS KMS) key encryption across organisational units.

Recovery: Disaster recovery supports backup and restore, pilot light, warm standby, or multi-site active/active strategies through underlying AWS services. Multi-Region governance capabilities enable consistent security controls and account structures across AWS Regions. Supports customer RPO objectives through automated backup plans and RTO objectives via standardised disaster recovery implementations.

Setup Process

Onboarding: Customers select their home AWS Region and complete automated pre-launch checks with sufficient service limits and proper AWS IAM permissions. Choose between full landing zone deployment or controls-dedicated experience. Quick start creates audit and log archive accounts in approximately 30 minutes.

Offboarding: Requires careful decommissioning by disabling controls, unregistering organisational units, and removing landing zone components. Organisations must manually delete Amazon S3 buckets and AWS CloudFormation stacks for complete removal.

Data Management

- **Storage Options:** Uses Amazon S3 for logging storage, AWS Config for configuration data, and AWS CloudTrail for audit trails
- **Data Removal:** Manual deletion required for Amazon S3 buckets containing logs and AWS CloudFormation stacks to ensure complete removal
- **Cross-Region Replication:** Multi-Region governance capabilities support consistent security controls and account structures across AWS Regions for resilience
- **Encryption:** Multi-Region AWS KMS keys provide encryption for backup functionality and centralised backup vault creation across organisational units

Optimise Cost and Consumption

AWS Control Tower provides selective detective controls configuration to prevent duplication across AWS Config items and AWS CloudTrail data events. Organisational-level CloudTrail trail configurations optimise logging charges while proper tagging strategies enable accurate cost attribution across accounts and organisational units. Centralised governance through shared security infrastructure and standardised resource provisioning prevents over-provisioning. Account Factory automation reduces operational overhead costs by streamlining baseline application processes.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/controltower/latest/userguide/what-is-control-tower.html
Getting Started: https://docs.aws.amazon.com/controltower/latest/userguide/getting-started-with-control-tower.html
API Reference: https://docs.aws.amazon.com/controltower/latest/APIReference/Welcome.html
Service Quotas: https://docs.aws.amazon.com/controltower/latest/userguide/limits.html

AWS Cost Explorer

General Information

AWS Cost Explorer is a comprehensive cost management and analysis tool that provides interactive dashboards, customisable reports, and AI-powered forecasting capabilities. It enables organisations to visualise AWS spending patterns and identify optimisation opportunities across services, accounts, and AWS Regions.

Main Benefits

AWS Cost Explorer delivers up to 75% compute cost reductions through automated Reserved Instance recommendations and savings via rightsizing analysis. Features 18-month AI-powered forecasting, 13 months historical data, and native AWS integration, enabling cost reductions for well-managed accounts.

Key Use Cases

- **Cost Analysis and Monitoring:** Track AWS spending patterns across services, accounts, and AWS Regions
- **Budget Planning and Forecasting:** Generate 18-month cost forecasts and create budgets with automated alerts
- **Cost Optimisation:** Identify underutilised resources and receive Reserved Instance recommendations for savings
- **Resource Rightsizing:** Analyse usage patterns to optimise Amazon Elastic Compute Cloud (Amazon EC2) instances and reduce waste

Core Features

- **Interactive Cost Visualisation:** Customisable charts and graphs for exploring cost and usage data with filtering options
- **18-Month Forecasting:** AI-powered cost forecasting with explainable insights based on historical usage patterns
- **Cost Comparison:** Automated cost comparison between time periods with detailed breakdown of cost drivers
- **Granular Data Access:** Daily, hourly, and resource-level cost data with support for multi-year historical analysis
- **Reserved Instance Reports:** Coverage and usage reporting with purchase recommendations for Reserved Instances and Savings Plans
- **Custom Reports and Dashboards:** Saveable custom reports and integration with Billing and Cost Management Dashboards
- **API Access:** Programmatic access to cost and usage data through comprehensive REST APIs
- **CSV Export:** Downloadable reports in CSV format for external analysis and integration.
- AWS Cost and Usage Reports (CUR): Provides detailed billing data with resource-level cost and usage information for analysis and optimization.

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/aws-cost-management/aws-cost-explorer/

FAQ: https://aws.amazon.com/aws-cost-management/aws-cost-explorer/faqs/

Technical Overview

Backup and Recovery Capabilities

Backup: AWS Cost Explorer provides data export functionality through CSV downloads and API access for external data preservation. Custom reports and dashboard configurations can be saved within the service, and organisations can implement their own backup strategies using the AWS Cost Explorer API to programmatically extract and archive cost data externally.

Recovery: AWS Cost Explorer is a fully managed service with inherent disaster recovery built into AWS infrastructure, automatically maintaining high availability and data consistency across AWS Regions. Organisations can maintain business continuity through API access and exported data during service interruptions. Supports minimal RPO through continuous data export and low RTO with 24-hour refresh cycles.

Setup Process

Onboarding: Cost Explorer activation is automatic upon AWS account creation, requiring no additional setup or configuration. Users access the service through the AWS Cost Management console with data available within 24 hours. AWS Identity and Access Management (AWS IAM) permissions can be configured for granular access control across teams and departments.

Offboarding: AWS Cost Explorer cannot be disabled once enabled, helping ensure continuous cost visibility. Organisations can remove saved reports and custom configurations through AWS IAM policies, but underlying cost data remains accessible as part of AWS's billing infrastructure following standard account closure procedures.

Data Management

- **Storage Options:** Managed service with 13 months historical data retention and optional extended multi-year analysis
- **Data Removal:** Cannot be disabled once enabled; data removal follows standard AWS account closure procedures
- **Cross-Region Replication:** Built-in high availability and data consistency automatically maintained across AWS Regions
- **Encryption:** Integrated with AWS billing infrastructure security with AWS IAM-controlled access policies and permissions

Optimise Cost and Consumption

AWS Cost Explorer's automated Reserved Instance and Savings Plans recommendations analyse usage patterns and can achieve up to 75% compute cost reductions. The service provides Amazon EC2 rightsizing recommendations that reduce oversized instance costs. AI-powered 18-month forecasting enables optimal capacity planning decisions while automated cost anomaly detection prevents budget overruns and unexpected charges across AWS environments.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/cost-management/latest/userguide/ce-what-is.html

Getting Started: https://docs.aws.amazon.com/cost-management/latest/userguide/ce-getting-started.html

API Reference: https://docs.aws.amazon.com/aws-cost-management/latest/APIReference/API_Operations_AWS_Cost_Explorer_Service.html

Service Quotas: https://docs.aws.amazon.com/cost-management/latest/userguide/management-limits.html

AWS Database Migration Service (AWS DMS)

General Information

AWS Database Migration Service (AWS DMS) is a fully managed service that enables migration of relational databases, data warehouses, and NoSQL databases to AWS or between environments. It automates discovery, schema conversion, and data migration with ongoing replication capabilities.

Main Benefits

AWS DMS reduces traditional migration complexity through fully managed infrastructure, automatic scaling, and built-in schema conversion. It delivers minimal downtime with continuous data replication, supports up to 35% cost savings with Database Savings Plans, and enables rapid disaster recovery with low RPO objectives.

Key Use Cases

- **Database Modernisation:** Migrate legacy databases to modern engines like Amazon Aurora or open-source databases
- **Cloud Migration:** Move on-premises databases to AWS managed services like Amazon Relational Database Service (Amazon RDS) with minimal downtime
- **Data Warehouse Migration:** Migrate data warehouses to Amazon Redshift for enhanced analytics and reporting capabilities
- **Disaster Recovery:** Set up cross-Region database replication for business continuity and data protection.

Core Features

- **Schema Conversion:** Automatically assesses and converts source schemas to target database engine format for migrations
- **Change Data Capture:** Replicates ongoing changes to keep sources and targets synchronised in real-time continuously
- **Homogeneous Migrations:** Serverless migrations between same database engines with automatic resource scaling and management
- **Heterogeneous Migrations:** Supports migrations between different database engines with built-in schema conversion capabilities
- **Automatic Failover:** Provides high availability with backup replication servers taking over during failures
- **Data Encryption:** Encrypts data at rest using AWS Key Management Service (AWS KMS) and in transit using SSL/TLS protocols
- **Multi-Availability Zone Deployment:** Supports replication instances across multiple Availability Zones for enhanced high availability protection.

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/dms/

FAQ: https://aws.amazon.com/dms/faqs/

Technical Overview

Backup and Recovery Capabilities

Backup: AWS DMS does not provide built-in backup capabilities and focuses on data migration and replication rather than backup functionality. Target databases created through AWS DMS migrations can be backed up using AWS Backup or native database backup features, with backup responsibilities handled by the target database services.

Recovery: AWS DMS enables disaster recovery through continuous data replication to secondary AWS Regions and cross-Region replication for maintaining synchronised standby databases. Automatic failover provides high availability with backup replication servers. Supports customer RPO objectives through continuous change data capture replicating changes in near real-time and RTO objectives through rapid failover scenarios with maintained synchronised replica databases.

Setup Process

Onboarding: Setup begins through AWS DMS console by creating prerequisite resources including Virtual Private Cloud, security groups, and target databases. Create replication instances, configure source and target endpoints, then create migration tasks.

Offboarding: Delete replication tasks to stop data synchronisation, terminate replication instances and delete Amazon Elastic Block Store (Amazon EBS) storage volumes. Manually remove target database data from Amazon RDS or Amazon Simple Storage Service. AWS CloudFormation templates can automate resource cleanup including endpoints, subnet groups, and certificates.

Data Management

- **Storage Options:** Uses Amazon EBS storage volumes for replication instances with 30,000 GB storage quota per AWS Region
- **Data Removal:** Delete replication tasks, terminate instances, remove Amazon EBS volumes, and manually delete target database data
- **Cross-Region Replication:** Supports cross-Region database replication for disaster recovery with synchronised standby databases maintained
- **Encryption:** Encrypts data at rest using AWS KMS and in transit using SSL/TLS protocols.

Optimise Cost and Consumption

AWS DMS serverless capabilities automatically scale based on workload demand, preventing over-provisioning costs. Database Savings Plans deliver up to 35% savings with one-year commitments. Right-sizing replication instances according to actual throughput requirements optimises resource consumption. Parallel loading for large tables and optimised change processing accelerate migrations, reducing billable duration and associated costs.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/dms/latest/userguide/Welcome.html

Getting Started: https://docs.aws.amazon.com/dms/latest/userguide/CHAP_GettingStarted.html

API Reference: https://docs.aws.amazon.com/dms/latest/APIReference/Welcome.html

Service Quotas: https://docs.aws.amazon.com/dms/latest/userguide/CHAP_Limits.html

AWS DataSync

General Information

AWS DataSync is an online data transfer service that simplifies and accelerates data migrations to AWS and movement between on-premises storage, cloud providers, and AWS services. It automates transfer processes using purpose-built protocols for high-performance, secure data movement.

Main Benefits

AWS DataSync delivers up to 10 times faster transfer speeds than open-source tools while eliminating expensive commercial network acceleration software costs. It automates data transfer processes with consumption-based pricing, supports up to 100 million objects per task, and provides built-in compression and incremental transfers.

Key Use Cases

- **Data Migration:** Transfer active datasets from on-premises or other cloud providers to AWS storage services
- **Data Archiving:** Move cold data to durable long-term storage classes like Amazon Simple Storage Service (Amazon S3) Glacier Deep Archive
- **Data Replication:** Copy data to AWS for business continuity, disaster recovery, and cross-Region backups
- **Hybrid Cloud Workflows:** Enable recurring transfers between on-premises and AWS for ongoing processing and analysis.

Core Features

- **Purpose-built Network Protocol:** AWS-designed transfer protocol with parallel, multi-threaded architecture for optimised data movement
- **Network Optimisation:** Performs incremental transfers, in-line compression, sparse file detection, and granular bandwidth controls
- **Multi-cloud Support:** Supports data movement between on-premises, AWS, Microsoft Azure, Google Cloud, and object storage systems
- **Enhanced and Basic Modes:** Enhanced mode offers higher performance while Basic mode provides broader compatibility
- **Security and Encryption:** Encrypts data in transit using TLS 1.3 with Kerberos authentication and data integrity validation
- **Metadata Preservation:** Preserves file permissions, timestamps, ownership, and other metadata during all data transfers
- **Scheduling and Automation:** Built-in scheduling mechanism for periodic data transfer tasks and automated transfer management
- **Monitoring and Auditing:** Provides JSON-formatted task reports, Amazon CloudWatch integration, and AWS CloudTrail auditing capabilities.

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/datasync/

FAQ: https://aws.amazon.com/datasync/faqs/

Technical Overview

Backup and Recovery Capabilities

Backup: AWS DataSync transfers data to Amazon S3 storage classes including Glacier Instant Retrieval, Glacier Flexible Retrieval, and Glacier Deep Archive for durable long-term storage. Backups can be created manually via AWS Management Console, CLI, or API, or automated through AWS Backup integration with scheduled recurring transfers and incremental backup capabilities.

Recovery: AWS DataSync enables cross-Region replication of Amazon Elastic File System (Amazon EFS) and Amazon FSx file systems with high-speed transfer capabilities up to 10 times faster than traditional tools. Supports aggressive recovery time targets through automation and reduces data restore timeframes. Supports customer RPO objectives with hourly incremental transfers and RTO objectives through accelerated transfer speeds and automated processes.

Setup Process

Onboarding: Setup involves creating AWS Identity and Access Management (AWS IAM) users with appropriate permissions and deploying DataSync agents on VMware ESXi, KVM, Nutanix AHV, Microsoft Hyper-V, or Amazon Elastic Compute Cloud for on-premises transfers. Configure transfer tasks with bandwidth controls and scheduling options via console, CLI, or API.

Offboarding: Delete DataSync tasks, locations, and agents through console or API. Remove associated AWS IAM roles and policies created for AWS DataSync operations. Transferred data remains in destination storage services and requires separate deletion procedures.

Data Management

- **Storage Options:** Supports Amazon S3 all storage classes, Amazon EFS, Amazon FSx Windows/Lustre/OpenZFS/NetApp ONTAP, NFS, SMB, HDFS, object storage
- **Data Removal:** DataSync does not store data permanently; delete tasks/locations/agents via console or API, clean destination separately
- **Cross-Region Replication:** Enables cross-Region replication of Amazon EFS and Amazon FSx file systems to different AWS Regions
- **Encryption:** Encrypts data in transit using TLS 1.3 with Kerberos authentication and data integrity validation.

Optimise Cost and Consumption

AWS DataSync provides granular bandwidth throttling controls to minimise network usage impact during transfers. Incremental transfer capabilities move only changed data, reducing transfer volumes and associated costs. Built-in compression automatically reduces data transfer sizes. Enhanced mode delivers better performance per dollar through parallel processing. Schedule transfers during off-peak hours to optimise network costs, and leverage Direct Connect for reduced data transfer charges on large migrations.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/datasync/latest/userguide/what-is-datasync.html
Getting Started: https://docs.aws.amazon.com/datasync/latest/userguide/setting-up.html
API Reference: https://docs.aws.amazon.com/datasync/latest/userguide/API_Reference.html
Service Quotas: https://docs.aws.amazon.com/datasync/latest/userguide/datasync-limits.html

AWS Device Farm

General Information

AWS Device Farm is an app testing service enabling developers to test Android, iOS, and web applications on real physical devices hosted by AWS. The service provides automated testing and remote access capabilities for comprehensive cross-device compatibility validation.

Main Benefits

AWS Device Farm reduces annual physical device lab costs while providing access to constantly updated real devices. The service offers parallel test execution, 15-month data retention, and Continuous Integration and Continuous Deployment (CI/CD) integration, accelerating development cycles through immediate feedback and comprehensive testing coverage.

Key Use Cases

- **Mobile App Testing:** Testing Android and iOS apps across multiple real devices for compatibility
- **Cross-Device Compatibility:** Verifying functionality across different screen sizes, operating systems, and device manufacturers
- **CI/CD Pipeline Integration:** Incorporating automated testing into continuous integration and deployment workflows
- **Bug Reproduction and Debugging:** Using remote access sessions to troubleshoot app behaviour in real-time.

Core Features

- **Real Device Testing:** Access to wide range of physical Android and iOS devices for authentic testing conditions
- **Remote Access Sessions:** Real-time interaction with devices through web browsers with 150-minute session limits
- **Automated Test Execution:** Parallel test execution across multiple devices using various testing frameworks
- **Managed Appium Endpoint:** Fully-managed endpoint for connecting interactive tests from local integrated development environments with live video streaming
- **Multiple Test Framework Support:** Support for Appium, XCTest, XCTest UI, Instrumentation, and built-in fuzz tests
- **Private Device Fleet:** Dedicated physical devices deployed exclusively for specific AWS accounts in secure data centres
- **Test Reporting and Analytics:** Comprehensive reports with crash information, logs, screenshots, and performance data retained 15 months
- **VPC Integration:** Secure connectivity to private endpoints and on-premises applications through VPC-ENI.

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/device-farm/
FAQ: https://aws.amazon.com/device-farm/faqs/

Technical Overview

Backup and Recovery Capabilities

Backup: AWS Device Farm provides limited native backup capabilities with automatic data retention policies. Test artifacts and files are stored for 30 days, while session logs and videos are retained for 400 days. Manual deletion is available through console or API. Customers must implement separate backup strategies for critical test assets.

Recovery: Limited disaster recovery due to single-Region deployment in us-west-2. Recovery relies on AWS infrastructure resilience within the AWS Region using multiple Availability Zones. Organisations must maintain alternative testing strategies for service disruptions. RPO and RTO objectives require customer-implemented backup processes for aggressive recovery requirements.

Setup Process

Onboarding: Setup requires AWS account creation and AWS Identity and Access Management user configuration with AWS Device Farm permissions. Create projects through the Device Farm console, upload applications (.apk/.ipa files), select testing approach (automated or remote access), and configure device pools. Free tier available for evaluation.

Offboarding: Data automatically deleted after retention periods (30 days for files, 400 days for logs). Manual deletion available via console or API. Private devices undergo secure hardware destruction, while public devices receive comprehensive cleanup with factory resets.

Data Management

- **Storage Options:** Test artifacts stored with Amazon Simple Storage Service integration, automatic lifecycle management for data retention
- **Data Removal:** Apps and files automatically deleted after 30 days, logs retained up to 15 months, manual deletion available
- **Cross-Region Replication:** Not supported, service only available in us-west-2 Region limiting cross-Region strategies
- **Encryption:** SSL/TLS connections mandatory for all communications, secure data centres for private device fleets.

Optimise Cost and Consumption

AWS Device Farm's device slot purchasing provides bulk pricing for frequent testing workflows, avoiding per-minute charges through reserved capacity allocation. Strategic device pool configuration targets specific device combinations, optimising test coverage while managing resource consumption. The 150-minute session timeout automatically controls usage duration, preventing runaway costs. Private devices offer fixed-cost dedicated resources without additional per-usage fees once deployed, enabling predictable testing budgets for high-volume scenarios.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/devicefarm/latest/developerguide/welcome.html
Getting Started: https://docs.aws.amazon.com/devicefarm/latest/developerguide/getting-started.html
API Reference: https://docs.aws.amazon.com/devicefarm/latest/APIReference/Welcome.html
Service Quotas: https://docs.aws.amazon.com/general/latest/gr/devicefarm.html

AWS Direct Connect

General Information

AWS Direct Connect establishes dedicated network connections from on-premises environments to AWS, bypassing the public internet. It provides consistent performance, reduced latency, and enhanced security for hybrid cloud architectures through standard Ethernet fiber-optic cables.

Main Benefits

AWS Direct Connect delivers consistent low latency, enhanced security through private connections, and reduced network costs with lower data transfer charges. Supporting speeds up to 400 Gbps, it provides predictable network behaviour and potential cost savings for high-volume data transfers.

Key Use Cases

- **Hybrid Cloud Connectivity:** Connecting on-premises data centers to AWS for hybrid architectures
- **High-Bandwidth Data Transfer:** Large-scale migration, backup, and synchronisation with reliable throughput and reduced costs
- **Low-Latency Applications:** Real-time applications, financial trading systems, and media streaming requiring predictable performance
- **Regulated Industries:** Private connections for compliance, security policies, and restricted public internet usage.

Core Features

- **Multiple Connection Types:** Dedicated connections (1-400 Gbps) and hosted connections (50 Mbps-10 Gbps) through AWS Partners
- **Virtual Interfaces:** Private, public, and transit VIFs for accessing Virtual Private Clouds (VPCs), public AWS services, and Transit Gateways
- **Direct Connect Gateway:** Globally available resource enabling connections to multiple VPCs and Transit Gateways across AWS Regions
- **SiteLink:** Private site-to-site connectivity between AWS Direct Connect locations using AWS global backbone infrastructure
- **Encryption Options:** MACsec encryption for dedicated connections and IPsec VPN support for additional security layers
- **Link Aggregation Groups:** Aggregating multiple connections for increased bandwidth capacity and enhanced redundancy
- **BGP Routing:** Dynamic routing with BGP and BGP MD5 authentication for reliable and secure path selection
- **Global Availability:** Available at numerous locations worldwide supporting connection speeds up to 400 Gbps.

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/directconnect/

FAQ: https://aws.amazon.com/directconnect/faqs/

Technical Overview

Backup and Recovery Capabilities

Backup: AWS Direct Connect does not provide native backup capabilities as a network connectivity service. However, it enhances backup operations by providing dedicated, high-bandwidth connections for reliable data transfer to Amazon Simple Storage Service, Amazon Elastic Block Store snapshots, and AWS storage services with consistent performance and reduced costs compared to internet-based transfers.

Recovery: AWS Direct Connect enables disaster recovery through stable, high-bandwidth connections for continuous data replication, failover operations, and recovery procedures. Multiple connection options and redundancy features minimise single points of failure. Supports customer RPO (frequent backup windows and faster data synchronisation) and RTO (predictable connectivity reducing network-related delays in failover procedures) objectives.

Setup Process

Onboarding: Order connections through AWS Management Console using Connection wizard, selecting bandwidth requirements (50 Mbps-400 Gbps) and resiliency levels. Download LOA-CFA, coordinate cross-connect installation, configure virtual interfaces, and establish Border Gateway Protocol (BGP) sessions with comprehensive AWS documentation and support.

Offboarding: Delete virtual interfaces, terminate BGP sessions, and remove physical cross-connects at AWS Direct Connect locations. Coordinate with network providers to disconnect cables and return equipment. Follow internal data governance procedures for connected AWS services.

Data Management

- **Storage Options:** Network connectivity service providing no data storage, only network transport for accessing AWS storage services
- **Data Removal:** Physical connection contains no customer data; deletion handled through standard AWS service processes for connected resources
- **Cross-Region Replication:** Enables multi-Region connectivity through AWS Direct Connect Gateway for accessing VPCs and Transit Gateways across Regions
- **Encryption:** MACsec encryption for dedicated connections and IPsec VPN support for additional security layers.

Optimise Cost and Consumption

Right-sizing connections based on bandwidth usage patterns optimises port hour costs. Hosted connections reduce expenses for lower bandwidth requirements. Link Aggregation Groups consolidate multiple connections for increased usage efficiency. AWS Direct Connect Gateway enables sharing connections across multiple VPCs and accounts. AWS Direct Connect SiteLink provides site-to-site connectivity reducing multiple internet circuits. Traffic engineering maximises existing connection usage before adding capacity.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/directconnect/latest/UserGuide/index.html
Getting Started: https://docs.aws.amazon.com/directconnect/latest/UserGuide/Welcome.html
API Reference: https://docs.aws.amazon.com/directconnect/latest/APIReference/index.html
Service Quotas: https://docs.aws.amazon.com/directconnect/latest/UserGuide/limits.html

AWS Directory Service

General Information

AWS Directory Service provides multiple ways to use Microsoft Active Directory and Lightweight Directory Access Protocol (LDAP) with AWS services. It offers three main options: AWS Managed Microsoft AD for full functionality, AD Connector for hybrid connectivity, and Simple AD for basic requirements.

Main Benefits

AWS Directory Service reduces operational overhead of maintaining domain controllers through fully managed infrastructure with automatic patching, monitoring, and multi-Availability Zone high availability. It integrates with multiple AWS services, supports Health Insurance Portability and Accountability Act and Payment Card Industry Data Security Standard compliance, and enables rapid 20-40 minute deployment.

Key Use Cases

- **Enterprise Directory Migration:** Migrate Active Directory-aware applications like SharePoint, SQL Server to AWS Cloud
- **Hybrid Cloud Integration:** Extend on-premises Active Directory infrastructure to AWS using AD Connector or trusts
- **Single Sign-On:** Enable users to access AWS Management Console using existing Active Directory credentials
- **AWS Service Authentication:** Authenticate and authorise access for Amazon WorkSpaces, Amazon Relational Database, Amazon Chime, and Amazon Connect

Core Features

- **AWS Managed Microsoft AD:** Fully managed Windows Server 2019 Active Directory with automatic patching and multi-AZ high availability
- **AD Connector:** Proxy service connecting AWS applications to existing on-premises Microsoft Active Directory without caching data
- **Simple AD:** Samba 4-based directory providing basic Active Directory compatibility for low-scale, cost-effective directory needs
- **Multi-Region Replication:** Enterprise Edition feature enabling directory replication across multiple AWS Regions for global workload support
- **Directory Sharing:** Share directories across multiple AWS accounts and Virtual Private Clouds (VPCs) within an organisation for centralised management
- **Trust Relationships:** Create trust relationships between AWS Managed Microsoft AD and existing on-premises or cloud domains
- **Schema Extensions:** Extend Active Directory schema to support custom attributes and object classes for specialised applications
- **LDAPS Support:** Secure LDAP communications with server-side and client-side LDAPS encryption for enhanced data protection.

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/directoryservice/

FAQ: https://aws.amazon.com/directoryservice/faqs/

Technical Overview

Backup and Recovery Capabilities

Backup: AWS Directory Service offers manual snapshots for AWS Managed Microsoft AD directories (up to 5 per directory with 180-day maximum retention) via AWS Management Console and automated daily snapshots managed by AWS. AD Connector and Simple AD have limited backup capabilities since they proxy to existing infrastructure.

Recovery: Point-in-time recovery through snapshot restoration, multi-Availability Zone automatic failover, and Enterprise Edition multi-Region replication for geographic disaster recovery. Built-in high availability across multiple zones ensures service continuity. Supports customer RPO (manual snapshots within 5-snapshot limit) and RTO (automatic failover and managed infrastructure) objectives.

Setup Process

Onboarding: Select directory type via AWS Management Console, create VPC with two subnets in different Availability Zones. Configure Domain Name System name, NetBIOS name, and administrator password. Service automatically creates domain controllers within 20-40 minutes. 30-day trial available with AWS Free Tier.

Offboarding: Deauthorise connected AWS applications, remove trust relationships, unjoin Amazon Elastic Compute Cloud instances from domain. Create manual snapshots before deletion. Directory deletion permanently removes all data and cannot be recovered.

Data Management

- **Storage Options:** Managed directory data stored across multiple Availability Zones with automated daily snapshots and manual snapshots
- **Data Removal:** Directory deletion permanently removes all users, groups, organisational units, and configuration settings without recovery option
- **Cross-Region Replication:** Enterprise Edition supports multi-Region directory replication across multiple AWS Regions for global workload management
- **Encryption:** Built-in encryption with LDAPS support providing server-side and client-side secure LDAP communications throughout service.

Optimise Cost and Consumption

Right-size directory editions by monitoring Amazon CloudWatch performance metrics to optimise domain controller counts based on actual usage patterns. Directory sharing across multiple AWS accounts reduces duplicate directory deployments within organisations. Avoid NETLOGON and SYSVOL shares for file storage in multi-Region deployments to minimise data transfer costs. Simple AD and AD Connector provide no-cost options when integrated with Amazon WorkSpaces, meeting minimum user requirements.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/directoryservice/latest/admin- guide/what_is.html

Getting Started: https://docs.aws.amazon.com/directoryservice/latest/admin-guide/ms_ad_getting_started.html

API Reference: https://docs.aws.amazon.com/directoryservice/latest/devguide/welcome.html

Service Quotas: https://docs.aws.amazon.com/general/latest/gr/ds_Region.html

AWS Elastic Beanstalk

General Information

AWS Elastic Beanstalk is a service that simplifies deploying and scaling web applications across multiple programming languages. It automatically handles infrastructure provisioning, load balancing, and health monitoring while maintaining full control over underlying AWS resources.

Main Benefits

AWS Elastic Beanstalk accelerates time-to-market by reducing infrastructure management complexity while maintaining full control over AWS resources. With zero additional service fees, automatic scaling capabilities, and support for multiple deployment strategies across eight programming languages, it delivers significant cost optimisation and rapid deployment capabilities.

Key Use Cases

- **Web Application Deployment:** Rapid deployment across multiple programming languages without infrastructure management
- **DevOps Acceleration:** Streamlined Continuous Integration and Continuous Deployment workflows with blue/green deployments and rollback capabilities
- **Microservices Architecture:** Deploy containerised applications and APIs with automatic scaling and balancing
- **Development Environment Provisioning:** Quick setup of development, testing, and staging environments.

Core Features

- **Multi-platform Support:** Supports Java, .NET, PHP, Node.js, Python, Ruby, Go, and Docker containers
- **Automatic Scaling:** Built-in auto-scaling capabilities using Amazon Elastic Compute Cloud (Amazon EC2) Auto Scaling with Amazon CloudWatch metrics and triggers
- **Load Balancing:** Integrated Elastic Load Balancing with support for Application, Network, and Classic Load Balancers
- **Health Monitoring:** Comprehensive application health monitoring with automatic instance replacement and detailed health dashboards
- **Application Versioning:** Complete application lifecycle management with version control, rollback capabilities, and blue/green deployments
- **Configuration Management:** Centralised configuration templates and integration with AWS Systems Manager Parameter Store and AWS Secrets Manager
- **Security Integration:** Native AWS Identity and Access Management integration, Virtual Private Cloud support, and compliance with Health Insurance Portability and Accountability Act (HIPAA), Payment Card Industry Data Security Standard (PCI DSS), and System and Organization Controls (SOC) standards
- **Infrastructure Automation:** Automatic provisioning of Amazon EC2 instances, security groups, load balancers, and AWS CloudFormation stacks.

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/elasticbeanstalk/

FAQ: https://aws.amazon.com/elasticbeanstalk/faqs/

Technical Overview

Backup and Recovery Capabilities

Backup: Application versions and configuration templates stored in Amazon Simple Storage Service (Amazon S3) with configurable lifecycle policies, preserving up to 1000 application versions by default. Backups can be created manually through AWS Management Console, CLI, or API, or automated via AWS Backup with scheduled policies for underlying infrastructure.

Recovery: Multi-Availability Zone deployments provide sub-minute failover capabilities, blue/green deployment strategies minimise downtime, and cross-Region deployment enables recovery from regional outages. Auto Scaling and health monitoring provide automatic instance failure recovery. Supports customer RPO objectives through continuous Amazon S3 version storage and RTO objectives via automated scaling and rapid environment recreation.

Setup Process

Onboarding: Access through AWS Management Console, CLI, or EB CLI by creating applications and uploading code via ZIP files, Git integration, or container images. Configuration options include customising instance types, scaling policies, and environment variables with automatic AWS resource provisioning.

Offboarding: Terminate environments through AWS Management Console, CLI, or API with automatic removal of Amazon EC2 instances, load balancers, and security groups. Manual deletion of Amazon S3-stored application versions or automated cleanup via lifecycle policies ensures complete data removal.

Data Management

- **Storage Options:** Application versions stored in Amazon S3 with configurable lifecycle policies, preserving up to 1000 versions
- **Data Removal:** Manual deletion via AWS Management Console, CLI, or API with automatic lifecycle policies for complete data cleanup
- **Cross-Region Replication:** Cross-Region application deployment capabilities supported for geographic redundancy and disaster recovery scenarios
- **Encryption:** Native integration with AWS Secrets Manager and compliance with HIPAA, SOC, and PCI DSS standards.

Optimise Cost and Consumption

AWS Elastic Beanstalk's automatic scaling capabilities reduce over-provisioning by adjusting capacity based on traffic patterns. Spot Instances provide significant savings for development and testing environments. Time-based scaling matches resource allocation to predictable workloads. Application version lifecycle policies automatically clean up old versions to reduce Amazon S3 storage costs. Single-instance environments optimise development scenario expenses while Savings Plans and Reserved Instances support predictable workloads.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/elasticbeanstalk/latest/dg/Welcome.html

Getting Started: https://docs.aws.amazon.com/elasticbeanstalk/latest/dg/GettingStarted.html

API Reference: https://docs.aws.amazon.com/elasticbeanstalk/latest/api/Welcome.html

Service Quotas: https://docs.aws.amazon.com/elasticbeanstalk/latest/api/API_ResourceQuotas.html

AWS Elastic Disaster Recovery (AWS DRS)

General Information

AWS Elastic Disaster Recovery (AWS DRS) minimises downtime and data loss through continuous block-level replication of on-premises and cloud applications. It enables rapid recovery by converting replicated servers to run natively on AWS within minutes during disasters.

Main Benefits

AWS DRS delivers superior recovery performance with RTOs measured in minutes and RPOs in seconds through continuous block-level replication. Its staging area architecture eliminates expensive idle recovery infrastructure while maintaining continuous protection, significantly reducing costs compared to traditional disaster recovery solutions.

Key Use Cases

- **Disaster Recovery:** Protect on-premises, Amazon Elastic Compute Cloud (Amazon EC2), and third-party cloud workloads with rapid failover
- **Ransomware Recovery:** Launch clean, unencrypted server versions from point-in-time snapshots before attacks
- **Cross-Region Protection:** Replicate applications between AWS Regions for protection against Regional outages
- **Migration and Modernisation:** Enable one-time migration of workloads from on-premises to AWS infrastructure.

Core Features

- **Block-level Continuous Replication:** Provides near real-time data replication with RPOs of seconds using advanced technology
- **Point-in-Time Recovery:** Creates automated snapshots every 10 minutes, hourly for 24 hours, daily for 7 days.
- **Launch Management Settings:** Configure default launch settings, bulk modifications, and recovery Availability Zone selection
- **Post-Launch Actions:** Enable automated actions after recovery instance launch using AWS Systems Manager commands
- **Automated Network Replication:** Replicates AWS network components including subnets, security groups, and route tables
- **Drill Testing:** Provides non-disruptive testing capabilities to validate disaster recovery plans without impacting production
- **Staging Area Design:** Reduces expensive idle recovery infrastructure while maintaining continuous protection and reducing costs
- **Multi-Environment Support:** Supports on-premises, Amazon EC2, and third-party cloud workloads with comprehensive failover capabilities.

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/disaster-recovery/

FAQ: https://aws.amazon.com/disaster-recovery/faqs/

Technical Overview

Backup and Recovery Capabilities

Backup: AWS DRS creates automated point-in-time snapshots with crash-consistent recovery points stored in AWS staging areas. Snapshots are automated every 10 minutes for the last hour, hourly for 24 hours, and daily for 7 days. Works complementarily with AWS Backup for comprehensive data protection strategies.

Recovery: Enables rapid recovery with automated server conversion and launch processes, cross-Region and cross-Availability Zone disaster recovery options, and non-disruptive drill testing capabilities. Supports customer RPO objectives (continuous block-level replication with seconds-level protection) and RTO objectives (minutes-level recovery through automated failover).

Setup Process

Onboarding: Initialise AWS DRS in target AWS Regions through AWS Management Console. Install AWS Replication Agent on source servers to begin automatic continuous replication. Configure launch settings and perform initial recovery drills to validate setup.

Offboarding: Stop replication on source servers and terminate recovery instances. Delete source server resources from AWS DRS console and clean up staging area resources including replication servers and Amazon Elastic Block Store (Amazon EBS) volumes.

Data Management

- **Storage Options:** Uses low-cost staging areas in AWS with Amazon EBS volumes for replicated data storage
- **Data Removal:** Stop replication, terminate recovery instances, delete source servers, and clean up staging area resources
- **Cross-Region Replication:** Supports replication between AWS Regions for protection against Regional outages and disasters
- **Encryption:** Integrates with AWS security services and uses secure communication through Virtual Private Cloud endpoints and AWS PrivateLink.

Optimise Cost and Consumption

AWS DRS features a staging area architecture that uses minimal compute resources during normal operations, scaling up only during recovery events. Right-sizing recovery instances and selecting appropriate Amazon EBS volume types optimise storage costs. Efficient drill testing schedules reduce testing expenses while automated snapshot retention policies manage storage consumption. Proper resource cleanup after recovery operations prevents ongoing charges for unused infrastructure components.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/drs/latest/userguide/what-is-drs.html
Getting Started: https://docs.aws.amazon.com/drs/latest/userguide/getting-started.html
API Reference: https://docs.aws.amazon.com/drs/latest/APIReference/Welcome.html
Service Quotas: https://docs.aws.amazon.com/general/latest/gr/drs.html

AWS Elemental Services – AWS Elemental MediaConnect

General Information

AWS Elemental MediaConnect is a reliable, secure transport service for live video that enables broadcasters to ingest content into AWS Cloud and distribute to multiple destinations. It provides quality-of-service over IP transport with packet recovery and industry-standard encryption.

Main Benefits

AWS Elemental MediaConnect offers faster deployment (minutes vs months) compared to traditional satellite and fibre solutions, with immediate source failover with less than 1 second of disruption and consumption-based pricing without upfront commitments. It provides professional-grade quality of service, comprehensive monitoring, and seamless integration with AWS Media Services ecosystem.

Key Use Cases

- **Contribution:** Ingest content from on-premises encoders into AWS Cloud using transport stream or Cloud Data Interface (CDI) flows
- **Distribution:** Deliver content to different geographical locations using multi-Region flows for global reach
- **Entitlements:** Share content securely with other AWS accounts through permission-based access controls
- **Replication and Monitoring:** Distribute video to multiple destinations and enable real-time monitoring capabilities.

Core Features

- **Multiple Protocol Support:** Supports Zixi, RIST, SRT, RTP, RTP-FEC, NDI, and other industry-standard protocols for reliable contribution
- **CDI and JPEG XS:** Provides uncompressed and visually-lossless video transport using AWS Cloud Digital Interface and compression
- **Source Failover:** Offers Merge and Failover modes using two redundant sources to minimise service disruption
- **End-to-End Encryption:** Built-in security with AES encryption, SPEKE key management, and IP allow listing capabilities
- **Real-Time Monitoring:** Provides network health metrics, content quality monitoring, and Amazon EventBridge integration for comprehensive oversight
- **Workflow Monitor:** Discover, visualise, and monitor live video resources across workflows for enhanced operational visibility
- **Router Integration:** Real-time routing of live video streams between sources and destinations in the cloud
- **Quality of Service:** Adds quality-of-service layer over standard IP transport with packet recovery for uninterrupted connections.

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/mediaconnect/
FAQ: https://aws.amazon.com/mediaconnect/faqs/

Technical Overview

Backup and Recovery Capabilities

Backup: AWS Elemental MediaConnect does not provide traditional backup capabilities as it is a live video transport service that does not persistently store video content. Backup strategies must be implemented at source or destination endpoints, such as using AWS Elemental MediaLive to create archive outputs to Amazon Simple Storage Service.

Recovery: Source failover capabilities switch between redundant sources with less than 1 second of disruption using Merge and Failover modes. Multi-Region deployment enables quick regional failover with built-in infrastructure redundancy and automatic recovery mechanisms. Supports customer RPO (near-zero through real-time replication) and RTO (sub-second through automated failover) objectives.

Setup Process

Onboarding: Setup involves creating AWS account, configuring AWS Identity and Access Management users, and accessing AWS MediaConnect console, CLI, or API. Create flows with source details including transport protocol and encryption type, configure destinations. Deployment occurs in minutes with minimal infrastructure requirements and step-by-step tutorials.

Offboarding: Stop and delete all flows to avoid continued charges. Delete flow configurations, source settings, output destinations, and entitlements through console, CLI, or API. Clean up associated resources including encryption keys in AWS Secrets Manager and monitoring configurations.

Data Management

- **Storage Options:** Live video transport service that does not persistently store customer video content during transmission
- **Data Removal:** Delete flow configurations, source settings, output destinations, and entitlements through AWS Management Console, CLI, or API
- **Cross-Region Replication:** Multi-Region flows enable content delivery to different geographical locations for global distribution workflows
- **Encryption:** Built-in AES encryption, SPEKE key management integration, and IP allow listing for comprehensive security.

Optimise Cost and Consumption

Reserved outbound bandwidth pricing with up to 12-month commitments provides discounted internet data transfer rates for predictable workloads. Flow scheduling enables stopping inactive flows to avoid charges since idle flows incur no costs. VPC outputs eliminate internet transfer charges for regional distribution. Bitrate and compression optimisation balances quality requirements with data transfer costs, while selecting appropriate flow types between transport stream and CDI optimises performance versus cost requirements.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/mediaconnect/latest/ug/what-is.html
Getting Started: https://docs.aws.amazon.com/mediaconnect/latest/ug/getting-started.html
API Reference: https://docs.aws.amazon.com/mediaconnect/latest/api/Welcome.html
Service Quotas: https://docs.aws.amazon.com/mediaconnect/latest/ug/quotas.html

AWS Elemental Services – AWS Elemental MediaConvert

General Information

AWS Elemental MediaConvert is a file-based video processing service that provides scalable transcoding to convert input videos into multiple formats and resolutions. It offers broadcast-grade features including professional codecs, High Dynamic Range (HDR) support, adaptive bitrate packaging, and advanced audio processing for premium content delivery across devices.

Main Benefits

AWS Elemental MediaConvert reduces video processing infrastructure complexity while delivering broadcast-grade capabilities that other cloud encoding services often lack. It offers accelerated transcoding up to 25 times faster, tiered pricing with volume-based discounts, and elastic scaling without capacity planning, enabling organisations to save up to 85% annually through optimised processing.

Key Use Cases

- **Video-on-Demand Content Creation:** Converting media files for broadcast and multiscreen delivery with adaptive streaming
- **Content Format Conversion:** Transcoding files for compatibility, compression, and distribution across different devices and platforms
- **Broadcast-Grade Processing:** Creating premium content with professional codecs, HDR support, and closed captioning capabilities
- **Media Library Migration:** Processing large archives with automated scaling and volume discounts for cost optimisation.

Core Features

- **Professional Broadcast Codecs:** Support for MPEG-2, AVC, AV1, HEVC with 10-bit 4:2:2 colour sampling and HDR
- **Adaptive Bitrate Packaging:** CMAF, HLS, DASH, and MSS output formats with automated adaptive bitrate configuration for multiscreen delivery
- **Advanced Audio Features:** Multi-language audio processing, audio loudness normalisation, and Dolby Atmos support for premium experiences
- **Content Protection:** Digital rights management integration and encryption capabilities with third-party Digital Rights Management (DRM) provider support
- **Graphics and Overlays:** Static graphic overlays, image insertion, motion graphics, and video overlays for enhanced content
- **Quality Features:** Quality-defined variable bitrate encoding and HDR processing including Dolby Vision for optimal output
- **Accelerated Transcoding:** Complete jobs up to 25 times faster for time-sensitive workflows and urgent content delivery
- **Elastic Scaling:** Automatic resource provisioning with built-in redundancy across multiple Availability Zones for reliability.

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/mediaconvert/

FAQ: https://aws.amazon.com/mediaconvert/faqs/

Technical Overview

Backup and Recovery Capabilities

Backup: AWS Elemental MediaConvert relies on customer-managed Amazon Simple Storage Service (Amazon S3) storage for media file backup with versioning and cross-Region replication capabilities. Backup can be configured manually (Amazon S3 console) or automated via AWS Backup with scheduled protection policies for comprehensive media file availability.

Recovery: AWS Elemental MediaConvert provides automatic component failure recovery through multi-AZ redundant infrastructure without job disruption, cross-Region job recreation using templates and presets, and high availability with automatic failover. Built-in fault tolerance automatically replaces degraded components. Supports customer RPO (Amazon S3 versioning and cross-Region replication) and RTO (multi-AZ infrastructure recovery) objectives.

Setup Process

Onboarding: Setup begins by helping ensure input files are accessible in Amazon S3, HTTP, or HTTPS locations and configuring AWS Identity and Access Management permissions. Create transcoding jobs through AWS MediaConvert console or API/SDK with system presets or custom templates. Jobs and file storage must be in the same AWS Region.

Offboarding: MediaConvert does not store customer content permanently, only processing files during transcoding jobs. Data removal involves deleting input and output files from customer-controlled Amazon S3 buckets to avoid ongoing storage charges.

Data Management

- **Storage Options:** Customer-controlled Amazon S3 buckets for input and output files, plus HTTP/HTTPS locations for input
- **Data Removal:** Delete input and output files from Amazon S3 storage locations; service does not store content permanently
- **Cross-Region Replication:** Supported through Amazon S3 cross-Region replication for source media protection and disaster recovery capabilities
- **Encryption:** Digital rights management integration and encryption capabilities with third-party DRM provider support available.

Optimise Cost and Consumption

AWS Elemental MediaConvert offers automated tiered volume discounts as monthly usage increases, while Quality-Defined Variable Bitrate encoding reduces file sizes while maintaining quality. HEVC encoding provides significant file size reduction for storage and delivery savings. Reserved pricing with 12-month commitments offers fixed rates for predictable workloads, and two-tier pricing structure enables choosing Basic tier for simple distribution or Professional tier only when advanced features are required.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/mediaconvert/latest/ug/what-is.html

Getting Started: https://docs.aws.amazon.com/mediaconvert/latest/ug/getting-started.html

API Reference: https://docs.aws.amazon.com/mediaconvert/latest/apireference/

Service Quotas: https://docs.aws.amazon.com/general/latest/gr/mediaconvert.html

AWS Elemental Services – AWS Elemental MediaLive

General Information

AWS Elemental MediaLive is a broadcast-grade live video processing service that encodes high-quality live video streams for delivery to broadcast televisions and multiscreen devices. It transforms live video content into multiple formats and packages, supporting industry-standard codecs like H.264/AVC and H.265/HEVC with automatic failover across Availability Zones.

Main Benefits

AWS Elemental MediaLive is designed to deliver broadcast-grade quality with 99.9% SLA and fully managed infrastructure that reduces hardware management overhead. The service provides automatic scaling, redundancy across Availability Zones (AZs), and up to 70% cost savings through Reserved Instances, while enabling rapid deployment with seamless AWS Media Services integration.

Key Use Cases

- **Live Streaming:** Processing real-time video into adaptive bitrate streams for multiscreen devices
- **Broadcast Television:** Encoding high-quality streams with ad markers, closed captions, and audio tracks
- **24x7 Linear Workflows:** Building continuous streaming workflows for around-the-clock channel operations
- **Event-Based Streaming:** Creating live streams with dynamic input switching and scheduling capabilities.

Core Features

- **Real-time Video Encoding:** Processes live video with H.264/AVC, H.265/HEVC, and AV1 codec support
- **Broadcast Capabilities:** Supports ad markers, closed captions, multiple language audio tracks, and loudness correction
- **High Availability:** Runs on redundant infrastructure with automatic failover and component replacement capabilities
- **Statistical Multiplexing:** Optimises bandwidth by combining multiple streams into single output for broadcast distribution
- **Input Switching:** Enables multiple inputs per channel with scheduled switching between different video sources
- **Motion Graphics:** Provides overlay capabilities for graphics, logos, and visual elements on live streams
- **Live Video Workflow Monitor:** Delivers comprehensive monitoring and alerting capabilities for video processing workflows
- **Automated Resource Provisioning:** Automatically deploys, scales, and monitors encoding resources based on current demand.

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/medialive/

FAQ: https://aws.amazon.com/medialive/faqs/

Technical Overview

Backup and Recovery Capabilities

Backup: AWS Elemental MediaLive does not provide traditional backup capabilities as it processes live streams in real-time. Channel configurations can be exported manually through the AWS Management Console or automated via AWS CloudFormation Infrastructure as Code tools. Content backup relies on integrated services like Amazon Simple Storage Service (Amazon S3) versioning and cross-Region replication.

Recovery: AWS Elemental MediaLive provides automatic failover across Availability Zones with component replacement within minutes without manual intervention. The service features redundant pipeline architecture helping ensure continuous operation and supports cross-Region disaster recovery through manual configuration replication. Supports customer RPO objectives through minimal data loss during failover and RTO objectives with automatic recovery mechanisms.

Setup Process

Onboarding: Setup through workflow wizard for simplified configuration, console tutorials for additional features, or custom workflow design from scratch. Requires AWS account with AWS Identity and Access Management (AWS IAM) permissions and preliminary setup completion. Supports progressive feature addition including input switching and motion graphics.

Offboarding: Delete channels, inputs, and associated resources via AWS Elemental MediaLive console or API. Content in integrated services like Amazon S3 requires separate deletion. No automatic data retention after resource deletion.

Data Management

- **Storage Options:** Integrates with Amazon S3 for asset storage, MediaStore for media-optimised storage, no persistent data storage
- **Data Removal:** Customer responsibility to delete channels, inputs, and resources via console or API manually
- **Cross-Region Replication:** Supported through manual channel configuration replication and multi-Region deployment setup for workflows
- **Encryption:** Secured through AWS IAM integration and Virtual Private Cloud private networking for input security and access control.

Optimise Cost and Consumption

MediaLive offers reserved instances with up to 70% savings for channels running over 180 hours monthly through 12-month commitments. Statistical multiplexing optimises bandwidth by combining multiple streams into single outputs for broadcast distribution efficiency. Single-pipeline channels reduce costs for non-critical workloads while dual-pipeline supports production requirements. Data-driven reservation management uses AWS Cost and Usage Reports with Amazon Quick dashboards for usage analysis and optimal resource sizing.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/medialive/latest/ug/what-is.html

Getting Started: https://docs.aws.amazon.com/medialive/latest/ug/getting-started.html

API Reference: https://docs.aws.amazon.com/medialive/latest/apireference/what-is.html

Service Quotas: https://docs.aws.amazon.com/general/latest/gr/medialive_Region.html

AWS Elemental Services – AWS Elemental MediaPackage

General Information

AWS Elemental MediaPackage is a cloud-based video packaging and origination service that delivers secure, scalable video streams to various devices and Content Delivery Networks (CDNs). It provides just-in-time packaging, dynamically converting live streams and Video on Demand (VOD) content into multiple delivery formats while applying Digital Rights Management (DRM) protection for broadcast-quality multiscreen delivery.

Main Benefits

AWS Elemental MediaPackage reduces packaging infrastructure complexity through just-in-time processing that reduces storage requirements while supporting multiple output formats from single sources. The managed service delivers broadcast-grade reliability across multiple Availability Zones (AZ) with automatic scaling, built-in input redundancy, and comprehensive DRM support, designed to work with CDNs to achieve cache hit ratios up to 99% for live content.

Key Use Cases

- **Live Streaming:** Real-time packaging for broadcast networks with adaptive bitrate streaming and DRM protection
- **Video-on-Demand:** Packaging and distribution of pre-recorded content with multiple format outputs and protection
- **OTT Video Services:** Secure multiscreen delivery to tablets, smartphones, TVs with time-shifted viewing capabilities
- **Live-to-VOD Creation:** Extract VOD assets from live streams through harvest jobs for on-demand viewing.

Core Features

- **Just-in-Time Packaging:** Dynamically packages content on-demand rather than pre-processing, reducing storage requirements and bandwidth usage
- **Multi-Format Output Support:** Supports MPEG-DASH, HLS, Microsoft Smooth Streaming, and CMAF packaging formats for broad device compatibility
- **Digital Rights Management:** Integrates with Google Widevine, Microsoft PlayReady, and Apple FairPlay DRM technologies for content protection
- **Adaptive Bitrate Streaming:** Automatically adjusts video quality based on viewer's network conditions for optimal viewing experience
- **Time-Shifted Viewing:** Enables DVR-like features including start-over, catch-up TV, and time delay for live content
- **Input Redundancy:** Creates two ingest URLs for automatic failover if primary input fails, helping ensure uninterrupted delivery
- **CDN Authorisation:** Supports secure content delivery through Amazon CloudFront and other CDNs with header-based authorisation
- **Live-to-VOD Asset Creation:** Extracts VOD assets from live streams through harvest jobs for later on-demand viewing.

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/mediapackage/

FAQ: https://aws.amazon.com/mediapackage/faqs/

Technical Overview

Backup and Recovery Capabilities

Backup: AWS Elemental MediaPackage does not provide traditional backup capabilities, focusing on real-time content packaging and delivery. Persistent content protection relies on Amazon Simple Storage Service (Amazon S3) integration through live-to-VOD harvest jobs, which can use Amazon S3's backup features and AWS Backup service. Content is temporarily stored for time-shifted viewing (maximum 14 days) for operational purposes.

Recovery: Recovery is achieved through high-availability architecture across multiple Availability Zones with automatic input redundancy and cross-Region failover capabilities. Instant endpoint activation and automatic scaling minimise downtime during disruptions. Supports low RPO through real-time processing and RTO optimisation through distributed infrastructure and seamless failover mechanisms.

Setup Process

Onboarding: Setup involves creating AWS account with AWS Identity and Access Management permissions, then using AWS Management Console, CLI, or API to create channel groups, channels, and endpoints. Configure upstream encoders with WebDAV authentication for live workflows or create packaging groups for VOD workflows with Amazon S3 integration.

Offboarding: Requires systematic deletion in hierarchical order: endpoints before channels, channels before channel groups. AWS Management Console, CLI, and API provide deletion capabilities with confirmation dialogs to prevent accidental removal.

Data Management

- **Storage Options:** Temporary storage for time-shifted viewing (maximum 14 days), Amazon S3 integration for live-to-VOD assets
- **Data Removal:** Systematic hierarchical deletion required: endpoints before channels, channels before channel groups, with confirmation dialogs
- **Cross-Region Replication:** Cross-Region failover capabilities enable content delivery from alternative AWS Regions during outages and disruptions
- **Encryption:** DRM integration with Google Widevine, Microsoft PlayReady, Apple FairPlay; AWS Secrets Manager stores encryption keys.

Optimise Cost and Consumption

Just-in-time packaging reduces storage requirements by dynamically converting single-format source content into multiple delivery formats on-demand. Amazon CloudFront integration can achieve 99% cache hit ratios for live content and 80% for VOD, reducing origination charges. Rightsizing input streams, efficient encoder settings, and adaptive bitrate streaming optimise bandwidth consumption. Automatic scaling prevents over-provisioning while annual commitment discounts benefit high-volume customers exceeding 10TB monthly usage.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/mediapackage/latest/userguide/what-is.html

Getting Started: https://docs.aws.amazon.com/mediapackage/latest/userguide/getting-started.html

API Reference: https://docs.aws.amazon.com/mediapackage/latest/apireference/what-is.html

Service Quotas: https://docs.aws.amazon.com/mediapackage/latest/userguide/quotas.html

AWS Elemental Services – AWS Elemental MediaTailor

General Information

AWS Elemental MediaTailor is a scalable ad insertion and channel assembly service that delivers targeted advertisements into live and Video on Demand (VOD) streams while maintaining broadcast quality. The service generates personalised manifest files for each viewer and creates linear streaming channels using existing content.

Main Benefits

AWS Elemental MediaTailor offers broadcast-quality server-side ad insertion that bypasses ad blockers while providing frame-accurate transitions and personalised targeting. The service scales automatically to handle up to 10,000 requests per second, reduces upfront commitments, and reduces operational complexity through automatic ad transcoding and beacon tracking.

Key Use Cases

- **Server-Side Ad Insertion:** Insert personalised advertisements into live and VOD streams with broadcast quality
- **Channel Assembly:** Create linear streaming channels using existing VOD content and live sources
- **Sports Event Monetisation:** Deliver targeted ads during live sports broadcasts with precise timing
- **Multi-Platform Ad Delivery:** Support cross-device ad insertion across web, mobile, TV, and set-top boxes.

Core Features

- **Server-Side Ad Insertion:** Inserts advertisements upstream before delivery, helping ensure playback without discrimination between content and ads
- **Channel Assembly:** Creates linear streaming channels using existing VOD content and live sources without touching content segments
- **Dynamic Transcoding:** Automatically matches ad quality and format to primary video content for optimal viewing experience
- **HLS Interstitials Support:** Enables insertion of interstitial advertisements directly into live streams using HLS Interstitials specification
- **Advanced Ad Tracking:** Provides server-side beaconing, beacon deduplication, adaptive throttling, and in-order beacon delivery
- **Time-Shifted Viewing:** Supports digital video recorder-like functions including start-over, pause, rewind, and fast-forward capabilities
- **Multi-Protocol Support:** Supports HLS, DASH, and CMAF streaming protocols for broad device compatibility
- **Personalisation:** Delivers targeted ads based on viewer demographics, interests, and location preferences.

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/mediatailor/
FAQ: https://aws.amazon.com/mediatailor/faqs/

Technical Overview

Backup and Recovery Capabilities

Backup: AWS Elemental MediaTailor operates as a manifest manipulation service without traditional backup capabilities since it does not store video content. Configuration data is maintained within AWS managed infrastructure with built-in redundancy. Customers can programmatically export configurations using APIs or Infrastructure as Code (IaC) tools for recreation purposes.

Recovery: Recovery is achieved through AWS multi-AZ infrastructure with automatic failover between Availability Zones (AZ). Cross-Region disaster recovery enables configuration recreation in alternate AWS Regions within minutes. The manifest-only architecture minimises data loss risks. Supports minimal RPO (configuration data quickly recreated) and low RTO objectives (automatic infrastructure failover).

Setup Process

Onboarding: Access MediaTailor console through AWS account and create playback configurations with content source URL prefix and ad decision server URL. Configure HLS or DASH compatible manifests, ad decision server integration, CDN integration, logging, and personalisation thresholds using comprehensive documentation and step-by-step guides.

Offboarding: Delete MediaTailor configurations, channels, and associated resources through AWS Management Console, CLI, or API. Clean up playbook configurations, channel assembly components, source locations, and VOD sources with configuration and metadata cleanup.

Data Management

- **Storage Options:** Manifest-only service storing configuration data in AWS managed infrastructure with built-in redundancy
- **Data Removal:** Delete configurations, channels, and resources through AWS Management Console, CLI, or API with metadata cleanup
- **Cross-Region Replication:** Export and recreate configurations programmatically in alternate AWS Regions using APIs or IaC
- **Encryption:** Requires HTTPS connections from well-known certificate authorities for secure manifest and content delivery.

Optimise Cost and Consumption

AWS Elemental MediaTailor offers included ad transcoding allowances covering the first 10 transcodes per 1,000 insertions at no charge. Channel assembly uses existing encoded content without re-transcoding to minimise processing costs. MediaTailor Vended Logs and Log Filtering enable selective log delivery to cost-effective Amazon S3 storage, reducing logging expenses. Strategic CDN cache configuration significantly reduces origin request costs and data transfer charges.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/mediatailor/latest/ug/what-is.html
Getting Started: https://docs.aws.amazon.com/mediatailor/latest/ug/getting-started.html
API Reference: https://docs.aws.amazon.com/mediatailor/latest/apireference/
Service Quotas: https://docs.aws.amazon.com/mediatailor/latest/ug/quotas.html

AWS Elemental Services – AWS End User Messaging

General Information

AWS End User Messaging allows businesses to communicate with end users across multiple channels including SMS, MMS, voice messages, push notifications, and WhatsApp messaging. The service provides application-to-person messaging capabilities with global scale and resiliency for delivering critical communications in applications.

Main Benefits

AWS End User Messaging provides a unified platform across more than 30 AWS Regions with built-in high availability and automatic scaling. Organisations benefit from a single API for all messaging channels, reducing operational overhead and helping achieve near-zero RPO and sub-hour RTO through multi-Region deployment strategies.

Key Use Cases

- **OTP Authentication:** One-time password and login verification for secure access
- **Marketing Campaigns:** Promotional messages with multimedia content and interactive features across channels
- **Critical Notifications:** Appointment reminders, delivery updates, and emergency alerts for timely communication
- **Customer Engagement:** Multi-channel messaging workflows combining SMS, push notifications, and WhatsApp

Core Features

- **SMS and MMS Messaging:** Global text and multimedia messaging with phone pools, sender IDs, and configuration sets
- **Voice Messaging:** Text-to-speech voice message delivery to phone numbers globally for audio communications
- **Push Notifications:** Multi-platform push notifications supporting Apple Push Notification Service (APNs), Firebase Cloud Messaging (FCM), Amazon Device Messaging (ADM), and Baidu Cloud Push services
- **WhatsApp Business Integration:** Built-in WhatsApp messaging with rich media, interactive buttons, and template support
- **Phone Number Management:** Automated provisioning of origination identities including short codes, long codes, and toll-free numbers
- **SMS Protect:** Advanced security with artificial traffic filtering and country-based blocking configurations for fraud prevention
- **Message Feedback Tracking:** Performance monitoring for conversion rates, click-through rates, and engagement metrics across channels
- **Event Integration:** Integration with Amazon CloudWatch, Amazon Simple Notification Service (Amazon SNS), and Amazon EventBridge for delivery receipts and real-time notifications

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://docs.aws.amazon.com/sms-voice/latest/userguide/registrations-sg-faq.html

FAQ: https://docs.aws.amazon.com/sms-voice/latest/userguide/registrations-sg-faq.html

Technical Overview

Backup & Recovery Capabilities

Backup: Customer configuration data, contact lists, and message templates can be backed up using Amazon S3 storage and Amazon DynamoDB for structured data. Message delivery logs are preserved in Amazon CloudWatch Logs per configured retention policies. Opt-out lists and phone number configurations can be exported manually via console or API.

Recovery: Multi-region architecture allows cross-Region failover with automated recovery mechanisms. Built-in redundancy through multi-Availability Zone deployment allows built-in high availability. Organisations can implement phone pool redundancy for resilient delivery. Supports customer RPO objectives with near-zero configuration data replication and RTO targets under 1 hour through automated failover.

Setup Process

Onboarding: Access through AWS Management Console, AWS CLI, or SDKs with sandbox environment for initial testing. Configure phone pools, origination identities, event destinations, and opt-out management. SMS simulators allow cost-free testing with automated progress alerts via Amazon EventBridge.

Offboarding: Delete messaging resources including phone numbers, sender IDs, and configuration sets through console or API. Export contact lists and message logs before termination. Service automatically handles cleanup with separate deletion required for Amazon CloudWatch Logs retention.

Data Management

- **Storage Options:** Use Amazon S3 for MMS media files and message archives, Amazon DynamoDB for contact management and opt-out lists
- **Data Removal:** Delete messaging resources via console or API including phone numbers, sender IDs, and configuration sets
- **Cross-Region Replication:** Message configurations and phone pools replicated across regions using Infrastructure as Code practices
- **Encryption:** Use AWS IAM to allow fine-grained access control and security policies for comprehensive data protection

Optimise Cost and Consumption

Phone pools allow automatic origination identity selection to optimise routing costs across regions. Message feedback tracking monitors conversion rates and click-through rates to reduce campaign waste and improve ROI. SMS Protect filters artificial traffic preventing fraud-related charges. Geographic routing strategies use local numbers to minimise international messaging costs. Character count management and message template optimisation reduce segmentation expenses while AWS Cost Management integration provides detailed cost analysis and budget alerts.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/sms-voice/latest/userguide/what-is-sms-mms.html

Getting Started: https://docs.aws.amazon.com/sms-voice/latest/userguide/getting-started.html

API Reference:
https://docs.aws.amazon.com/pinpoint/latest/apireference_smsvoicev2/Welcome.html

Service Quotas: https://docs.aws.amazon.com/general/latest/gr/end-user-messaging.html

AWS Elemental Services – AWS Entity Resolution

General Information

AWS Entity Resolution is a fully-managed service that matches, links, and enhances related records across multiple applications and data stores. It creates unified views using rule-based, ML, and provider-based matching techniques without requiring specialised expertise.

Main Benefits

AWS Entity Resolution eliminates the complexity of building custom entity matching systems with its no-code, fully-managed approach. It processes datasets from thousands to billions of records with automatic scaling, delivers near real-time matching within seconds, and requires no ML expertise.

Key Use Cases

- **Customer Identity Resolution:** Creating unified customer profiles by reducing duplicate identities across systems
- **Healthcare Patient Matching:** Building unified patient indexes linking records across multiple healthcare providers
- **Product Catalog Management:** Matching product information using Stock Keeping Unit and Universal Product Code codes for consistency
- **Fraud Detection:** Identifying duplicate accounts and verifying customer identities across financial institutions.

Core Features

- **Rule-Based Matching:** Ready-to-use deterministic rules with fuzzy algorithms including Levenshtein Distance and Soundex
- **ML Matching:** Pre-configured ML models analysing entire records with confidence scores and missing field handling
- **Near Real-Time APIs:** GetMatchId and GenerateMatchId APIs providing synchronous entity matching within seconds
- **Flexible Data Preparation:** Supports up to 20 data inputs with customisable schema mapping and normalisation
- **Data Service Provider Integration:** Licensed datasets and third-party provider integration for enhanced data enrichment capabilities
- **Automated Processing:** Manual bulk processing and automatic incremental processing for updated matches from Amazon Simple Storage Service (Amazon S3)
- **Built-in Data Protection:** AWS Key Management Service (AWS KMS) encryption support with data Regionalisation and secure PII hashing
- **Multi-Technique Matching:** Rule-based, ML, and provider-led matching approaches in single unified service.

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/entity-resolution/

FAQ: https://aws.amazon.com/entity-resolution/faqs/

Technical Overview

Backup and Recovery Capabilities

Backup: AWS Entity Resolution does not provide native backup capabilities. Data backup responsibilities lie with underlying Amazon S3 storage services using Amazon S3 versioning, replication, or AWS Backup policies. Workflow configurations can be recreated using AWS CloudFormation or programmatically via APIs.

Recovery: Workflows can be rapidly recreated and executed without infrastructure provisioning delays due to serverless architecture. Cross-Region disaster recovery requires Amazon S3 data replication and workflow recreation using Infrastructure as Code (IaC). Supports customer RPO (cross-Region S3 replication) and RTO (sub-second matching capabilities) objectives.

Setup Process

Onboarding: Setup begins by preparing data in Amazon S3 and using AWS Glue crawlers for catalogue entries. Create schema mappings via console, then configure matching workflows with pre-built templates and ready-to-use rules. Comprehensive SDKs support Java, Python, and JavaScript.

Offboarding: Delete Amazon S3 input/output buckets, Entity Resolution workflows, and schema mappings. Optionally remove AWS Glue Data Catalog entries. Intermediate processing data automatically cleans up after job completion.

Data Management

- **Storage Options:** Data stored in Amazon S3 with AWS Glue Data Catalog format required
- **Data Removal:** Delete Amazon S3 buckets, workflows, schema mappings; intermediate data automatically cleaned after processing
- **Cross-Region Replication:** Requires manual Amazon S3 data replication and workflow recreation using IaC
- **Encryption:** Built-in AWS KMS encryption support with data Regionalisation and secure PII hashing capabilities.

Optimise Cost and Consumption

Data preprocessing reduces record volumes before processing, while incremental processing prevents reprocessing unchanged data. Rule-based matching offers cost advantages over ML-based approaches for appropriate use cases. Amazon S3 storage class optimisation and lifecycle policies manage temporary processing data efficiently. Right-sizing processing jobs prevent unnecessary resource consumption, while regular monitoring of processing volumes and match rates identifies rule optimisation opportunities.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/entityresolution/latest/userguide/

Getting Started: https://docs.aws.amazon.com/entityresolution/latest/userguide/what-is-service.html

API Reference: https://docs.aws.amazon.com/entityresolution/latest/apireference/

Service Quotas: https://docs.aws.amazon.com/entityresolution/latest/userguide/quotas.html

AWS Fargate

General Information

AWS Fargate is a serverless compute engine for containers that runs containerised applications without managing servers or clusters. Compatible with Amazon Elastic Container Service (Amazon ECS) and Amazon Elastic Kubernetes Service (Amazon EKS), it provides task-level isolation with dedicated resources and automatic scaling capabilities.

Main Benefits

AWS Fargate reduces infrastructure management overhead, enabling teams to focus on application development. Features include task-level isolation for enhanced security, per-second billing, automatic scaling, and up to 70% savings with Fargate Spot for interrupt-tolerant workloads, delivering faster time-to-market.

Key Use Cases

- **Microservices Architecture:** Deploy and scale microservices without infrastructure management overhead
- **Batch Processing Workloads:** Run batch jobs and data processing tasks with automatic resource provisioning
- **Development and Test Environments:** Quickly spin up containerised environments for development and testing
- **Web Applications and APIs:** Host containerised web applications with seamless scaling based on demand

Core Features

- **Serverless Container Compute:** Run containers without managing servers or clusters, with automatic infrastructure provisioning
- **Task Isolation:** Each task has its own isolation boundary with dedicated CPU, memory, and network resources
- **Automatic Scaling:** Scale applications based on demand without manual intervention or capacity planning
- **Fargate Spot:** Run interrupt-tolerant workloads at up to 70% discount using spare AWS compute capacity
- **Platform Version Support:** Support for Amazon Linux 2, Bottlerocket, and Windows 2019 Server platforms
- **Load Balancer Integration:** Integration with Application Load Balancer and Network Load Balancer
- **Per-Second Billing:** Pay only for resources consumed with per-second billing granularity
- **ECS and EKS Compatibility:** Compatible with both Amazon ECS and Amazon EKS.

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/fargate/

FAQ: https://aws.amazon.com/fargate/faqs/

Technical Overview

Backup and Recovery Capabilities

Backup: AWS Fargate does not provide native backup capabilities since containers are ephemeral. Data backup must be implemented at the application level using external AWS storage services like Amazon Elastic File System (Amazon EFS), Amazon Simple Storage Service, or Amazon Relational Database Service. AWS Backup integration is available for supported resources like Amazon EFS file systems.

Recovery: Disaster recovery is achieved through multi-Region deployments, container image replication across AWS Regions, and Infrastructure as Code (IaC) practices. Fargate supports automatic failover through Application Load Balancer health checks and rapid task replacement across Availability Zones. Supports low RTO objectives with task startup within minutes and RPO objectives through application-level data replication strategies.

Setup Process

Onboarding: Create Fargate-enabled Amazon ECS or Amazon EKS clusters through AWS Management Console, AWS CLI, or IaC tools. Configure AWS Identity and Access Management (AWS IAM) execution roles, define task definitions with CPU/memory requirements, and set up Virtual Private Cloud networking with security groups.

Offboarding: Stop running tasks and services, delete task definitions, and remove associated resources like load balancers. Clean up AWS IAM roles, Amazon CloudWatch log groups, and delete persistent storage resources like Amazon EFS file systems.

Data Management

- **Storage Options:** Ephemeral container storage with Amazon EFS for persistent storage; Amazon EBS volumes cannot be mounted
- **Data Removal:** Container data automatically deleted when tasks stop; persistent data requires manual deletion of external services
- **Cross-Region Replication:** Supported through container image replication across Regions and multi-Region deployment patterns for disaster recovery
- **Encryption:** Task-level isolation with dedicated resources; encryption depends on integrated AWS services like Amazon EFS and Amazon Elastic Container Registry.

Optimise Cost and Consumption

AWS Fargate Spot delivers up to 70% savings on interrupt-tolerant workloads using spare compute capacity. Per-second billing with no minimum commitments helps ensure precise resource consumption tracking. AWS Compute Optimiser provides right-sizing recommendations to identify overprovisioned tasks. Compute Savings Plans offer discounted rates for consistent usage patterns, while Graviton2-powered instances deliver enhanced price-performance ratios for optimised workload efficiency.

Technical Resources

Developer Guide:
https://docs.aws.amazon.com/AmazonECS/latest/developerguide/AWS_Fargate.html

Getting Started: https://docs.aws.amazon.com/AmazonECS/latest/developerguide/getting-started-fargate.html

API Reference: https://docs.aws.amazon.com/AmazonECS/latest/APIReference/

Service Quotas: https://docs.aws.amazon.com/AmazonECS/latest/developerguide/service-quotas-manage.html

AWS Fault Injection Service (AWS FIS)

General Information

AWS Fault Injection Service (AWS FIS) is a fully managed chaos engineering service that performs fault injection experiments on AWS workloads. It creates controlled disruptive events like instance failures and network disruptions to test application resilience and validate recovery procedures.

Main Benefits

AWS FIS reduces the complexity of building custom chaos engineering tools while providing production-safe experimentation with built-in guardrails and automatic rollback capabilities. The service reduces MTTR, prevents costly outages, and offers pay-per-use pricing that scales with experiment duration rather than requiring upfront infrastructure investment.

Key Use Cases

- **Application Resilience Testing:** Inject faults like instance failures and network disruptions to validate recovery procedures
- **Chaos Engineering Implementation:** Build confidence in system reliability through regular testing under real-world failure conditions
- **Disaster Recovery Validation:** Simulate regional failures and service outages to verify backup systems and RTO objectives
- **Multi-Account Testing:** Run fault injection experiments across multiple AWS accounts from centralised orchestrator account.

Core Features

- **Preconfigured Actions:** Ready-to-use fault injection activities for Amazon Elastic Compute Cloud, Amazon Elastic Container Service, Amazon Elastic Kubernetes Service, Amazon Relational Database Service, AWS Lambda, Amazon Simple Storage Service and other AWS services
- **Stop Conditions:** Amazon CloudWatch alarm-based safety mechanisms that automatically halt experiments when predefined thresholds are exceeded
- **Multi-Account Experiments:** Run fault injection experiments across multiple AWS accounts from centralised orchestrator with cross-account AWS Identity and Access Management (AWS IAM) roles
- **Systems Manager Integration:** Execute custom fault injection scripts and automation documents for complex multi-step experimental scenarios
- **Experiment Templates:** Reusable blueprints defining actions, targets, stop conditions and configurations for consistent testing scenarios
- **Automatic Rollback:** Built-in guardrails and automatic rollback capabilities ensure production-safe experimentation with comprehensive safety controls
- **Fine-Grained Targeting:** Precise resource targeting using AWS resource tags to minimise unnecessary resource consumption during testing
- **Deep AWS Integration:** Integration with compute, storage, database and networking components across the entire AWS environment.

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/fis/
FAQ: https://aws.amazon.com/fis/faqs/

Technical Overview

Backup and Recovery Capabilities

Backup: AWS FIS does not provide traditional backup capabilities as it is a fault injection testing service rather than a data storage service. The service does not integrate with AWS Backup since it operates on existing resources to test their resilience. Experiment data is automatically retained for 120 days before deletion.

Recovery: AWS FIS serves as a critical tool for testing disaster recovery scenarios by simulating regional failures, Availability Zone disruptions, and service outages to validate recovery mechanisms. The service helps customers validate RPO and RTO objectives through controlled chaos experiments that measure actual recovery times and verify backup restoration procedures.

Setup Process

Onboarding: Create AWS IAM roles with necessary permissions and experiment templates through AWS Management Console, CLI, or APIs. Start with non-production environments, implement Amazon CloudWatch alarms as stop conditions, and gradually expand scope using pre-built templates.

Offboarding: Delete experiment templates and ensure no active experiments are running. Remove IAM roles and cleanup experiment artifacts. Experiment data automatically deletes after 120 days retention period.

Data Management

- **Storage Options:** Stores experiment logs and templates regionally with 120-day automatic retention period per Region
- **Data Removal:** Experiment data automatically deleted after 120 days; manual deletion of templates and Amazon CloudWatch alarms available
- **Cross-Region Replication:** Data stored Regionally without cross-Region replication; experiments must be configured per target AWS Region
- **Encryption:** Standard AWS encryption applies to stored experiment data and templates following Regional security practices.

Optimise Cost and Consumption

Experiment templates provide reusable blueprints that reduce repetitive setup costs across multiple testing scenarios. Fine-grained resource targeting using AWS tags minimises unnecessary resource consumption during fault injection experiments. Shorter experiment durations optimise action-minute billing charges. Multi-account centralised orchestration consolidates chaos engineering activities, reducing operational overhead across teams while enabling shared learning and standardised testing practices.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/fis/latest/userguide/what-is.html
Getting Started: https://docs.aws.amazon.com/fis/latest/userguide/getting-started-planning.html
API Reference: https://docs.aws.amazon.com/fis/latest/APIReference/
Service Quotas: https://docs.aws.amazon.com/fis/latest/userguide/fis-quotas.html

AWS Firewall Manager

General Information

AWS Firewall Manager is a centralised security management service that simplifies administration of firewall protections across multiple AWS accounts and resources. It automatically applies consistent security policies organisation-wide for AWS WAF (Web Application Firewall), AWS Shield Advanced, AWS Network Firewall, Amazon Virtual Private Cloud (Amazon VPC) security groups, and Amazon Route 53 Domain Name System (DNS) Firewall.

Main Benefits

AWS Firewall Manager reduces security management complexity by automatically applying consistent policies across up to 2,500 accounts per policy scope. The service reduces administrative overhead and human error while supporting hierarchical rule management that balances centralised control with account-specific customisations for optimal organisational security governance.

Key Use Cases

- **Centralised Security Policy Management:** Deploy consistent firewall rules across multiple AWS accounts from single point
- **Compliance Enforcement:** Automatically monitor new resources and ensure mandatory security policy compliance from deployment
- **Multi-layer Protection Orchestration:** Coordinate AWS WAF, AWS Shield Advanced, AWS Network Firewall, and VPC security groups organisation-wide
- **Security Audit and Remediation:** Audit existing configurations for over-permissive rules and automatically remediate non-compliance.

Core Features

- **Multi-Service Protection Management:** Centrally manages AWS WAF, AWS Shield Advanced, VPC security groups, AWS Network Firewall, and DNS Firewall
- **Cross-Account Policy Enforcement:** Automatically applies security policies across multiple AWS accounts within AWS Organizations structure
- **Automatic Resource Protection:** Automatically applies protections to new resources as created, helping ensure continuous compliance
- **Hierarchical Rule Management:** Supports organisational rules while permitting account-specific customisations for balanced control and flexibility
- **Compliance Dashboard:** Provides centralised visibility into compliance status with notifications for non-compliant resources and monitoring
- **Third-Party Firewall Integration:** Supports integration with third-party solutions from AWS Marketplace including Palo Alto Networks
- **AWS Config Integration:** Requires AWS Config for continuous resource monitoring and compliance tracking across accounts
- **Centralised Security Findings:** Integrates with AWS Security Hub for centralised security findings management and oversight.

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/firewall-manager/
FAQ: https://aws.amazon.com/firewall-manager/faqs/

Technical Overview

Backup and Recovery Capabilities

Backup: AWS Firewall Manager does not provide built-in backup capabilities as it is a security policy management service rather than data storage. Policy configurations should be backed up through Infrastructure as Code (IaC) practices, version control systems, or documented configurations that can be recreated manually via AWS Management Console or APIs.

Recovery: Disaster recovery relies on AWS infrastructure resilience and the ability to recreate policies in alternate AWS Regions. The service operates regionally, requiring separate policies per AWS Region. Organisations can achieve low RTO by maintaining policy configurations in IaC and using automatic policy application for rapid security restoration.

Setup Process

Onboarding: Requires AWS Organizations with all features enabled and designated Firewall Manager administrator account. AWS Config must be enabled across all member accounts for continuous monitoring. Create security policies by selecting protection type, configuring rules, defining scope, and choosing automatic or manual remediation modes.

Offboarding: Delete policies using DeletePolicy API with DeleteAllPolicyResources parameter set to true for complete cleanup. DisassociateThirdPartyFirewall API removes third-party firewall integrations. Customer data is removed when policies are deleted.

Data Management

- **Storage Options:** Maintains policy configurations and compliance state information as a management service, not data storage
- **Data Removal:** DeletePolicy API with DeleteAllPolicyResources parameter removes policies and associated resources completely upon deletion
- **Cross-Region Replication:** Operates regionally requiring separate policies per Region; no automatic cross-Region policy replication capabilities
- **Encryption:** Relies on underlying AWS infrastructure encryption; focuses on security policy management rather than data protection

Optimise Cost and Consumption

AWS Shield Advanced customers receive Firewall Manager WAF policies at no additional charge, providing immediate cost benefits. Organisations can optimise spending by carefully scoping policies to avoid unnecessary resource inclusion and using manual remediation mode during testing phases. Consolidating similar protection requirements into fewer policies reduces management overhead while using AWS Marketplace managed rules provides cost-effective security coverage with economies of scale.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/waf/latest/developerguide/fms-chapter.html
Getting Started: https://docs.aws.amazon.com/waf/latest/developerguide/fms-prereq.html
API Reference: https://docs.aws.amazon.com/fms/2018-01-01/APIReference/Welcome.html
Service Quotas: https://docs.aws.amazon.com/waf/latest/developerguide/fms-limits.html

AWS Global Accelerator

General Information

AWS Global Accelerator is a networking service that improves application performance and availability by routing traffic through the AWS global network. It provides static IP addresses as fixed entry points and automatically directs users to optimal endpoints based on performance and health policies.

Main Benefits

AWS Global Accelerator delivers up to 60% performance improvement through its global network of over 113 Points of Presence. It provides static IP addresses, instant failover capabilities across multiple AWS Regions and Availability Zones, and works with both TCP and UDP protocols (including non-HTTP traffic), offering deterministic routing with fast regional failover compared to Domain Name System (DNS)-based solutions that depend on DNS caching.

Key Use Cases

- **Latency-Sensitive Applications:** Gaming, media, mobile apps, ad-tech, and financial applications requiring reduced latency and jitter
- **Multi-Region Disaster Recovery:** Detects endpoint failures and instantly reroutes traffic to available AWS Regions for business continuity
- **Global Application Scaling:** IoT, retail, media, automotive, and healthcare applications requiring global reach without frequent updates
- **Custom Routing Applications:** Voice over Internet Protocol, online gaming, and EdTech applications needing custom logic for specific endpoint mapping.

Core Features

- **Static Anycast IP Addresses:** Provides two static IPv4 and IPv6 addresses as fixed entry points with Bring Your Own IP (BYOIP) support
- **Global Performance-Based Routing:** Routes TCP and UDP traffic to healthy endpoints based on user location and health
- **Health Monitoring:** Continuously monitors endpoint health using TCP, HTTP, and HTTPS checks with automatic traffic redirection
- **Traffic Control:** Fine-grained control with traffic dials and endpoint weights for load distribution across AWS Regions
- **Distributed Denial-of-Service (DDoS) Protection:** Protected by AWS Shield Standard with optional AWS Shield Advanced for enhanced protection and support
- **Custom Routing Accelerators:** Enables custom application logic to route traffic to specific Amazon Elastic Compute Cloud instance destinations
- **Client Affinity:** Consistently routes users to the same endpoint for applications requiring stateful connections
- **Fault-Tolerant Design:** Two static IP addresses serviced by independent network zones with automatic failover capabilities.

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/global-accelerator/
FAQ: https://aws.amazon.com/global-accelerator/faqs/

Technical Overview

Backup and Recovery Capabilities

Backup: AWS Global Accelerator configurations can be backed up through Infrastructure as Code (IaC) practices using AWS CloudFormation templates or API scripts (AWS Management Console or API). Static IP addresses are maintained while accelerators exist but permanently lost upon deletion. Customers should maintain configuration backups through code repositories and automation scripts.

Recovery: Provides instant traffic failover to healthy endpoints within seconds, automated network-level failover, and integration with AWS Application Recovery Controller for coordinated multi-Region management. Supports near-zero RTO (instant failover capabilities) and works with underlying services to minimise RPO (data loss during failover scenarios).

Setup Process

Onboarding: Setup involves three steps through AWS Management Console or API from US West (Oregon) Region: create accelerator with routing type selection, configure listeners and endpoint groups, then register endpoints with weights and health checks. Optional BYOIP support available.

Offboarding: Disable accelerator through AWS Management Console or API, remove listeners and endpoint groups before deletion. Static IP addresses are permanently lost upon deletion. Cross-account attachments can be deleted separately with irreversible cleanup process.

Data Management

- **Storage Options:** Network routing service with configuration storage through CloudFormation templates and API scripts for IaC
- **Data Removal:** Disable accelerator, remove listeners and endpoint groups before deletion; static IP addresses permanently lost upon deletion
- **Cross-Region Replication:** Automatically routes traffic across multiple AWS Regions based on performance, health, and endpoint policies
- **Encryption:** Traffic routing service with built-in DDoS protection via AWS Shield Standard and optional Shield Advanced.

Optimise Cost and Consumption

Traffic dials enable fine-grained control over traffic percentage distribution across AWS Regions for gradual shifting and cost management. Strategic endpoint placement in lower-cost AWS Regions combined with traffic weights optimises load distribution efficiency. BYOIP functionality reduces IP address expenses. Automatic traffic routing to nearest endpoints minimises data transfer costs while performance improvements reduce compute requirements and infrastructure overhead across endpoints.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/global-accelerator/latest/dg/what-is-global-accelerator.html
Getting Started: https://docs.aws.amazon.com/global-accelerator/latest/dg/getting-started.html
API Reference: https://docs.aws.amazon.com/global-accelerator/latest/api/Welcome.html
Service Quotas: https://docs.aws.amazon.com/global-accelerator/latest/dg/limits-global-accelerator.html

AWS Glue

General Information

AWS Glue is a serverless data integration service that simplifies data discovery, preparation, movement, and integration for analytics and machine learning. It supports over 70 data sources with automated Extract, Transform, and Load (ETL) capabilities and centralised cataloguing.

Main Benefits

AWS Glue reduces infrastructure management overhead through its serverless architecture. AWS Glue DataBrew reduces data preparation development by up to 80% using visual interfaces and over 250 built-in transformations. The service delivers significant cost reductions compared to traditional ETL platforms, with customer case studies showing 20-80% cost savings. AWS Glue supports over 70 diverse data sources with over 100 native and AWS Marketplace connectors available, featuring automatic scaling and pay-per-use pricing with no upfront commitments.

Key Use Cases

- **Data Lake Creation:** Building modern data platforms with unified metadata repository and automated discovery
- **ETL Pipeline Development:** Creating complex batch, micro-batch, and streaming data processing pipelines with visual interfaces
- **Data Warehousing Analytics:** Preparing and loading data from various sources into warehouses like Amazon Redshift
- **Machine Learning Preparation:** Cleaning, transforming, and preparing data for ML workflows with built-in quality features.

Core Features

- **Data Catalog:** Persistent metadata store with automatic schema discovery, version history, and centralised data asset repository
- **AWS Glue Studio:** Visual drag-and-drop interface for creating, running, and monitoring data integration jobs without coding
- **Crawlers and Classifiers:** Automatically scan repositories, classify data, extract schema information, and populate metadata catalog
- **ETL Jobs System:** Managed infrastructure running Scala or PySpark scripts with automatic scaling and event-driven triggers
- **Streaming ETL:** Real-time data processing using Apache Spark Structured Streaming for continuous data transformation capabilities
- **DataBrew:** No-code visual data preparation offering over 250 ready-made transformations for cleaning and normalisation
- **Data Quality Management:** Built-in quality rules, anomaly detection, and automated cleansing powered by ML algorithms
- **Generative AI Integration:** AI-powered assistance for ETL code generation, Apache Spark job modernisation, and troubleshooting

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/glue/

FAQ: https://aws.amazon.com/glue/faqs/

Technical Overview

Backup and Recovery Capabilities

Backup: Data Catalog metadata is automatically backed up with built-in redundancy and version control. Job definitions, crawler configurations, and connection details are stored with high durability. Additional backup can be created manually through APIs or automated via AWS CloudFormation templates.

Recovery: Serverless multi-AZ architecture enables rapid job restart and automatic scaling without infrastructure provisioning delays. Cross-Region disaster recovery through metadata replication to secondary Regions. Supports customer RPO objectives through automatic metadata backup and version control, and RTO objectives through serverless architecture enabling quick failover.

Setup Process

Onboarding: Setup AWS Identity and Access Management (AWS IAM) permissions and access AWS Glue console or Studio. Create default database in Data Catalog, configure network access to data sources, and optionally setup encryption. Begin with crawlers to discover data, then create ETL jobs using visual or code-based approaches.

Offboarding: Delete ETL jobs, crawlers, and metadata from Data Catalog through AWS Management Console or API. Remove table definitions, connection configurations, and CloudWatch logs separately. Complete cleanup requires removing AWS IAM roles, security groups, and Virtual Private Cloud (VPC) endpoints.

Data Management

- **Storage Options:** Data Catalog metadata stored with built-in redundancy and high durability across multiple Availability Zones
- **Data Removal:** Delete ETL jobs, crawlers, and metadata through AWS Management Console or API; Amazon CloudWatch logs require separate deletion
- **Cross-Region Replication:** Manual setup required through APIs or infrastructure as code for cross-Region metadata catalog replication
- **Encryption:** AWS Key Management Service integration for encryption with optional setup during onboarding and VPC network isolation support

Optimise Cost and Consumption

AWS Glue's autoscaling capabilities dynamically adjust DPU allocation based on workload demands while flexible job execution enables scheduling non-urgent workloads during off-peak periods. Job bookmarks process only new data to reduce processing time and costs. Right-sizing through appropriate worker types (G.1X, G.2X, G.025X) optimises resource allocation. Data partitioning with efficient file formats like Parquet minimises processing overhead and consumption costs.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/glue/latest/dg/what-is-glue.html

Getting Started: https://docs.aws.amazon.com/glue/latest/dg/setting-up.html

API Reference: https://docs.aws.amazon.com/glue/latest/webapi/

Service Quotas: https://docs.aws.amazon.com/general/latest/gr/glue.html

AWS Ground Station

General Information

AWS Ground Station is a fully managed service that enables customers to control satellite communications and process satellite data without building ground station infrastructure. The service provides secure, fast satellite communications using AWS-managed antennas across a global network.

Main Benefits

AWS Ground Station delivers up to 80% cost savings compared to traditional ground stations by reducing infrastructure investment and maintenance overhead. The service processes satellite data within seconds using a low-latency global fibre network, with pay-per-minute pricing and AWS environment integration.

Key Use Cases

- **Real-time Data Downlinking:** Downlink satellite data in real-time and store directly in Amazon Simple Storage Service (Amazon S3)
- **Satellite Command and Control:** Uplink commands and perform Telemetry, Tracking, and Command functions for missions
- **Earth Observation and Remote Sensing:** Process satellite imagery for weather monitoring, agriculture, and environmental analysis
- **Defence and Government Operations:** Support secure, scalable ground segment solutions for national security applications.

Core Features

- **Global Ground Station Network:** AWS-managed antennas and ground stations integrated across multiple regions worldwide
- **Multiple Frequency Band Support:** Supports X-band, S-band, and Ku-band satellite protocols with various modulation schemes
- **Pay-per-Minute Pricing:** Purchase antenna time in one-minute increments with no upfront costs or commitments
- **Wideband Digital Intermediate Frequency:** High-bandwidth signal processing capabilities for advanced satellite communications at select sites
- **Digital Twin Testing:** Test contact orchestration workflows against virtual ground stations for cost-effective validation
- **Automatic Contact Scheduling:** Automated satellite communication session scheduling and management with 30 days advance booking
- **Customer-Provided Ephemerides:** Support for custom ephemeris data upload for antenna targeting during operations
- **AWS Service Integration:** Direct integration with Amazon S3, Amazon Elastic Compute Cloud (Amazon EC2), AWS Lambda, Amazon SageMaker for comprehensive data processing

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/ground-station/

FAQ: https://aws.amazon.com/ground-station/faqs/

Technical Overview

Backup and Recovery Capabilities

Backup: AWS Ground Station integrates with AWS Backup for protecting associated Amazon EC2 instances and Amazon Elastic Block Store volumes. Satellite data stored in Amazon S3 uses S3 versioning, cross-region replication, and Amazon S3 Glacier for long-term archival. Mission configurations should be backed up using Infrastructure as Code (IaC) practices.

Recovery: Mission profiles and configurations can be recreated in alternate regions using IaC. The distributed ground station network provides built-in redundancy through multiple locations and automated failover mechanisms. Supports near-zero RPO through cross-region data replication and improved RTO via multiple communication paths.

Setup Process

Onboarding: Setup requires AWS account creation with AWS Identity and Access Management permissions, satellite registration through data collection and technical validation, spectrum licensing, and integration testing. Create mission profiles with tracking configurations, configure Virtual Private Cloud, Amazon S3, Amazon EC2 resources, and use digital twin testing for validation.

Offboarding: Cancel scheduled contacts, delete mission profiles and satellite configurations, remove dataflow endpoint groups. Delete satellite data from Amazon S3 buckets and terminate Amazon EC2 instances through AWS Management Console, CLI, or APIs with no contract termination required.

Data Management

- **Storage Options:** Direct integration with Amazon S3 for automatic data storage, archiving, and Amazon S3 Glacier for long-term archival
- **Data Removal:** Delete satellite data from Amazon S3 buckets and terminate Amazon EC2 instances via AWS console, CLI, or APIs
- **Cross-Region Replication:** Automatic satellite data replication across AWS Regions using Amazon S3 cross-Region replication for disaster recovery
- **Encryption:** AWS Key Management Service provides encryption services for data security with built-in security and compliance capabilities.

Optimise Cost and Consumption

Pay-per-minute pricing enables granular cost control by purchasing antenna time in one-minute increments. Reserved scheduling provides discounted rates for regular users with monthly commitments over 12 months. Digital twin testing allows cost-effective workflow validation without consuming actual antenna time. Off-peak contact scheduling and narrow-band communications reduce operational costs, while the global network enables selection of the most cost-effective ground station locations for optimised data paths.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/ground-station/latest/ug/index.html

Getting Started: https://docs.aws.amazon.com/ground-station/latest/ug/getting-started.html

API Reference: https://docs.aws.amazon.com/ground-station/latest/APIReference/Welcome.html

Service Quotas: https://docs.aws.amazon.com/general/latest/gr/gs.html

AWS Health Dashboard

General Information

AWS Health Dashboard provides ongoing visibility into the performance and availability of AWS resources and services. It offers real-time alerts for both account-specific events affecting resources and public AWS service issues. The service enables proactive incident management through automated Amazon EventBridge integration.

Main Benefits

AWS Health Dashboard provides authoritative health information directly from AWS at no additional cost, reducing the need for third-party monitoring solutions. It offers 90 days of event history, real-time notifications, and deep Amazon EventBridge integration for automated incident response workflows, accelerating troubleshooting and reducing operational overhead.

Key Use Cases

- **Proactive Incident Management:** Monitor AWS Health events to identify potential issues before they impact users
- **Scheduled Maintenance Planning:** View upcoming AWS service changes and maintenance windows to prepare applications
- **Multi-Account Organisational Monitoring:** Aggregate health events across all AWS accounts for centralised organisational visibility
- **Automated Response Integration:** Trigger AWS Lambda functions and notifications through EventBridge for incident management workflows

Core Features

- **Account Events Dashboard:** Displays events specific to your AWS account including open, recent, and scheduled changes
- **Public Events View:** Shows general AWS service issues and interruptions not specific to individual accounts
- **Calendar View:** Provides monthly calendar format for viewing AWS Health events and scheduled changes
- **Event Filtering:** Filter events by status, start time, AWS Region, service, and other parameters for focused monitoring
- **Resource Export:** Download affected resources list in CSV or JSON format for external analysis
- **EventBridge Integration:** Automatically trigger actions and notifications through Amazon EventBridge rules for automated responses
- **Organisational View:** Centralised view of health events across all accounts in AWS Organizations for enterprise visibility
- **Timeline View:** Historical view of events and issues that have affected resources over time.

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://docs.aws.amazon.com/health/latest/ug/what-is-aws-health.html
FAQ: https://docs.aws.amazon.com/health/latest/ug/what-is-aws-health.html

Technical Overview

Backup and Recovery Capabilities

Backup: AWS Health Dashboard does not provide traditional backup capabilities as it is a monitoring service displaying real-time health events rather than storing persistent data. The service maintains 90 days of event history for reference. Customers can export affected resources data in CSV or JSON format for external backup purposes if needed for compliance requirements.

Recovery: AWS Health Dashboard does not support traditional recovery operations but assists with disaster recovery planning by monitoring health events and integrating with Amazon EventBridge to trigger automated responses. The service is highly available across AWS Regions. Supports customer RPO objectives through real-time visibility and RTO objectives through automated incident response workflows via Amazon EventBridge integration.

Setup Process

Onboarding: AWS Health Dashboard requires no setup and is immediately accessible through the AWS Management Console with any AWS account. Configure Amazon EventBridge rules for automated responses and enable organisational view for multi-account visibility. Advanced API features require AWS Business or AWS Enterprise Support plans.

Offboarding: No explicit offboarding process required as the service displays real-time status information without storing persistent customer data. Events automatically purge after 90 days through standard retention policy with no manual cleanup needed.

Data Management

- **Storage Options:** Displays real-time health events without storing persistent data, maintains 90 days of event history
- **Data Removal:** Events automatically purged after 90 days through standard retention policy, no manual deletion required
- **Cross-Region Replication:** Highly available across AWS Regions, some events appear in US East Virginia by default
- **Encryption:** Standard AWS encryption applied to health event data transmitted through the service infrastructure.

Optimise Cost and Consumption

Implement efficient Amazon EventBridge rule filtering to minimise unnecessary downstream service triggers, reducing AWS Lambda, Amazon Simple Notification Service (SNS), and Amazon Simple Queue Service (Amazon SQS) costs. The organisational view centralises multi-account health monitoring without additional charges, reducing operational overhead. Event filtering by status, AWS Region, and service parameters focuses monitoring resources on relevant alerts, optimising notification workflows and preventing costly downtime through proactive maintenance planning and early issue detection.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/health/latest/ug/index.html
Getting Started: https://docs.aws.amazon.com/health/latest/ug/getting-started-health-dashboard.html
API Reference: https://docs.aws.amazon.com/health/latest/APIReference/Welcome.html
Service Quotas: https://docs.aws.amazon.com/health/latest/ug/index.html

AWS HealthImaging

General Information

AWS HealthImaging is a Health Insurance Portability and Accountability Act (HIPAA)-eligible cloud service that enables healthcare customers to store, analyse, and share medical images at petabyte scale. It provides cloud-native Picture Archiving and Communication System (PACS) capabilities with DICOM P10 support and APIs for low-latency image retrieval.

Main Benefits

AWS HealthImaging delivers industry-leading performance with HTJ2K encoding that is an order of magnitude faster than JPEG2000. It reduces infrastructure management overhead while providing petabyte-scale storage, subsecond image retrieval, and automatic cost optimisation through intelligent tiering.

Key Use Cases

- **Enterprise Imaging:** Store and stream medical imaging data with low-latency performance and high availability
- **Long-term Image Archival:** Save costs on archival while maintaining subsecond image retrieval access
- **AI/ML Development:** Run AI and ML inference over imaging archives
- **Multimodal Analysis:** Combine imaging data with AWS HealthLake and AWS HealthOmics for precision medicine insights.

Core Features

- **HTJ2K Image Decoding:** SIMD-accelerated encoding delivers order-of-magnitude faster performance than JPEG2000
- **Developer-Friendly DICOM Metadata:** Returns metadata with human-friendly keywords rather than hexadecimal numbers
- **Pixel Data Verification:** Built-in verification using International Federation of Red Cross and Red Crescent Societies checksums helps ensure accuracy during import process
- **Scalable DICOM Imports:** Cloud-native parallel import capabilities for multiple studies without impacting workloads
- **Service-Managed Data Hierarchy:** Automatically organises DICOM P10 data by Patient, Study, and Series levels
- **DICOMweb API Compatibility:** Offers conformant APIs for existing applications plus cloud-native APIs
- **Intelligent Tiering:** Automatic movement between access tiers based on usage patterns to optimise costs
- **Petabyte-Scale Storage:** Cloud-native PACS capabilities supporting DICOM P10 format with low-latency retrieval.

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://docs.aws.amazon.com/healthimaging/latest/devguide/what-is.html

FAQ: https://docs.aws.amazon.com/healthimaging/latest/devguide/what-is.html

Technical Overview

Backup and Recovery Capabilities

Backup: AWS HealthImaging relies on AWS's underlying infrastructure for data durability with automatic storage across multiple Availability Zones. Customers can implement backup strategies using copy and export capabilities combined with Amazon Simple Storage Service (Amazon S3) cross-Region replication. Built-in pixel data verification helps ensure data integrity.

Recovery: Supports subsecond image retrieval access from both Amazon S3 Frequent Access and Archive Instant Access tiers. Multi-AZ deployment and progressive resolution capabilities enable recovery. Customers can implement cross-Region disaster recovery strategies using AWS HealthImaging APIs and Amazon EventBridge integration for automated failover workflows. Supports low RTO and RPO objectives.

Setup Process

Onboarding: Setup involves AWS account creation, AWS Identity and Access Management (AWS IAM) configuration, creating Amazon S3 buckets for data staging, establishing AWS HealthImaging data stores, and configuring AWS IAM roles. Access via AWS Management Console, CLI, or SDKs with comprehensive documentation and sample projects provided.

Offboarding: Data removal through AWS Management Console, CLI, or APIs with systematic cleanup processes. Customers can delete specific image sets or entire data stores using AWS SDKs while maintaining healthcare compliance requirements.

Data Management

- **Storage Options:** Intelligent tiering with Amazon S3 Frequent Access and Archive Instant Access tiers for automatic cost optimisation
- **Data Removal:** Delete image sets and data stores through AWS Management Console, CLI, or APIs with systematic cleanup processes
- **Cross-Region Replication:** Implement cross-Region strategies using HealthImaging APIs combined with Amazon S3 cross-Region replication capabilities
- **Encryption:** HIPAA-eligible service with built-in security features and AWS underlying infrastructure encryption support.

Optimise Cost and Consumption

AWS HealthImaging's intelligent tiering automatically moves image sets between Frequent Access and Archive Instant Access tiers based on 30-day access patterns, reducing storage costs while maintaining millisecond access latencies. HTJ2K compression efficiency reduces storage requirements, and indexed storage billing separates metadata costs from image storage costs for granular optimisation.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/healthimaging/latest/devguide/what-is.html

Getting Started: https://docs.aws.amazon.com/healthimaging/latest/devguide/getting-started.html

API Reference: https://docs.aws.amazon.com/healthimaging/latest/APIReference/index.html

Service Quotas: https://docs.aws.amazon.com/healthimaging/latest/devguide/endpoints-quotas.html

AWS HealthLake

General Information

AWS HealthLake is a Health Insurance Portability and Accountability Act (HIPAA)-eligible, serverless service that stores, searches, and analyses healthcare data using FHIR R4 specification. It automatically extracts clinical insights through natural language processing and transforms unstructured medical data into analytics-ready formats for modern healthcare applications.

Main Benefits

AWS HealthLake reduces healthcare data infrastructure complexity while ensuring HIPAA, HITRUST CSF, and System and Organization Controls (SOC) 2 compliance. It delivers sub-second query latency, zero-ETL analytics with petabyte-scale capabilities, and automatic data standardisation using ICD-10-CM and RxNorm, accelerating time-to-market for healthcare applications.

Key Use Cases

- **Population Health Analytics:** Analyse healthcare data to identify trends and support value-based care initiatives
- **Clinical Decision Support:** Build AI-powered patient profiles with comprehensive histories and evidence-based insights
- **Healthcare Data Integration:** Unify clinical data across health systems using standardised FHIR formats
- **Research and Clinical Trials:** Support participant screening, longitudinal analysis, and research data management.

Core Features

- **High-Performance FHIR Datastore:** Enterprise-grade FHIR R4 repository with sub-second latency and seamless data ingestion
- **Integrated Natural Language Processing:** HIPAA-eligible Natural Language Processing (NLP) libraries extracting medical entities, conditions, medications from unstructured text
- **Zero-ETL Analytics:** Direct SQL querying through Amazon Athena with automatic FHIR-to-SQL transformation at petabyte scale
- **FHIR API Infrastructure:** Comprehensive FHIR search capabilities and transactional server supporting standard FHIR operations
- **SMART on FHIR Authorisation:** Support for medical applications authorisation alongside AWS Signature Version 4
- **Real-time Subscriptions:** FHIR Subscriptions monitoring data changes through Amazon EventBridge or REST-Hook endpoints
- **Data Standardisation:** Automatic standardisation using medical nomenclatures like ICD-10-CM and RxNorm for consistency
- **Regulatory Compliance:** HIPAA eligibility, HITRUST CSF certification, and SOC 2 compliance for healthcare requirements.

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/healthlake/

FAQ: https://aws.amazon.com/healthlake/faqs/

Technical Overview

Backup and Recovery Capabilities

Backup: FHIR data export functionality to Amazon Simple Storage Service (Amazon S3) with built-in regional replication and encryption at rest using customer-managed keys can be created manually (AWS Management Console, CLI, APIs) or automated via AWS Lambda functions with custom scheduling frequencies. Exported data maintains FHIR R4 compliance for re-import capabilities.

Recovery: Multi-Region disaster recovery through Amazon S3 Cross-Region Replication and concurrent import jobs to accelerate restoration, application-level recovery patterns, and automatic scaling to minimise complexity. Supports customer RPO (regular export scheduling based on data change frequency) and RTO (concurrent import jobs and multi-Region active-passive architectures) objectives.

Setup Process

Onboarding: Create FHIR datastores through AWS Management Console, CLI, or SDKs after establishing AWS Identity and Access Management users, Amazon S3 buckets, and service roles. Configure AWS Signature Version 4 or SMART on FHIR authorisation. Upload FHIR R4 files to Amazon S3 and initiate import jobs.

Offboarding: Delete individual FHIR resources or entire datastores through AWS Management Console, CLI, or APIs. Export data for backup before deletion as required for regulatory compliance. Revoke customer-managed encryption keys to render data inaccessible.

Data Management

- **Storage Options:** Enterprise-grade FHIR R4 repository with automatic Regional replication and Amazon S3 integration
- **Data Removal:** FHIR versioned deletion hides resources from analysis while maintaining compliance; datastore deletion removes everything
- **Cross-Region Replication:** Export to Amazon S3 enables cross-Region replication strategies for multi-Region disaster recovery architectures
- **Encryption:** Customer-managed keys for encryption at rest and in-transit encryption with key revocation capabilities.

Optimise Cost and Consumption

Customers can optimise query patterns to stay within required queries per hour limit, preventing overage charges. NLP cost management through data preprocessing removes unnecessary text before analysis, reducing character-based charges. FHIR subscription delivery method selection between EventBridge and REST-Hook endpoints controls event-based pricing. Data lifecycle policies and scheduled off-peak exports optimise storage and transfer costs while Amazon S3 lifecycle management enables long-term retention savings.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/healthlake/latest/devguide/

Getting Started: https://docs.aws.amazon.com/healthlake/latest/devguide/getting-started.html

API Reference: https://docs.aws.amazon.com/healthlake/latest/APIReference/

Service Quotas: https://docs.aws.amazon.com/healthlake/latest/devguide/reference-healthlake-endpoints-quotas.html

AWS HealthOmics

General Information

AWS HealthOmics is a Health Insurance Portability and Accountability Act (HIPAA)-eligible service that accelerates clinical diagnostic testing, drug discovery, and agricultural research by fully managing bioinformatics workflows. It supports industry-standard workflow languages and scales infrastructure to handle tens of thousands of genomic tests daily with predictable cost-per-sample pricing.

Main Benefits

AWS HealthOmics reduces bioinformatics infrastructure complexity while delivering enterprise-scale performance that scales to over 100,000 concurrent vCPUs. Customers benefit from predictable cost-per-sample pricing, HIPAA-eligible compliance features, comprehensive data provenance tracking, and integration with AWS services environment, enabling focus on scientific breakthroughs rather than technical management.

Key Use Cases

- **Clinical Diagnostics:** Build scalable diagnostic testing workflows with predictable costs and variant calling capabilities
- **Drug Discovery:** Accelerate therapeutic research by orchestrating biological foundation models across millions of candidates
- **Agricultural Research:** Enhance crop traits like drought tolerance through AI-powered genomic analysis workflows
- **Bioinformatics Processing:** Process genomics data using industry-standard workflow languages with fully managed infrastructure.

Core Features

- **Workflows:** Process genomics data using WDL, Nextflow, or CWL with Ready2Run pre-built workflows
- **Purpose-Built Storage:** Specialised storage for bioinformatics file formats with sequence and reference stores
- **Auto-Scaling:** Seamlessly scales across over 100,000 concurrent vCPUs supporting tens of thousands of daily tests
- **Data Collaboration:** Secure sharing capabilities with fine-grained access controls, tagging, and permissions management
- **HIPAA Compliance:** Built-in compliance features with comprehensive audit trails and data provenance tracking
- **Predictable Pricing:** Cost-per-sample pricing model with no upfront costs for bioinformatics workflows
- **Service Integration:** Extensive integration with AWS services including Amazon Simple Storage Service (Amazon S3), Amazon Elastic Container Registry, Amazon SageMaker, and Amazon CloudWatch.

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/healthomics/
FAQ: https://aws.amazon.com/healthomics/faqs/

Technical Overview

Backup and Recovery Capabilities

Backup: AWS HealthOmics provides built-in data durability through purpose-built storage architecture with automatic ETag generation for content integrity validation. The service preserves data relationships for provenance tracking and supports data export capabilities to Amazon S3 for external backup strategies via AWS Management Console, CLI, or APIs.

Recovery: Multi-Region availability enables cross-Region data operations with distributed architecture across AWS Availability Zones. Data export features to Amazon S3 enable cross-Region replication strategies. Supports customer RPO requirements through export capabilities and RTO objectives via distributed architecture for faster recovery times.

Setup Process

Onboarding: Setup requires proper AWS Identity and Access Management permissions and service roles. Begin using Ready2Run workflows through the AWS HealthOmics console, selecting pre-built workflows and configuring output locations. Supports Amazon Quick CLI for natural language workflow management and GitHub integration via AWS CodeConnections.

Offboarding: Delete resources through AWS Management Console, CLI, or APIs including sequence stores, reference stores, and workflows. Data follows 30-day minimum retention with prorated early deletion charges. Requires proper cross-account permission management.

Data Management

- **Storage Options:** Purpose-built storage with sequence stores for genomics data and reference stores for genome references.
- **Data Removal:** Delete resources via AWS Management Console, CLI, or APIs with 30-day minimum retention and prorated early deletion charges.
- **Cross-Region Replication:** Multi-Region availability with cross-Region data operations and Amazon S3 export capabilities for replication strategies.
- **Encryption:** HIPAA-eligible infrastructure with built-in encryption, comprehensive audit trails, and data provenance tracking features.

Optimise Cost and Consumption

AWS HealthOmics provides run caching mechanisms to reduce redundant computations across identical workflow tasks. Automated storage tiering transitions infrequently accessed data to archive storage after 30 days, reducing long-term retention costs. Ready2Run workflows offer predictable per-run pricing for standardised bioinformatics tasks. Organisations can optimise workflow resource requests to match computational requirements and implement data lifecycle policies for automatic cost-efficient storage management.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/omics/latest/dev/index.html
Getting Started: https://docs.aws.amazon.com/omics/latest/dev/getting-started.html
API Reference: https://docs.aws.amazon.com/omics/latest/api/Welcome.html
Service Quotas: https://docs.aws.amazon.com/general/latest/gr/healthomics-quotas.html

AWS Identity and Access Management (IAM)

General Information

AWS Identity and Access Management (IAM) is a web service that securely controls access to AWS resources by managing permissions for users and entities. IAM provides centralised identity management, fine-grained access control through JSON policies, and security features including multi-factor authentication and identity federation across all AWS services.

Main Benefits

IAM provides centralised identity management with no additional charges, offering built-in integration with more than 250 AWS services. Unlike third-party solutions, IAM offers automatic scaling, built-in Payment Card Industry Data Security Standard (PCI DSS) compliance, and high availability through multi-server replication, enhancing security through least-privilege principles and audit capabilities.

Key Use Cases

- **Enterprise User Management:** Connect corporate Active Directory to centralised access control across accounts
- **Application Security:** Grant Amazon Elastic Compute Cloud (Amazon EC2) applications secure AWS service access using temporary credentials
- **Cross-Account Access:** Allow secure access between AWS accounts for partners and vendors
- **Compliance and Governance:** Implement least-privilege controls and audit trails for regulatory requirements

Core Features

- **IAM Users and Groups:** Create individual users and organise into groups with specific permissions
- **IAM Roles:** Provide temporary credentials for federated access, cross-account access, and Amazon EC2 applications
- **Policies and Permissions:** Define granular permissions using JSON policy language with identity-based and resource-based policies
- **Multi-Factor Authentication:** Support passkeys, security keys, virtual authenticator apps, and hardware TOTP tokens
- **Identity Federation:** Allow corporate network users to access AWS using existing passwords and temporary credentials
- **IAM Access Analyzer:** Generate least-privilege policies, verify public access, and validate IAM policy configurations
- **Permissions Boundaries:** Set maximum permissions for identity-based policies to establish guardrails for delegated administration
- **Service-Linked Roles:** Predefined IAM roles created automatically when AWS services need resource access permissions
- AWS IAM Identity Center: Centralized workforce identity management with single sign-on access to multiple AWS accounts and applications.

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/iam/

FAQ: https://aws.amazon.com/iam/faqs/

Technical Overview

Backup & Recovery Capabilities

Backup: IAM does not provide traditional backup capabilities as it is a managed service with built-in redundancy and automatic multi-server replication. Organisations should implement infrastructure-as-code practices using AWS CloudFormation or Amazon Terraform templates to version control and recreate IAM configurations for documentation purposes.

Recovery: IAM inherently supports disaster recovery through its eventually consistent, globally replicated architecture across multiple servers and regions. Cross-Region IAM role assumptions and federated access continue functioning during regional outages. Supports effectively zero RPO due to automatic replication and low RTO through cross-Region roles.

Setup Process

Onboarding: Begin by securing the root user with multi-factor authentication, then create administrative users for everyday tasks. Set up IAM users and groups with appropriate AWS managed policies, configure roles for applications, and implement multi-factor authentication (MFA). Follow least-privilege principles starting with minimal permissions.

Offboarding: Systematically delete IAM identities by removing access keys, detaching policies, and removing users from groups before final deletion. Use IAM Access Analyzer to identify unused access before removal and allow time for eventual consistency propagation.

Data Management

- **Storage Options:** IAM is a managed service with built-in redundancy and automatic multi-server replication for high availability
- **Data Removal:** IAM users can systematically delete users, roles, policies, and groups with IAM Access Analyzer identifying unused access
- **Cross-Region Replication:** IAM supports automatic replication across AWS Regions with eventual consistency and built-in disaster recovery capabilities
- **Encryption:** Identity management data is inherently protected through AWS managed service architecture with built-in security

Optimise Cost and Consumption

IAM Access Analyzer identifies unused permissions and generates least-privilege policies to prevent unintended resource usage. Permissions boundaries establish spending guardrails for delegated administrators by limiting maximum resource access capabilities. Policy conditions restrict actions based on time, IP address, or location to control when expensive operations occur. Cross-account roles help to remove duplicate resource provisioning across accounts while maintaining security controls.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/IAM/latest/UserGuide/introduction.html

Getting Started: https://docs.aws.amazon.com/serverless/latest/devguide/starter-iam.html

API Reference: https://docs.aws.amazon.com/IAM/latest/APIReference/welcome.html

Service Quotas: https://docs.aws.amazon.com/IAM/latest/UserGuide/reference_iam-quotas.html

AWS IoT Services – AWS IoT Core

Feature of AWS IoT Services

General Information

AWS IoT Core is a managed cloud service that supports secure bi-directional communication between connected devices and cloud applications. It provides device gateway capabilities, message broker functionality, and secure device management with authentication and encrypted communication.

Main Benefits

AWS IoT Core delivers scalability that can support billions of devices and trillions of messages across more than 20 global Regions. Enterprise-grade security with end-to-end encryption and certificate-based authentication helps ensure data protection. Built-in integration with more than 20 AWS services and SDKs help accelerate development taking away the infrastructure management overhead.

Key Use Cases

- **Connected Home Automation:** Smart devices, appliances, security systems with Alexa integration for monitoring
- **Industrial IoT and Predictive Maintenance:** Asset condition monitoring, predictive analytics, and equipment performance optimisation
- **Vehicle Connectivity and Fleet Management:** Connected mobility solutions with telematics and real-time analytics
- **LoRaWAN Applications:** Low-power, long-range connectivity for sensor networks and remote monitoring

Core Features

- **Device Gateway:** Secure bi-directional communication bridge supporting MQTT, HTTP, and WebSocket protocols
- **Message Broker:** Scalable messaging service for real-time communication between devices and AWS services
- **Device Shadow:** Virtual device representation maintaining state information for online/offline communication and synchronisation
- **Rules Engine:** SQL-like service for filtering, transforming, and routing device data to AWS services
- **Thing Registry:** Device metadata database storing certificates, attributes, and thing types for organised management
- **Authentication and Authorisation:** Fine-grained security with X.509 certificates, AWS Identity and Access Management (AWS IAM) integration, and IoT policies
- **LoRaWAN Support:** Managed LoRaWAN Network Server for low-power, long-range connectivity with network support
- **Device Advisor:** Testing service for validating IoT devices before deployment with compatibility assessments

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/iot-core/
FAQ: https://aws.amazon.com/iot-core/faqs/

Technical Overview

Backup & Recovery Capabilities

Backup: Device registry data, certificates, and policies are replicated to secondary Regions through Amazon DynamoDB streams and AWS Lambda functions. Custom solutions provide backup for device shadow state and metadata. AWS Backup delivers comprehensive data protection for connected services.

Recovery: The system features cross-Region replication with automated failover powered by Amazon Route 53 health checks, plus a dedicated disaster recovery solution using AWS Step Functions and Amazon DynamoDB Global Tables. The multi-Region architecture supports fault tolerance and supports customer recovery objectives: RPO of minutes through continuous replication, and RTO ranging from 2-4 hours for backup/restore operations to sub-hour performance for active/active architectures.

Setup Process

Onboarding: Access AWS IoT console to create thing objects, generate security certificates and keys, configure device policies, and install Device SDKs. Accelerate deployment using quick connect tutorials, interactive demos, and documentation. Validate compatibility before production using Device Advisor.

Offboarding: Delete thing objects, revoke certificates, remove associated policies and thing groups, clear all device shadow data, disable rules engine configurations, and delete any custom authorisers. Note that data stored in connected AWS services requires separate deletion.

Data Management

- **Storage Options:** Thing Registry for device metadata, Device Shadow for state information, integrated AWS services for data storage.
- **Data Removal:** Users can delete thing objects, revoke certificates, remove policies, clear device shadows, disable rules engine configurations manually.
- **Cross-Region Replication:** Device registry data, certificates, and policies are automatically replicated across Regions using Amazon DynamoDB Global Tables combined with AWS Lambda functions.
- **Encryption:** Data protection through end-to-end encryption using TLS protocols, X.509 certificate-based device authentication, and fine-grained access controls for authorisation management.

Optimise Cost and Consumption

Message batching reduces per-message costs while payload compression minimises data transfer charges. Basic Ingest feature provides free message processing using reserved topics. Device shadows minimise frequent state updates, and Rules Engine filtering reduces downstream service calls. Persistent sessions optimise device connections by reducing reconnection overhead. LoRaWAN connectivity offers low-cost, low-power communication options for sensor networks.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/iot/latest/developerguide/what-is-aws-iot.html
Getting Started: https://docs.aws.amazon.com/iot/latest/developerguide/iot-gs.html
API Reference: https://docs.aws.amazon.com/iot/latest/apireference/
Service Quotas: https://docs.aws.amazon.com/general/latest/gr/iot-core.html

AWS IoT Services – AWS IoT Device Management

Feature of AWS IoT Services

General Information

AWS IoT Device Management aids secure, scalable management of IoT devices throughout their lifecycle. It provides bulk registration, fleet organisation, remote updates, secure tunneling, and monitoring capabilities for devices from sensors to complex systems.

Main Benefits

Provides fully managed device lifecycle operations with enterprise-grade security and 99.9% SLA reliability, removing infrastructure overhead. Integrates with more than 200 AWS services for advanced analytics and machine learning. Pay-as-you-use pricing requires no upfront costs while automated bulk operations reduce per-device management expenses.

Key Use Cases

- **Fleet Management:** Organise large device fleets with hierarchical grouping and automated lifecycle management
- **Remote Device Operations:** Send commands, retrieve logs, and perform remote diagnostics on field devices
- **Over-the-Air Updates:** Push software updates with controlled rollout schedules and rollback capabilities
- **Device Troubleshooting:** Establish secure tunneling connections for remote validation and corrective actions

Core Features

- **Bulk Registration:** Register multiple devices using templates with device information, certificates, and configurations
- **Fleet Indexing and Search:** Query device records and aggregate statistics based on attributes, states, and connectivity
- **Device Jobs:** Run and monitor software updates and remote operations on single devices or entire fleets
- **Secure Tunneling:** Create secure communications sessions to individual devices for troubleshooting without firewall changes
- **Device Commands:** Send low-latency remote commands to devices with real-time execution status tracking
- **Software Package Management:** Track and monitor software package versions with centralised dashboard and targeted updates
- **Device Organisation:** Group devices hierarchically based on function, security requirements, or other categories
- **Managed Integrations:** Pre-built integrations for ZigBee, Z-Wave, and Wi-Fi devices with unified device control

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/iot-device-management/

FAQ: https://aws.amazon.com/iot-device-management/faq/

Technical Overview

Backup & Recovery Capabilities

Backup: Cross-Region replication of device registries, certificates, and policies using AWS Lambda and Amazon DynamoDB streams. AWS Backup protects associated data in integrated services like Amazon Simple Storage Service (Amazon S3) and Amazon DynamoDB. Custom backup strategies required for device configurations using infrastructure-as-code practices.

Recovery: Cross-Region replication capabilities through Amazon DynamoDB global tables and AWS Lambda functions, automatic failover via Amazon Route 53 health checks, and disaster recovery through secondary Region deployment. Supports customer RPO objectives (real-time replication of device registry data) and RTO objectives (2-4 hours with automated failover mechanisms).

Setup Process

Onboarding: Access AWS IoT Device Management console to create IoT things with certificates and policies. Bulk registration allows template uploads for large-scale deployments. Configure device grouping and fleet indexing. Prerequisites include AWS CLI installation and AWS Identity and Access Management (AWS IAM) permissions.

Offboarding: Delete device things, certificates, and policies using delete-managed-thing command or console interface. Remove device shadow data and historical logs through service APIs. Separate deletion processes required for integrated services data.

Data Management

- **Storage Options:** Device data and firmware packages stored in Amazon DynamoDB and Amazon S3 with scalable storage
- **Data Removal:** Delete device things, certificates, policies using delete-managed-thing command or AWS Management Console interface APIs
- **Cross-Region Replication:** Device registries, certificates, policies replicated using Amazon DynamoDB global tables and AWS Lambda functions
- **Encryption:** Enterprise-grade security with encrypted communications and secure tunneling capabilities for device management

Optimise Cost and Consumption

Batch device operations to reduce per-operation costs and optimise bulk registration efficiency. Control fleet indexing expenses by limiting custom fields and query complexity. Optimise device jobs through targeted deployments to specific device groups rather than fleet-wide operations. Minimise secure tunneling costs by limiting session duration and concurrent connections. Implement device data compression and efficient data transfer protocols to reduce bandwidth consumption costs.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/iot/latest/developerguide/what-is-aws-iot.html

Getting Started: https://docs.aws.amazon.com/iot/latest/developerguide/iot-gs.html

API Reference: https://docs.aws.amazon.com/iot/latest/developerguide/iot-remote-command.html

Service Quotas:
https://docs.aws.amazon.com/general/latest/gr/iot_device_management.html

AWS IoT Services – AWS IoT Greengrass

Feature of AWS IoT Services

General Information

AWS IoT Greengrass is an open-source edge runtime and cloud service that extends AWS IoT Core capabilities to edge devices. It allows local data processing, machine learning inference, and secure device communication with offline operation capabilities.

Main Benefits

AWS IoT Greengrass delivers sub-millisecond latency for critical applications through local processing and ML inference. It reduces bandwidth costs by up to 90% through edge data filtering, allows offline operation during connectivity issues, and provides centralised fleet management with automated over-the-air updates.

Key Use Cases

- **Industrial IoT:** Deploy predictive maintenance, factory automation, and equipment monitoring applications
- **Smart Cities:** Allow autonomous traffic systems, smart lighting, and environmental monitoring
- **Connected Vehicles:** Support autonomous operations, fleet management, and real-time diagnostics
- **Healthcare Monitoring:** Deploy patient monitoring systems and medical device management solutions

Core Features

- **Local Processing:** Run AWS Lambda functions locally with containers, device shadows, and messaging capabilities
- **Machine Learning Inference:** Deploy and run ML models at edge with SageMaker integration
- **Stream Manager:** Manage data streams from multiple sources with local buffering and synchronisation
- **Over-the-Air Updates:** Deploy software updates, configuration changes, and components remotely to devices
- **Security Integration:** Hardware Security Module integration, secrets management, and end-to-end encryption for operations
- **Client Device Support:** Act as secure MQTT endpoint for local devices with messaging
- **Component System:** Modular software architecture using pre-built components for extending device functionality
- **Offline Operation:** Continue local processing, messaging, and data caching during intermittent connectivity

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/greengrass/

FAQ: https://aws.amazon.com/greengrass/faqs/

Technical Overview

Backup & Recovery Capabilities

Backup: Component artifacts are stored in Amazon Simple Storage Service (Amazon S3) with versioning, while configurations backup via AWS CloudFormation templates. Local data storage and Stream Manager provide local buffering with data persistence. Custom components support device data backup using Amazon S3 versioning capabilities.

Recovery: Multi-region device registry replication with automated failover using AWS Route 53 health checks support cross-Region recovery. Local shadows maintain device state, while persistent MQTT sessions ensure message delivery upon reconnection. Supports near-zero RPO through local buffering and RTO of minutes to hours via automated deployment mechanisms.

Setup Process

Onboarding: Create AWS IoT thing, install Greengrass Core software on Linux/Windows devices with 512MB RAM minimum. Automated setup script installs dependencies and certificates. Fleet provisioning aids zero-touch deployment for large-scale implementations.

Offboarding: Delete components, deployments, and uninstall Core software from devices. Remove IoT things, certificates, policies, and local cached data. Revoke certificates and clean IAM roles for complete removal.

Data Management

- **Storage Options:** Local data caching, Amazon S3 for component artifacts, Amazon DynamoDB for edge storage
- **Data Removal:** Manual deletion of local data, component artifacts, logs, and cached data
- **Cross-Region Replication:** Multi-Region device registry replication with automated failover using Amazon Route 53
- **Encryption:** End-to-end encryption, Hardware Security Module integration, and secure secrets management

Optimise Cost and Consumption

Stream Manager's data aggregation and compression capabilities minimise bandwidth usage by processing multiple data sources locally before cloud transmission. The component-based architecture allows selective deployment of only required functionality, optimising resource consumption. Local data filtering and preprocessing reduce cloud storage costs by transmitting only relevant data. Pay-per-active-core pricing aligns costs with actual device usage patterns.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/greengrass/v2/developerguide/what-is-iot-greengrass.html

Getting Started: https://docs.aws.amazon.com/greengrass/v2/developerguide/getting-started.html

API Reference: https://docs.aws.amazon.com/greengrass/v2/APIReference/Welcome.html

Service Quotas: https://docs.aws.amazon.com/general/latest/gr/greengrassv2.html

AWS IoT Services – AWS IoT SiteWise

Feature of AWS IoT Services

General Information

AWS IoT SiteWise is a managed service that allows industrial enterprises to collect, store, organise, and monitor data from industrial equipment at scale. It provides asset modeling frameworks to represent devices and processes, with capabilities for real-time monitoring and predictive analytics.

Main Benefits

AWS IoT SiteWise takes away the complexity of industrial data collection with managed infrastructure, automatic scaling, and built-in security. Features three-tier automated storage cost optimisation, edge processing for hybrid architectures, AI-powered assistant for accelerated troubleshooting, and pre-built asset modelling frameworks that reduce time to value.

Key Use Cases

- **Manufacturing:** Monitor operations across facilities, compute OEE and MTBF metrics, prevent equipment issues
- **Food and Beverage:** Optimise production processes, support quality control, maintain compliance with safety standards
- **Energy and Utilities:** Monitor power generation equipment, optimise asset performance, predict equipment failures
- **Predictive Maintenance:** Use anomaly detection and AI assistant for equipment insights and troubleshooting

Core Features

- **Data Ingestion:** Collect data through OPC UA, MQTT, edge gateways, AWS IoT Core rules, and APIs
- **Asset Modelling:** Create representations of industrial devices, processes, and facilities with hierarchies and property definitions
- **Edge Processing:** Process, store, and analyse data locally using AWS IoT SiteWise Edge with data collection and processing packs
- **Time Series Storage:** Three-tier storage system with intelligent data lifecycle management for automatic cost optimisation
- **Analytics and Monitoring:** SQL queries, real-time dashboards, Amazon CloudWatch alarms, and anomaly detection for predictive maintenance
- **Visualisation:** No-code web applications through AWS IoT SiteWise Monitor and Grafana plugin for operational data visualisation
- **AI Assistant:** Generative AI-powered industrial assistant providing insights, troubleshooting, and operational guidance
- **Multi-Protocol Support:** Support up to 100 OPC UA server connections per edge gateway for device connectivity

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/iot-sitewise/
FAQ: https://aws.amazon.com/iot-sitewise/faqs/

Technical Overview

Backup & Recovery Capabilities

Backup: Data backup is achieved through property value notifications to AWS IoT Core with rules routing to Amazon Simple Storage Service (Amazon S3) and Amazon DynamoDB. Historical asset property data and resource definitions can be backed up manually using query APIs or describe APIs for disaster recovery scenarios.

Recovery: Recovery is supported through multi-Availability Zone (multi-AZ) architecture and integration with durable AWS services. Cross-Region disaster recovery requires custom solutions using Amazon DynamoDB Global Tables and AWS Lambda functions to replicate device registries and configurations. Supports customer RPO (measured in minutes through data replication) and RTO (hours through automated replication of asset configurations) objectives.

Setup Process

Onboarding: Begin with quick start demo creating sample asset models and AWS IoT SiteWise Monitor portal through AWS Management Console. Choose data ingestion methods including OPC UA connections, edge gateways, or MQTT. Console-based asset modeling allows equipment representations with immediate visualisation capabilities.

Offboarding: Delete resources including assets, asset models, gateways, and portals through console or APIs. Export historical data using query APIs or bulk export to Amazon S3. Quick start demo auto-deletes while production resources require manual cleanup.

Data Management

- **Storage Options:** Three-tier storage system (hot, warm, cold) with intelligent data lifecycle management for cost optimisation
- **Data Removal:** Delete resources through console or APIs; export historical data using query APIs or bulk export to Amazon S3
- **Cross-Region Replication:** Custom solutions using Amazon DynamoDB Global Tables and AWS Lambda functions for device registry and configuration replication
- **Encryption:** Built-in security with managed infrastructure; integrates with AWS security features for data protection

Optimise Cost and Consumption

AWS IoT SiteWise provides intelligent three-tier storage tiering that automatically moves historical data to cost-optimised warm and cold tiers based on user-defined policies. Buffered ingestion allows data velocity configuration for different use cases, optimising messaging costs. SiteWise Edge processes data locally to reduce cloud compute and data transfer expenses, while Amazon S3 export capabilities offer additional cost-effective archival options for long-term analytics workloads.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/iot-sitewise/latest/userguide/what-is-sitewise.html
Getting Started: https://docs.aws.amazon.com/iot-sitewise/latest/userguide/getting-started.html
API Reference: https://docs.aws.amazon.com/iot-sitewise/latest/APIReference/Welcome.html
Service Quotas: https://docs.aws.amazon.com/iot-sitewise/latest/userguide/endpoints-and-quotas.html

AWS Key Management Service

General Information

AWS Key Management Service (AWS KMS) is a fully managed cryptographic service that allows customers to create and control encryption keys for data protection. Operating with FIPS 140-3 validated hardware security modules, AWS KMS provides centralised key management across AWS services and applications.

Main Benefits

AWS KMS removes cryptographic infrastructure complexity while providing FIPS 140-3 validated security through hardware security modules. With integration across more 100 AWS services, 99.999999999% durability, and 20,000 free monthly requests, AWS KMS delivers enterprise-grade encryption with pay-per-use efficiency and automatic key lifecycle management.

Key Use Cases

- **Data Encryption at Rest:** Protecting sensitive data in Amazon Simple Storage Service (Amazon S3), Amazon Elastic Block Store (Amazon EBS), Amazon Relational Database Service (Amazon RDS), and Amazon DynamoDB services
- **Client-Side Encryption:** Using envelope encryption with AWS Encryption SDK before sending data to services
- **Digital Signing and Verification:** Creating asymmetric keys for signing documents, APIs, and code authentication
- **Compliance Requirements:** Helps to meet Payment Card Industry Data Security Standard (PCI DSS) and the General Data Protection Regulation (GDPR) standards through centralised key management

Core Features

- **Hardware Security Modules:** FIPS 140-3 Security Level 3 validated HSMs with tamper-evident and tamper-resistant protection
- **Symmetric Encryption Keys:** 256-bit AES keys for direct encryption up to 4KB or envelope encryption
- **Asymmetric Keys:** RSA, ECC, ML-DSA, and SM2 key pairs for encryption, signing, and key agreement
- **Automatic Key Rotation:** Annual rotation for AWS-managed keys and configurable rotation for customer-managed keys
- **Multi-Region Keys:** Keys replicated across multiple AWS Regions for global applications and disaster recovery
- **Custom Key Stores:** Integration with AWS CloudHSM clusters or external key managers for additional control
- **Fine-Grained Access Controls:** IAM policies, key policies, and temporary grants for delegation and access management
- **AWS CloudTrail Integration:** Detailed logging of all key usage and management operations for audit compliance

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/kms/

FAQ: https://aws.amazon.com/kms/faqs/

Technical Overview

Backup & Recovery Capabilities

Backup: AWS KMS keys are automatically backed up across multiple Availability Zones with 99.999999999% durability storage. Historical key versions are maintained through automatic and on-demand rotation. Multi-Region keys provide cross-region backup capabilities, while AWS Backup integration allows customer-managed key encryption of backup vaults.

Recovery: Sub-second recovery times for most operations through automatic failover mechanisms and globally distributed endpoints. Multi-Region key replication allows disaster recovery with AWS Application Recovery Controller integration. Supports near-zero RPO (real-time replication) and RTO targets of minutes (automatic failover) objectives.

Setup Process

Onboarding: Minimal setup through AWS Management Console, AWS CLI, or SDKs. Create keys by selecting symmetric encryption, asymmetric signing, or HMAC types. Integration achieved by specifying AWS KMS key ARNs in service configurations with AWS SDK examples available.

Offboarding: Secure key deletion with mandatory 7-30 day waiting period preventing accidental data loss. Keys can be disabled for immediate cessation or recovered during waiting period. Multi-Region keys require primary Region deletion.

Data Management

- **Storage Options:** FIPS 140-3 validated HSMs with 99.999999999% durability across multiple Availability Zones automatically
- **Data Removal:** Secure key deletion with mandatory 7-30 day waiting period preventing accidental data loss
- **Cross-Region Replication:** Multi-Region keys replicated across AWS Regions for global applications and disaster recovery
- **Encryption:** 256-bit AES symmetric keys, RSA/ECC asymmetric keys, and HMAC keys for data integrity

Optimise Cost and Consumption

The AWS Encryption SDK's data key caching capability dramatically reduces API calls—by as much as 99%—through strategic key reuse. Amazon Simple Storage Service (Amazon S3) bucket keys further optimise performance by distributing a single encryption key across multiple objects, decreasing the volume of per-object AWS KMS requests. Additional efficiency gains come from request batching and connection pooling, which minimise API overhead. Organisations can also consolidate multiple applications under a smaller number of keys, preserving security boundaries while lowering operational expenses.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/kms/latest/developerguide/

Getting Started: https://docs.aws.amazon.com/kms/latest/developerguide/example_kms_Scenario_Basics_section.html

API Reference: https://docs.aws.amazon.com/kms/latest/APIReference/

Service Quotas: https://docs.aws.amazon.com/general/latest/gr/kms.html

AWS Lake Formation

General Information

AWS Lake Formation is a fully managed service that simplifies building, securing, and managing data lakes on AWS. It provides centralised governance with fine-grained access control for data stored in Amazon Simple Storage Service (Amazon S3) and metadata in AWS Glue Data Catalog. The service automates complex data lake setup tasks including data ingestion, cataloging, transformation, and security management.

Main Benefits

AWS Lake Formation simplifies complex data lake setup by automating permissions, cataloging, and security management with centralised governance. It provides fine-grained access control at database, table, column, row, and cell levels without complex identity and access management (IAM) policies. The service reduces operational overhead, accelerates time-to-value for data initiatives, and takes away the overhead of third-party data governance tool licensing fees.

Key Use Cases

- **Data Lake Governance:** Centrally govern enterprise data lakes with fine-grained access control at multiple levels
- **Cross-Account Data Sharing:** Securely share data across AWS accounts with tag-based and named resource controls
- **Analytics and Machine Learning:** Allow secure data access for Amazon Athena, Amazon Redshift, Amazon Elastic MapReduce (Amazon EMR), and Amazon SageMaker services
- **Compliance and Audit:** Implement data access monitoring and regulatory compliance for sensitive data

Core Features

- **Fine-grained Access Control:** Database, table, column, row, and cell-level permissions with tag-based and attribute-based access control
- **Cross-Account Data Sharing:** Secure data sharing across AWS accounts and external organisations with centralised governance
- **Data Catalog Management:** Centralised metadata management using AWS Glue Data Catalog with automatic schema discovery
- **Audit and Compliance:** Comprehensive data access logging, monitoring, and audit capabilities for regulatory compliance
- **Blueprint Workflows:** Pre-built ETL workflows for database snapshots, incremental updates, and log file processing
- **Hybrid Access Mode:** Selective onboarding of AWS Lake Formation permissions alongside existing IAM controls
- **Governed Tables:** Automatic storage layout optimisation to improve query performance and reduce costs
- **JDBC Connectivity:** Support for external data sources through JDBC connections for database integration

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/lake-formation/
FAQ: https://aws.amazon.com/lake-formation/faqs/

Technical Overview

Backup & Recovery Capabilities

Backup: AWS Lake Formation relies on underlying AWS services for data protection. Amazon S3 data uses cross-Region replication and versioning, while data catalog metadata can be exported using AWS Glue APIs. AWS Lake Formation permissions and configurations can be backed up via AWS Config or exported through APIs.

Recovery: Recovery involves Amazon S3 data replication and Data Catalog reconstruction with cross-region data access capabilities. AWS CloudFormation templates allow automated infrastructure deployment for permission and configuration recreation. Supports customer RPO objectives with metadata replication in minutes and RTO minimisation through cross-region access and automated deployment.

Setup Process

Onboarding: Setup begins through AWS Management Console or AWS CloudFormation templates by creating a data lake administrator and configuring default permissions. Register Amazon S3 locations as data lake storage, create Data Catalog databases, and configure user permissions using pre-built blueprints.

Offboarding: Remove Lake Formation by deleting Data Catalog tables and databases, revoking all user permissions, and unregistering Amazon S3 locations. Delete cross-account sharing configurations and custom IAM roles while auditing all permissions for complete access removal.

Data Management

- **Storage Options:** Data stored in Amazon S3 with metadata managed through AWS Glue Data Catalog
- **Data Removal:** Delete Data Catalog tables, revoke permissions, unregister Amazon S3 locations, and remove cross-account sharing configurations
- **Cross-Region Replication:** Supports cross-region data access but requires customer-implemented Amazon S3 cross-Region replication strategies for backup
- **Encryption:** Inherits Amazon S3 encryption capabilities for data at rest and in transit protection mechanisms

Optimise Cost and Consumption

AWS Lake Formation's Storage API pricing model charges only for bytes scanned, encouraging columnar formats like Parquet and ORC to minimise scanning costs. Governed Tables automatically optimise storage layout through compaction to improve query performance. Cross-account data sharing capabilities reduce data duplication and transfer charges across organisations. The service provides no additional charges for basic permissions management and governance features.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/lake-formation/latest/dg/what-is-lake-formation.html
Getting Started: https://docs.aws.amazon.com/lake-formation/latest/dg/getting-started-setup.html
API Reference: https://docs.aws.amazon.com/lake-formation/latest/dg/aws-lake-formation-api.html
Service Quotas: https://docs.aws.amazon.com/general/latest/gr/lake-formation.html#limits_lake-formation

AWS Lambda

General Information

AWS Lambda is a serverless compute service that runs code without provisioning or managing servers. It automatically scales execution environments, supports multiple programming languages, and charges based on actual usage for event-driven applications.

Main Benefits

AWS Lambda takes away server management overhead and automatically scales from zero to enterprise scale with pay-per-use pricing. It integrates with more than 200 AWS services, delivers sub-second cold start performance with SnapStart, and allows significant cost savings for variable workloads.

Key Use Cases

- **File Processing:** Automatically process files when uploaded to Amazon Simple Storage Service (Amazon S3) or other storage services
- **Real-time Stream Processing:** Process data streams from Amazon Kinesis and Amazon DynamoDB for analytics and monitoring
- **Web and Mobile Backends:** Create scalable API backends using Amazon API Gateway integration that automatically adjust
- **IoT Backends:** Handle IoT device requests and process sensor data at enterprise scale

Core Features

- **Automatic Scaling:** Scales execution environments from zero to thousands based on demand with configurable concurrency controls
- **Multi-language Runtime Support:** Supports Node.js, Python, Java, Go, .NET, Ruby, and custom runtimes with managed environments
- **Event-driven Architecture:** Integrates with more than 200 AWS services as event sources including Amazon S3, Amazon DynamoDB, and Amazon API Gateway
- **SnapStart:** Reduces cold start times to sub-second startup performance by creating snapshots during initialisation
- **Lambda Layers:** Share code libraries and dependencies across multiple functions to optimise code reuse and maintenance
- **Provisioned Concurrency:** Pre-initialise execution environments to remove cold starts for latency-sensitive applications
- **Container Image Support:** Deploy functions using container images up to 10GB for complex dependencies and custom runtimes
- **VPC Integration:** Access resources in Amazon VPC with network isolation and comprehensive security controls

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/lambda/

FAQ: https://aws.amazon.com/lambda/faqs/

Technical Overview

Backup & Recovery Capabilities

Backup: AWS Lambda functions are automatically replicated across Availability Zones with built-in durability. Deployment packages can be stored in Amazon S3 with versioning based (console, CLI, API) or automated via CI/CD pipelines. Version control systems and Infrastructure as Code templates provide configuration backup capabilities.

Recovery: Multi-Region deployment enables sub-minute recovery times, cross-Region replication for disaster recovery, and Application Recovery Controller orchestration. Built-in durability provides sub-second recovery for regional failures. Supports aggressive RPO (approaching zero for stateless functions) and RTO (minutes to sub-minute) objectives.

Setup Process

Onboarding: Setup through AWS Management Console, AWS CLI, or Infrastructure as Code tools using built-in templates or sample code. Configure triggers from AWS services and deploy via ZIP files, container images, or CI/CD pipelines.

Offboarding: Function deletion removes all versions, aliases, and associated resources via AWS Management Console, CLI, or API. Separately delete event source mappings, AWS IAM roles, Amazon CloudWatch logs, and S3 deployment artifacts for complete cleanup.

Data Management

- **Storage Options:** Ephemeral /tmp directory provides 512 MB to 10 GB storage with automatic replication across Availability Zones
- **Data Removal:** Function deletion permanently removes source code, deployment packages, versions, aliases, and associated resources immediately
- **Cross-Region Replication:** Multi-Region deployment supported via Infrastructure as Code with Amazon S3 cross-Region replication for deployment packages
- **Encryption:** Code signing ensures integrity and security compliance by verifying signed code before function execution

Optimise Cost and Consumption

AWS Compute Optimiser delivers ML-driven memory configuration recommendations reducing costs by 10-50%. Graviton processors provide 20% better price performance. SnapStart minimises cold start duration costs while Provisioned Concurrency removes cold starts for predictable workloads. AWS Lambda Layers optimise code reuse across functions, and AWS Lambda Managed Instances offer Amazon Elastic Compute Cloud (Amazon EC2) pricing for steady-state workloads. Savings Plans provide additional discounts for consistent usage patterns.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/lambda/latest/dg/welcome.html
Getting Started: https://docs.aws.amazon.com/lambda/latest/dg/getting-started.html
API Reference: https://docs.aws.amazon.com/lambda/latest/api/
Service Quotas: https://docs.aws.amazon.com/lambda/latest/dg/gettingstarted-limits.html

AWS License Manager

General Information

AWS License Manager helps organisations manage and track software licenses from vendors like Microsoft, IBM, Oracle, and SAP across AWS Regions and accounts. It provides centralised visibility, automated compliance enforcement, and comprehensive reporting to prevent licensing violations and reduce audit risks.

Main Benefits

AWS License Manager delivers significant cost savings with up to 77% reduction for Windows Server licenses and 45% for SQL Server licenses through optimised BYOL management. The service provides automated compliance enforcement, prevents costly licensing violations, and offers centralised visibility across cloud and on-premises environments at no additional charge.

Key Use Cases

- **BYOL Management:** Manage existing software licenses for cloud resources while maintaining compliance
- **License Compliance Tracking:** Monitor and enforce license usage across AWS and on-premises environments
- **Automated License Discovery:** Automatically discover and track software on Amazon Elastic Compute Cloud (Amazon EC2), Amazon Relational Database Service (Amazon RDS), and AWS Elastic Beanstalk
- **Vendor Audit Preparation:** Provide comprehensive reporting for software vendor audits and compliance verification

Core Features

- **Self-Managed Licenses:** Create and manage licensing rules based on enterprise agreements for compliance
- **License Asset Groups:** Organisation-wide license management across multiple regions and accounts with centralised oversight
- **Automated Discovery:** Automatically discover and track software installed on AWS resources including uninstalled software
- **User-Based Subscriptions:** Associate licenses with specific users for Microsoft Office and Visual Studio subscriptions
- **Linux Subscriptions Management:** Manage and track Red Hat and SUSE subscriptions with conversion capabilities
- **License Type Conversions:** Convert instances between license types and switch between different licensing models
- **Built-in Dashboards:** Monitor license usage and compliance through comprehensive dashboards with real-time visibility
- **Managed Entitlements:** Allow ISVs to manage and distribute software licenses through centralised entitlement management

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/license-manager/

FAQ: https://aws.amazon.com/license-manager/faqs/

Technical Overview

Backup & Recovery Capabilities

Backup: AWS License Manager does not provide traditional backup capabilities as it manages license metadata rather than customer data. License configurations are automatically maintained by AWS service infrastructure with built-in redundancy across availability zones. Customers should document license rules using Infrastructure as Code approaches with AWS CloudFormation templates for reproducibility.

Recovery: AWS License Manager relies on AWS's underlying infrastructure resilience for availability. AWS License configurations remain available during infrastructure failures through built-in service resilience. For disaster recovery, customers must recreate license rules in alternate regions using AWS CloudFormation templates, as cross-region tracking is not supported. RPO and RTO objectives are not applicable since the service manages configuration metadata.

Setup Process

Onboarding: Setup begins through the AWS License Manager console with automatic IAM role configuration. Create self-managed licenses using enterprise agreements, associate with AMIs, launch templates, or CloudFormation. Allow AWS Organizations trusted access for cross-account discovery and centralised management.

Offboarding: Deactivate licenses to stop tracking, then systematically disassociate from Amazon EC2 instances and Amazon RDS databases. Remove organisational integrations by disabling trusted access via AWS Organizations CLI or API. License configurations automatically disassociate from deregistered AMIs every 8 hours.

Data Management

- **Storage Options:** License metadata stored within AWS service infrastructure with built-in redundancy across availability zones
- **Data Removal:** Deactivate licenses, disassociate from resources, remove organisational integrations via AWS Organizations CLI or API
- **Cross-Region Replication:** Not supported; separate license configurations required per region with no cross-region tracking
- **Encryption:** Uses open industry standards for license representation and cryptographic verification for security

Optimise Cost and Consumption

AWS License Manager optimises costs through BYOL management that repurposes existing license inventory for cloud resources. vCPU-based licensing integration with Optimise CPU capability tracks custom usage patterns for reduced licensing costs. Automated license usage tracking identifies unused or underused licenses for reallocation. License type conversion allows switching between licensing models for optimal cost efficiency. Proactive license limit enforcement prevents costly overages and compliance violations.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/license-manager/latest/userguide/license-manager.html
Getting Started: https://docs.aws.amazon.com/license-manager/latest/userguide/getting-started.html
API Reference: https://docs.aws.amazon.com/license-manager/latest/APIReference/
Service Quotas: https://docs.aws.amazon.com/general/latest/gr/licensemanager.html

AWS Network Firewall

General Information

AWS Network Firewall is a managed, stateful network firewall and intrusion prevention service for Amazon Virtual Private Cloud (Amazon VPC). It filters traffic at VPC perimeters using Suricata-based inspection with deep packet analysis, domain filtering, and TLS inspection capabilities.

Main Benefits

Removes infrastructure management overhead with fully managed service scaling up to 100 Gbps per Availability Zone. Provides NAT Gateway cost discounts when deployed in service chains, enterprise-grade Suricata-based threat protection, and quick integration with AWS Transit Gateway and AWS Firewall Manager for centralised multi-account security governance.

Key Use Cases

- **VPC Perimeter Security:** Filter network traffic at VPC boundaries to protect against external threats
- **Compliance Requirements:** Meet security standards through centralised firewall management and detailed traffic monitoring
- **Advanced Threat Protection:** Block malicious URLs, command-and-control channels, and other threat vectors using IPS
- **Multi-Account Governance:** Centrally manage firewall policies across multiple AWS accounts using Firewall Manager

Core Features

- **Stateful and Stateless Inspection:** Context-aware traffic flow inspection and individual packet inspection for comprehensive filtering
- **Suricata-based IPS:** Open source intrusion prevention system for advanced threat detection with compatible rule support
- **TLS Inspection:** Decrypts and inspects SSL/TLS encrypted traffic for enhanced security analysis and threat detection
- **Active Threat Defense:** Managed rule groups providing real-time threat intelligence and protection against emerging threats
- **Domain and URL Filtering:** Web filtering capabilities to control access to specific domains and URLs
- **Flow Management:** Flow capture and flush operations for network visibility, troubleshooting, and security policy enforcement
- **Central Management:** Integration with AWS Firewall Manager for centralised configuration across multiple AWS accounts
- **Transit Gateway Integration:** Built-in integration for simplified centralised network security deployment across multi-VPC architectures

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/network-firewall/
FAQ: https://aws.amazon.com/network-firewall/faqs/

Technical Overview

Backup & Recovery Capabilities

Backup: Configuration backup achieved through Infrastructure as Code practices using AWS CloudFormation templates or Terraform (APIs and Console) with AWS Config for configuration history tracking. AWS manages underlying infrastructure backup automatically.

Recovery: Cross-Region deployment using AWS CloudFormation or Terraform achieves RTO of minutes to hours, with multi-Availability Zone deployment providing near-zero RTO for AZ failures and built-in high availability. Supports effectively zero RPO (version-controlled configurations) and aggressive RTO (automated Infrastructure as Code deployment) objectives.

Setup Process

Onboarding: Deploy via AWS Management Console with step-by-step guidance, AWS CloudFormation templates, or APIs. Create rule groups, firewall policies, associate with Amazon VPC subnets, and update route tables. Requires dedicated firewall subnets per Availability Zone.

Offboarding: Reverse Amazon VPC route table changes, delete endpoint associations, and systematically remove firewall resources. Manually delete Amazon CloudWatch Logs and Amazon S3 data per retention policies.

Data Management

- **Storage Options:** Firewall logs stored in Amazon CloudWatch Logs and Amazon S3 based on organisational data retention policies
- **Data Removal:** Manual deletion of firewall logs from Amazon CloudWatch Logs and Amazon S3 required during offboarding process
- **Cross-Region Replication:** Cross-Region deployment achieved through Infrastructure as Code approaches using AWS CloudFormation templates or Terraform
- **Encryption:** TLS inspection capabilities decrypt and inspect SSL/TLS encrypted traffic using certificates from AWS Certificate Manager

Optimise Cost and Consumption

NAT Gateway hourly and data processing charges are waived when deployed in same-region Availability Zone service chains. Secondary VPC endpoints offer reduced hourly rates. Tiered pricing for Advanced Inspection and Active Threat Defense allows selective capability activation. Transit Gateway attachments provide flexible cost allocation across accounts based on actual usage, while centralised deployment models reduce multiple firewall instances requirements.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/network-firewall/latest/developerguide/what-is-aws-network-firewall.html

Getting Started: https://docs.aws.amazon.com/network-firewall/latest/developerguide/getting-started.html

API Reference: https://docs.aws.amazon.com/network-firewall/latest/APIReference/

Service Quotas: https://docs.aws.amazon.com/network-firewall/latest/developerguide/quotas.html

AWS Organizations

General Information

AWS Organizations is a centralised account management and governance service that helps businesses manage multiple AWS accounts at scale. It allows automated account creation, hierarchical organisation through organisational units, and policy-driven governance across your entire AWS environment.

Main Benefits

AWS Organizations delivers unmatched centralised governance for up to 10,000 accounts while being completely free. It integrates with more than 40 AWS services, allowing up to 72% compute cost savings through shared Reserved Instances, and provides automated policy enforcement with enterprise-grade compliance capabilities.

Key Use Cases

- **Automated Account Management:** Quickly create and categorise AWS accounts with instant security policy application
- **Compliance Governance:** Apply service control policies to enforce security and compliance requirements organisation-wide
- **Centralised Security Monitoring:** Provide security teams read-only access across all accounts for threat identification
- **Resource Sharing:** Share critical resources like AWS Directory Service and Amazon VPC subnets across accounts

Core Features

- **Multi-account Management:** Create, organise, and manage multiple AWS accounts programmatically through AWS CLI, SDKs, or APIs
- **Organisational Units:** Group accounts into hierarchical organisational units to organise workflows and apply policies at scale
- **Service Control Policies:** Define and enforce preventive guardrails controlling access to AWS services, resources, and regions
- **Consolidated Billing:** Centralise billing and payment for all member accounts with volume discounts and single bill management
- **Trusted Access Integration:** Allow AWS services to perform tasks on behalf of the organisation through service-linked roles
- **Resource Control Policies:** Centrally prevent unintended use of AWS resources with fine-grained access controls across accounts
- **Tag Policies:** Standardise resource tagging across accounts for classification, tracking, and attribute-based access control
- **Centralised Security Monitoring:** Provide security teams read-only access and centrally manage Amazon GuardDuty, AWS Security Hub, and AWS Identity and Access Management (AWS IAM) Access Analyzer

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/organizations/
FAQ: https://aws.amazon.com/organizations/faqs/

Technical Overview

Backup & Recovery Capabilities

Backup: AWS Organizations metadata and policies are automatically maintained by AWS with high availability and durability across multiple availability zones. The service supports centralised backup governance through AWS Backup integration, allowing organisation-wide backup policies to be applied across accounts and OUs via AWS Management Console, CLI, and APIs.

Recovery: Organizations itself is highly available across AWS Regions with automatic failover capabilities. The service supports disaster recovery governance through centralised policies and resource sharing across accounts in different Regions. Supports customer RPO (centralized backup schedule enforcement) and RTO (multi-region account organization) objectives for member account workloads.

Setup Process

Onboarding: Create your organization using the AWS Management Console, AWS CLI, or APIs with administrative access and MFA enabled. Choose all features option, create organisational units hierarchically, configure service control policies for governance, and allow trusted access for integrated services.

Offboarding: Remove member accounts voluntarily or via the management account, transfer billing responsibilities, delete organizational units and policies. Complete organization deletion requires removing all accounts first with 90-day retention period.

Data Management

- **Storage Options:** Managed service configurations automatically maintained by AWS with high availability and durability
- **Data Removal:** Remove member accounts voluntarily or via management account with 90-day retention period
- **Cross-Region Replication:** Highly available across AWS regions with automatic failover capabilities for service metadata
- **Encryption:** HTTPS connectivity required for API operations with IAM service-linked roles for secure access

Optimise Cost and Consumption

AWS Organizations delivers consolidated billing that aggregates usage across member accounts to reach volume discount tiers faster. Reserved Instances and Savings Plans automatically share across the organization, maximizing utilization rates. Tag policies allow centralized cost allocation using Cost Categories for departmental chargeback. Service control policies prevent expensive resource deployment, while AWS Resource Access Manager reduces duplicate infrastructure costs through centralised resource sharing.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/organizations/latest/userguide/orgs_introduction.html
Getting Started: https://docs.aws.amazon.com/organizations/latest/userguide/orgs_getting-started.html
API Reference: https://docs.aws.amazon.com/organizations/latest/APIReference/Welcome.html
Service Quotas: https://docs.aws.amazon.com/organizations/latest/userguide/orgs_reference_limits.html

AWS Outposts

General Information

AWS Outposts is a fully managed service that extends AWS infrastructure, services, and APIs to customer premises. It allows organisations to run AWS Cloud services locally while maintaining integration with AWS Regions for hybrid cloud operations. AWS Outposts Family members include AWS Outposts rack and AWS Outposts servers.

Main Benefits

AWS Outposts delivers sub-millisecond latencies for real-time applications while taking away the infrastructure management overhead through fully managed services. Organisations achieve guaranteed data residency for compliance, reduce data transfer costs through local processing, and maintain identical APIs with AWS Regions for hybrid operations.

Key Use Cases

- **Low-Latency Applications:** Workloads requiring sub-millisecond latencies and real-time processing near end users
- **Data Residency and Compliance:** Applications keeping data on-premises for regulatory requirements and sovereignty mandates
- **Local Data Processing:** Processing large data volumes locally before sending insights to cloud
- **Edge Computing:** Running compute and storage services at edge locations for IoT and manufacturing

Core Features

- **Compute Services:** Amazon Elastic Compute Cloud (Amazon EC2) instances including M5, C5, R5, M7i, C7i, R7i, G4dn for specialised workloads
- **Storage Services:** Amazon EBS volumes with gp2 storage, Amazon Simple Storage Service (Amazon S3) on AWS Outposts buckets, and Amazon Elastic Block Store (Amazon EBS) Local Snapshots
- **Database Services:** Amazon Relational Database Service (Amazon RDS), Amazon DynamoDB, Amazon DocumentDB, and Amazon ElastiCache nodes running locally on AWS Outposts infrastructure
- **Container Services:** Amazon Elastic Container Service (Amazon ECS), Amazon Elastic Kubernetes Service (Amazon EKS), and AWS Fargate support for containerised applications with local orchestration
- **Networking:** VPC extension, local gateway connectivity, Application Load Balancers, and customer-owned IP addresses
- **Management Tools:** AWS Systems Manager, AWS Config, AWS Resource Access Manager with standard AWS APIs and SDKs
- **Hybrid Consistency:** Identical APIs, tools, and services as AWS Regions removing operational complexity differences
- **Fully Managed Service:** AWS handles automatic updates, patches, maintenance, monitoring, and hardware lifecycle management

Key Resources

Service Page: https://aws.amazon.com/outposts/rack/

FAQ: https://aws.amazon.com/outposts/rack/faqs/

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace. The Outposts configuration price includes delivery, installation, and servicing. Except as otherwise noted, Outposts prices are exclusive of applicable taxes and duties, including VAT and applicable sales tax.

Operating Systems

Prices do not include operating system (OS) charges. OS charges will be applied based on usage per instance.

AWS Services

AWS Services running on Outposts are priced based on usage by the hour per instance and exclude underlying EC2 instance and EBS storage charges which are already included in the Outposts configuration prices. This means you will not be double charged for the EC2 instances and EBS volumes that these services use.

Data Transfer

Data transfer from your Outpost to the parent AWS Region incurs no charge. Similarly, data transfer to and from your Outpost to your local network or to the internet also incurs no charge. Similar to data transfer charges in AWS Regions, data transfer out charges from the parent availability zone to the Outpost is the same as in the AWS Regions. Intra-AZ, inter-AZ and VPC data transfer charges in the AWS Region, remain the same.

Technical Overview

Setup Process

Onboarding: AWS Outposts deployment begins in the AWS Management Console, where customers create site configurations and specify their required compute and storage capacity. AWS manages delivery, installation, and activation process, while customers provide the necessary power, cooling, and network connectivity infrastructure. An Enterprise support subscription and appropriate IAM permissions are prerequisites for deployment. For custom configurations of AWS Outposts rack, customers should contact their AWS account manager. For Outposts rack, the order process follows a checklist of activities from capacity configuration, network assessment, site validation by AWS employees, through to order acceptance and installation by AWS. Outposts servers are shipped to customers with full instructions for self-service installation, or installation by an AWS preferred third-party contractor. Once installation is complete, AWS will remotely provision compute and storage resources and activate the customer's Outpost. Outposts service billing will begin the day after the Outpost is activated.

Customers will need to support the process and provide site access and necessary information including but not limited to networking information to finalise the order placement and service activation. AWS reserves the right to cancel the order process.

Offboarding: Customers must first terminate all Amazon EC2 instances, delete associated Amazon EBS volumes and Amazon S3 buckets, and remove VPC associations. Data security is supported through AWS Nitro Security Key destruction, which cryptographically shreds all information on the hardware. Automated cleanup options are available to streamline the retrieval process.

Customer Responsibilities

- AWS Outposts have an updated shared responsibility model underlying security. AWS is responsible for protecting infrastructure for Outposts similar to how it secures infrastructure in the cloud today. Customers are responsible for securing applications

running on Outposts as well as the physical security of the Outpost and ensuring consistent networking to them.

- Customers must subscribe and remain enrolled in AWS Support at the Enterprise level during the entire period of their use of Outposts.
- For the duration of their use of Outpost, customers are responsible for ensuring at all times, that their site meets the minimum requirements necessary to support the installation, maintenance, use, and removal of Outposts as described in the Outpost user guide here: https://docs.aws.amazon.com/outposts/latest/userguide/outposts-requirements.html (for Outposts rack) and https://docs.aws.amazon.com/outposts/latest/server-userguide/outposts-requirements.html (for Outposts servers) or as may be provided to the customer during the Outpost order process.
- Customers are governed by the service responsibilities as set out in the AWS service terms here: https://aws.amazon.com/service-terms/. AWS may terminate a customer's use of Outposts and remove the Outposts Equipment if the customer breaches the service terms with respect to Outposts.

Service Levels and Compliance Standards

The Service Terms for any Services that run locally on AWS Outposts also apply to your use of those Services on AWS Outposts. There are inherent differences between Services running locally on AWS Outposts from those Services running at AWS operated facilities because the Outposts Equipment is physically located at the Designated Facility where you are responsible for physical security and access controls, as well as all power, networking, and environmental conditions. Due to these differences:

a) The Service Level Agreements for any Services that run locally on AWS Outposts do not apply to your use of those Services on AWS Outposts.
b) Any AWS commitments in the Agreement that depend on AWS's operation of such physical security and access controls, or power, networking, and environmental conditions, do not apply to AWS Outposts or any Services running locally on AWS Outposts.
c) The specific compliance and assurance programs for which AWS Outposts are in scope are listed here: https://aws.amazon.com/compliance/services-in-scope/. Note that, for other Services listed here, those Services are not in scope when running locally on AWS Outposts unless AWS Outposts is also separately listed for the specific compliance or assurance program.

Service Constraints

Connectivity

An Outpost requires connectivity to the parent AWS Region. Outposts are not designed for disconnected operations or environments with limited to no connectivity. We recommend that customers have redundant and highly available network connections back to their AWS Region.

Please refer to the networking section of the AWS Outposts High Availability Design and Architecture Considerations whitepaper to learn how to architect for redundant and highly available network connections. Please visit the Outposts Rack FAQs and Outposts Servers FAQs to understand what will happen if your external network connectivity is unavailable.

Site Assessment

A site and network assessment must be completed and any gaps discovered by these assessments must be addressed prior to acceptance of your Outposts order by AWS.

Backup & Recovery Capabilities

Backup: Amazon EBS snapshots stored locally on Amazon S3 on AWS Outposts or in AWS Regions for geographic redundancy can be created manually through APIs, CLI, AWS Management Console or automated via AWS Backup with configurable scheduling policies. Amazon S3 object versioning and Cross-Region Replication provide data protection.

Recovery: AWS Elastic Disaster Recovery continuously replicates instances between AWS Outposts and AWS Regions, multi-AZ deployments across multiple AWS Outposts provide local redundancy, and Application Load Balancers allow automatic failover scenarios. Supports customer RPO objectives measured in minutes through local snapshots and RTO targets of minutes to hours via automatic failover mechanisms.

Data Management

Storage Options: Amazon EBS volumes with gp2 storage, Amazon S3 on AWS Outposts buckets, and third-party block storage integration

Data Removal: Terminate Amazon EC2 instances, delete Amazon EBS volumes and Amazon S3 buckets, destroy AWS Nitro Security Key for cryptographic shredding

Cross-Region Replication: Amazon S3 Cross-Region Replication protects object data between local Amazon S3 on AWS Outposts and Regional buckets

Encryption: Enterprise-grade security with AWS Nitro Security Key allowing complete data destruction through cryptographic shredding capabilities

Optimise Cost and Consumption

Right-sizing AWS Outposts configurations to match workload requirements prevents over-provisioning and optimises capacity usage. Local data processing reduces data transfer costs to AWS Regions while local Amazon S3 storage minimises Regional storage expenses. AWS Resource Access Manager allows cost sharing across multiple accounts, maximising shared infrastructure investments. Consolidating multiple workloads on single AWS Outposts improves hardware usage rates and reduces per-workload operational costs.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/outposts/latest/userguide/what-is-outposts.html
Getting Started: https://docs.aws.amazon.com/outposts/latest/userguide/get-started-outposts.html
API Reference: https://docs.aws.amazon.com/outposts/latest/APIReference/Welcome.html
Service Quotas: https://docs.aws.amazon.com/outposts/latest/userguide/outposts-limits.html

AWS Parallel Computing Service (PCS)

General Information

AWS Parallel Computing Service (AWS PCS) is a managed service that simplifies running and scaling high performance computing workloads using Slurm. It integrates AWS compute, storage, and networking resources while providing built-in management and observability capabilities for scientific and engineering applications.

Main Benefits

AWS PCS removes operational overhead through fully managed Slurm infrastructure with automatic scaling and maintenance. Organisations achieve faster time-to-science with instant access to virtually unlimited computational resources, avoiding traditional HPC queue times. Spot instance support reduces compute costs by up to 90% for suitable workloads.

Key Use Cases

- **Computational Chemistry:** Chemical formulation discovery for pharmaceutical research and materials science applications
- **Scientific Simulations:** Computational fluid dynamics, structural analysis, and engineering modelling workloads
- **Healthcare Research:** Genomic sequencing, medical imaging analysis, and clinical research simulations
- **Aerospace Modelling:** Mission planning, space operations, and vehicle innovation simulation workloads

Core Features

- **Managed Slurm Scheduler:** Fully managed Slurm workload manager with automated scaling and maintenance
- **Elastic Compute Node Groups:** Automatically scaling Amazon Elastic Compute Cloud (Amazon EC2) instances with customisable instance types and purchase options
- **Queue Management:** Job scheduling and execution across compute node groups with customisable policies
- **Built-in Observability:** Job completion logs, monitoring, and cluster management through the AWS console
- **Slurm Customisation:** Over 60 configurable Slurm parameters for scheduling, resource allocation, and access control
- **Managed Accounting:** Slurm accounting feature for usage tracking, resource limits, and chargeback reporting
- **Capacity Blocks Integration:** Support for Amazon EC2 Capacity Blocks to reserve GPU instances up to 8 weeks
- **REST API Support:** HTTP interface for programmatic interaction with Slurm workload manager

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/pcs/
FAQ: https://aws.amazon.com/pcs/faqs/

Technical Overview

Backup & Recovery Capabilities

Backup: Customer data on Amazon Elastic File Store (Amazon EFS) and Amazon FSx for Lustre can be protected through Amazon Elastic Block Store (Amazon EBS) snapshots and AMI creation (AWS Management Console/AWS CLI) or automated via AWS Backup policies with scheduled frequencies. Job metadata and logs are durably stored in Amazon CloudWatch Logs or Amazon Simple Storage Service (Amazon S3) with built-in redundancy.

Recovery: New clusters can be provisioned within minutes using pre-configured AWS CloudFormation templates, cross-Region deployment for disaster recovery scenarios, and AWS Elastic Disaster Recovery integration. Supports customer RPO objectives (frequent automated snapshots of shared storage) and RTO objectives (cluster deployment and elastic scaling capabilities).

Setup Process

Onboarding: Setup through AWS Management Console or AWS CLI interfaces with VPC creation, security group configuration, and shared storage setup using Amazon EFS and Amazon FSx for Lustre. AWS CloudFormation templates automate cluster deployment with customisable Slurm scheduler and compute node group parameters.

Offboarding: Delete cluster resources via AWS Management Console or AWS CLI automatically terminating Amazon EC2 instances and removing configurations. Customer data in Amazon EFS and Amazon FSx file systems requires manual deletion with irreversible cleanup process.

Data Management

- **Storage Options:** Amazon EFS and Amazon FSx for Lustre provide shared storage for home directories and high-performance workloads
- **Data Removal:** Customer data in Amazon EFS and Amazon FSx requires manual deletion; job logs automatically removed when clusters deleted
- **Cross-Region Replication:** Supported through Amazon EFS and Amazon FSx cross-Region backup and restore capabilities for disaster recovery scenarios
- **Encryption:** Use AWS built-in encryption capabilities through integrated Amazon EFS, Amazon FSx, and Amazon EBS storage services

Optimise Cost and Consumption

AWS PCS provides elastic auto-scaling that dynamically adds and removes compute nodes based on job queue demand. Amazon EC2 Spot instance support delivers up to 90% cost reduction for fault-tolerant workloads. Capacity Blocks integration allows GPU instance reservations up to 8 weeks in advance with discounted rates. Mixed instance types within compute node groups optimise price-performance ratios. Managed accounting features provide detailed usage tracking for chargeback and resource optimisation.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/pcs/latest/userguide/what-is-service.html
Getting Started: https://docs.aws.amazon.com/pcs/latest/userguide/getting-started.html
API Reference: https://docs.aws.amazon.com/pcs/latest/APIReference/Welcome.html
Service Quotas: https://docs.aws.amazon.com/pcs/latest/userguide/service-endpoints-quotas.html

AWS Payment Cryptography

General Information

AWS Payment Cryptography is a managed service providing cryptographic functions and key management for payment processing in Payment Card Industry (PCI) compliance. It allows payment organisations to perform operations like card verification code (CVV) generation, PIN translation, and transaction authorisation using certified HSM hardware without managing on-premises infrastructure.

Main Benefits

AWS Payment Cryptography takes away the burden of managing costs for on-premises HSM infrastructure while delivering industry-leading availability SLAs and automatic failover across three availability zones. The service provides elastic scaling with tiered pricing, built-in PCI compliance using certified hardware, and electronic key exchange that replaces manual paper-based processes for faster payment application deployment.

Key Use Cases

- **Card Issuing Operations:** Europay, Mastercard, Visa (EMV) PIN changes, CVV/CVV2 generation and validation, ARQC processing for authorisation
- **Payment Terminal Processing:** Data decryption using Triple Data Encryption Standard (TDES), Advanced Encryption Standard (AES), Derived Unique Key Per Transaction (DUKPT) with PCI Point-to-Point Encryption (PCI P2PE) compliance verification
- **Key Exchange Operations:** Electronic key sharing with partners using TR-34, TR-31, ECDH protocols
- **PIN Translation Services:** Secure PIN translation between encryption keys and MAC verification operations

Core Features

- **Key Management:** Create, manage, import, export symmetric and asymmetric keys with alias management and tagging
- **Cryptographic Operations:** Encrypt, decrypt, re-encrypt data, generate/validate CVV/CVV2/ARQC, PIN translation, MAC generation and verification
- **Multi-Region Replication:** Automatic key replication across AWS Regions for high availability and disaster recovery
- **PCI Compliance:** Built-in compliance with PCI PIN, PCI P2PE, PCI DSS using certified hardware
- **Electronic Key Exchange:** Support for TR-34, TR-31, ECDH protocols for secure partner key sharing
- **Post-Quantum Security:** Hybrid post-quantum TLS support using ML-KEM cryptography for future-proof API protection
- **Certified Hardware:** PCI PIN Transaction Security (PCI PTS) HSM V3 and FIPS 140-2 Level 3 certified infrastructure
- **High Availability:** Automatic failover across three availability zones with industry-leading availability SLAs

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/payment-cryptography/
FAQ: https://aws.amazon.com/payment-cryptography/faqs/

Technical Overview

Backup & Recovery Capabilities

Backup: AWS Payment Cryptography provides inherent backup through HSM-based protection and encrypted storage with durability matching AWS's highest durability services. Multi-Region replication automatically distributes keys across regions using TR-34 and TR-31 protocols. Key material is protected by HSM AES 256 main keys with automatic replication across availability zones.

Recovery: Built-in disaster recovery through multi-Region key replication with near real-time synchronisation and automatic availability zone failover with zero planned downtime. Cross-region key distribution supports business continuity during regional outages. Supports aggressive customer RPO objectives through near real-time replication and RTO objectives through automatic failover capabilities.

Setup Process

Onboarding: Setup requires IAM permissions and AWS CLI, SDKs, or API access. Create cryptographic keys using Control Plane API, specifying algorithms and usage. Import existing keys from on-premises HSMs using TR-31/TR-34 protocols or create new keys with immediate cryptographic operation access.

Offboarding: Schedule keys for deletion with configurable waiting period (3-180 days, default 7 days). Keys enter DELETE_PENDING state and become unusable. Deletion permanently removes key material and metadata. Multi-Region key deletion automatically removes replica keys across all Regions.

Data Management

- **Storage Options:** HSM-based protection with encrypted databases and in-memory replication across three availability zones
- **Data Removal:** Scheduled key deletion with configurable waiting period, irreversibly removing key material and metadata
- **Cross-Region Replication:** Automatic key replication across AWS Regions with eventual consistency for disaster recovery
- **Encryption:** HSM AES 256 main keys protect key material with highest durability service standards

Optimise Cost and Consumption

AWS Payment Cryptography provides tiered key pricing models that scale with volume, offering cost-effective scaling for high-volume cryptographic operations. Elastic scaling charges only for actual usage rather than provisioned capacity, optimising costs based on transaction volumes and seasonal payment processing patterns. Multi-Region replication consolidates disaster recovery infrastructure requirements while electronic key exchange capabilities streamline partner integration workflows.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/payment-cryptography/latest/userguide/what-is.html
Getting Started: https://docs.aws.amazon.com/payment-cryptography/latest/userguide/getting-started.html
API Reference: https://docs.aws.amazon.com/payment-cryptography/latest/APIReference/Welcome.html
Service Quotas: https://docs.aws.amazon.com/payment-cryptography/latest/userguide/quotas.html

AWS PrivateLink

General Information

AWS PrivateLink allows private connectivity between VPCs, AWS services, and on-premises networks without using the public internet. It creates secure connections using private IP addresses within Amazon's network infrastructure, removing internet gateways and NAT devices and reducing the attack surface.

Main Benefits

AWS PrivateLink delivers security by removing internet exposure and reducing attack surface. It provides up to 92% cost savings compared to traditional connectivity while automatically scaling to handle millions of requests per second through Amazon's high-speed backbone network with consistent low-latency performance.

Key Use Cases

- **SaaS Application Access:** Privately access third-party SaaS applications without internet exposure or complex configurations
- **Shared Services Architecture:** Allow common services consumption across separate VPCs without VPC peering limitations
- **Hybrid Cloud Integration:** Extend services to on-premises resources via network load balancers and data centers
- **Microservices Communication:** Facilitate secure inter-service communication with controlled access between development teams

Core Features

- **Interface VPC Endpoints:** Elastic network interfaces providing private connectivity to AWS services and third-party applications
- **VPC Endpoint Services:** Allow service providers to expose applications behind Network Load Balancers for private consumption
- **Resource VPC Endpoints:** Direct private connectivity to specific VPC resources like databases without requiring load balancing
- **Cross-Region Connectivity:** Extends private connectivity to AWS services across Regions without internet or inter-Region peering
- **Security Group Integration:** Apply VPC security groups at network interface level for granular access control policies
- **Endpoint Policies:** Fine-grained access control policies restricting access to specific AWS principals, actions, and resources
- **DNS Resolution:** Private DNS names for endpoints offering application integration without requiring code changes
- **High Availability:** Horizontally scalable and redundant across multiple Availability Zones providing automatic fault tolerance

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/privatelink/

FAQ: https://aws.amazon.com/privatelink/faqs/

Technical Overview

Backup & Recovery Capabilities

Backup: AWS PrivateLink does not have inherent backup capabilities as a networking connectivity service. Configuration settings for VPC endpoints and endpoint services can be captured through AWS Config and replicated via AWS CloudFormation templates for infrastructure-as-code deployments.

Recovery: AWS PrivateLink supports disaster recovery through cross-Region connectivity and multi-Availability Zone (multi-AZ) deployments for high availability. The service supports failover scenarios via DNS-based routing and applies inherent redundancy across availability zones. Supports customer RPO objectives (low-latency connectivity with automatic scaling) and RTO objectives (predictable recovery timeframes through consistent private network performance).

Setup Process

Onboarding: Set up VPC endpoints using the AWS Management Console, AWS CLI, or AWS CloudFormation templates. Configure your DNS settings, security groups, and endpoint policies to control access. Plan your subnet distribution across multiple Availability Zones to maintain high availability.

Offboarding: Remove VPC endpoints using the AWS Management Console, AWS CLI, or API to instantly stop connectivity. The deletion process automatically cleans up associated endpoint policies, security group rules, and DNS configurations, eliminating all connections.

Data Management

- **Storage Options:** AWS PrivateLink is a networking connectivity service that does not store data itself
- **Data Removal:** Deleting VPC endpoints immediately terminates connectivity with no residual data remaining in AWS PrivateLink
- **Cross-Region Replication:** Supports cross-Region connectivity allowing private communication over Amazon's network infrastructure without internet peering
- **Encryption:** Traffic remains within Amazon's private network infrastructure using secure private IP addresses for connectivity

Optimise Cost and Consumption

Reduce endpoint charges by consolidating multiple service connections through single VPC endpoints, which also simplifies your network architecture. Use Resource VPC endpoints for cost-effective database access that takes away the need for load balancers. Implement endpoint policies that restrict access to specific principals and resources to avoid unnecessary data processing costs. For enterprise applications, use shared endpoint services to distribute expenses across multiple consumers.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/vpc/latest/privatelink/what-is-privatelink.html

Getting Started: https://docs.aws.amazon.com/vpc/latest/privatelink/getting-started.html

API Reference: https://docs.aws.amazon.com/AWSEC2/latest/APIReference/operation-list-privatelink.html

Service Quotas: https://docs.aws.amazon.com/vpc/latest/privatelink/vpc-limits-endpoints.html

AWS Resilience Hub

General Information

AWS Resilience Hub is a centralised service for managing and improving application resilience on AWS. It continuously assesses applications against RTO/RPO targets, provides AI-powered recommendations, and integrates chaos engineering to transform reactive disaster recovery into proactive resilience management.

Main Benefits

AWS Resilience Hub transforms reactive disaster recovery into proactive resilience engineering through centralised, automated assessments and AI-powered recommendations. It provides unified visibility across multi-account architectures, cost-optimised resilience strategies, and integrated chaos engineering to prevent costly outages while reducing operational overhead.

Key Use Cases

- **Application Resilience Assessment:** Continuously assess applications against RTO/RPO targets across failure scenarios
- **Proactive Disaster Recovery Planning:** Define and validate disaster recovery strategies using chaos engineering experiments
- **Multi-Account Resilience Management:** Centrally manage resilience across distributed, multi-account cloud architectures using hub-and-spoke model
- **CI/CD Pipeline Integration:** Integrate resilience validation into development lifecycle to identify weaknesses before production

Core Features

- **Resilience Dashboard:** Visual dashboard providing overview of application resilience status, compliance, and resiliency scores
- **Application Discovery and Management:** Automatic discovery and mapping of AWS resources from AWS CloudFormation, AWS Resource Groups, Amazon Elastic Kubernetes Service (Amazon EKS) clusters
- **Resiliency Policy Management:** Define and manage customisable resilience policies with specific RTO and RPO targets
- **Automated Assessments:** Continuous resilience assessments with scheduled daily evaluations, drift detection capabilities, and compliance reporting
- **AWS FIS Integration:** Integrated chaos engineering capabilities using AWS Fault Injection Service (AWS FIS) with tailored fault injection experiments to validate application resilience
- **Operational Recommendations:** AI-powered recommendations for infrastructure improvements, alarm configurations, and standard operating procedures based on the AWS Well-Architected Framework
- **Multi-Account Support:** Hub-and-spoke architecture allowing centralised resilience management across multiple AWS accounts with cross-account roles
- **Resiliency Scoring:** Quantitative scoring system measuring application readiness to handle disruptions based on recommendations and tests

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/resilience-hub/

FAQ: https://aws.amazon.com/resilience-hub/faqs/

Technical Overview

Backup & Recovery Capabilities

Backup: AWS Resilience Hub evaluates and recommends AWS Backup integration for Amazon Relational Database Service (Amazon RDS), Amazon Elastic File Storage (Amazon EFS), Amazon Elastic Block Store (Amazon EBS), and Amazon DynamoDB resources. The service assesses backup plan configurations and identifies coverage gaps through AWS Management Console or API operations. It generates specific recommendations for implementing backup strategies that align with defined RPO requirements.

Recovery: AWS Resilience Hub supports disaster recovery through warm standby, pilot light, and backup-and-restore approaches with cross-Region replication and multi-Availability Zone (multi-AZ) deployments. Chaos engineering experiments validate recovery procedures across application, infrastructure, AZ, and Regional failure scenarios. Supports customer RPO and RTO objectives.

Setup Process

Onboarding: Setup requires configuring IAM permissions including AWSResilienceHubAssessmentRole and adding applications from AWS CloudFormation stacks, AWS Resource Groups, Amazon EKS clusters, or Terraform state files. Create resiliency policies with RTO/RPO targets and configure notification service using Amazon SNS topics for drift notifications. Six-month free trial available for first three applications.

Offboarding: Delete applications through console or DeleteApp API operation, which permanently removes all versions, assessments, and metadata. Export required metrics before deletion as the process cannot be undone. AWS Resilience Hub automatically cleans up resources and stops charging upon application deletion.

Data Management

- **Storage Options:** Assessment data, policies, and metrics stored in AWS managed storage with Amazon Simple Storage Service (Amazon S3) integration
- **Data Removal:** Delete applications via AWS Management Console or DeleteApp API; data retained 365 days before automatic deletion
- **Cross-Region Replication:** Supports multi-Region architectures but limited cross-region assessments for Amazon EKS clusters in opt-in Regions
- **Encryption:** Uses built-in AWS encryption standards; specific encryption details follow standard AWS service encryption practices

Optimise Cost and Consumption

AWS Resilience Hub's multi-dimensional recommendation engine provides four distinct optimisation strategies: cost optimisation, minimal changes, AZ recovery, and Region recovery. The service prevents over-provisioning by delivering precise recommendations based on actual RTO/RPO requirements and right-sizing disaster recovery resources. Automated drift detection reduces operational overhead while optimising backup strategies and resource usage across multi-account architectures.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/resilience-hub/latest/userguide/what-is.html

Getting Started: https://docs.aws.amazon.com/resilience-hub/latest/userguide/getting-started.html

API Reference: https://docs.aws.amazon.com/resilience-hub/latest/APIReference/Welcome.html

Service Quotas: https://docs.aws.amazon.com/general/latest/gr/resiliencehub.html

AWS Resource Access Manager

General Information

AWS Resource Access Manager (AWS RAM) allows secure sharing of AWS resources across accounts, organisations, and identity and access management (IAM) roles without duplicating infrastructure. It reduces operational overhead by centralising resource management while maintaining security through managed permissions and resource-based policies.

Main Benefits

AWS RAM takes away duplicate resource provisioning across accounts, reducing infrastructure costs and operational overhead with no additional service charges. It supports more than 40 AWS services, supporting centralised governance, and provides built-in resource integration with audit capabilities through AWS CloudTrail and Amazon CloudWatch integration.

Key Use Cases

- **Multi-Account Resource Sharing:** Share transit gateways, subnets, and Amazon Route 53 Resolver rules across accounts without duplication
- **Centralised Infrastructure Management:** Create resources once in central account and share with others for consistency
- **Cross-Account Network Connectivity:** Share VPC resources and networking components for secure connectivity with centralised control
- **License Management:** Share AWS License Manager configurations and licensed software across multiple accounts efficiently

Core Features

- **Resource Sharing:** Share AWS resources across accounts, organisations, organisational units, and IAM roles/users securely
- **Managed Permissions:** Control access with AWS managed permissions or create custom customer managed permissions
- **Resource-based Policies:** Automatically generate and manage resource-based policies for all shared resources
- **Cross-account Visibility:** View shared resources as built-in resources in consuming accounts with full visibility
- **Organisations Integration:** Integration with AWS Organizations for organisation-wide resource sharing and management
- **Regional and Global Resources:** Support for sharing both Regional and global AWS resources across accounts
- **Invitation Management:** Handle resource share invitations for accounts outside the organisation with email notifications
- **Audit and Monitoring:** Integration with AWS CloudTrail and Amazon CloudWatch for logging and monitoring resource sharing activities

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/ram/
FAQ: https://aws.amazon.com/ram/faqs/

Technical Overview

Backup & Recovery Capabilities

Backup: AWS RAM does not provide inherent backup capabilities as it is a resource sharing service. However, RAM supports sharing of AWS Backup vaults, backup resources, and recovery points across accounts, facilitating centralised backup management strategies through shared backup infrastructure.

Recovery: RAM facilitates disaster recovery by allowing sharing of disaster recovery resources across accounts and regions, supporting multi-account disaster recovery architectures. RAM indirectly supports customer RPO objectives (through shared backup resources) and RTO objectives (by reducing recovery complexity via simplified resource access patterns).

Setup Process

Onboarding: Setup through AWS RAM console, AWS CLI, or SDKs by creating resource shares, specifying resources, selecting managed permissions, and defining principals. Enable AWS Organizations sharing first for organisation-wide access. Service automatically creates service-linked roles and resource-based policies.

Offboarding: Stop sharing by removing resources from shares or deleting resource shares entirely, immediately revoking all principal access. Leave unwanted resource shares to revoke access while resource owners retain full control.

Data Management

- **Storage Options:** AWS RAM manages resource sharing metadata; actual resources remain stored in their original AWS services
- **Data Removal:** Delete resource shares to immediately revoke all principal access while resource owners retain full control
- **Cross-Region Replication:** Supports sharing both regional and global AWS resources across accounts and Regions as needed
- **Encryption:** Inherits encryption capabilities from shared resources; complies with Payment Card Industry Data Security Standard (PCI DSS), Service Organization Controls (SOC), and ISO standards

Optimise Cost and Consumption

AWS RAM optimises costs through resource consolidation strategies that prevent duplicate provisioning of expensive infrastructure like transit gateways, dedicated hosts, and Amazon Aurora clusters across accounts. Resource tagging allows granular cost allocation tracking and visibility into shared resource consumption patterns. Organisations achieve better usage rates by sharing underused resources, avoiding over-provisioning scenarios while maintaining centralised governance and operational efficiency.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/ram/latest/userguide/what-is.html
Getting Started: https://docs.aws.amazon.com/ram/latest/userguide/getting-started.html
API Reference: https://docs.aws.amazon.com/ram/latest/APIReference/Welcome.html
Service Quotas: https://docs.aws.amazon.com/ram/latest/userguide/service-quotas.html

AWS Secrets Manager

General Information

AWS Secrets Manager is a secrets management service offering a secure and scalable method to store and manage secrets to protect access to your applications, services, and IT resources. The service automatically rotates, manages, and retrieves database credentials, API keys, and other secrets throughout their lifecycle. Using AWS Secrets Manager, you can secure and manage secrets used to access resources in the AWS Cloud, on third-party services, and on-premises.

Main Benefits

Automatically manages and updates your credentials without requiring custom code or manual work. Delivers 99.95% SLA with 40,000 requests/second capacity, cross-Region replication, and enterprise-grade encryption. Reduces security risks while providing in-depth auditing and AWS service integration.

Key Use Cases

- **Database Credential Management:** Store and automatically rotate Amazon Relational Database Service (Amazon RDS), Amazon DocumentDB, and Amazon Redshift credentials without disruption
- **Application Security:** Replace hard-coded API keys and OAuth tokens with dynamic secret retrieval
- **Third-Party SaaS Integration:** Manage and rotate credentials for external services like Salesforce and Snowflake
- **Multi-Region Disaster Recovery:** Replicate secrets across AWS Regions for business continuity and low-latency access

Core Features

- **Automatic Rotation:** Schedule-based or on-demand rotation of secrets without disrupting applications for AWS and third-party services
- **Encryption:** End-to-end encryption at rest using AWS Key Management Service (AWS KMS) keys and in transit using TLS with customer-managed key support
- **Cross-Region Replication:** Automatic replication of secrets to multiple AWS Regions for disaster recovery and cross-regional redundancy
- **Fine-Grained Access Control:** AWS Identity and Access Management (AWS IAM) integration for granular permissions, resource-based policies, and attribute-based access control support
- **Managed External Secrets:** Managed rotation for third-party SaaS providers with predefined formats and integrated partner applications.
- **In-Depth Auditing:** Complete logging through AWS CloudTrail, monitoring via Amazon CloudWatch, and integration with security services like Amazon GuardDuty
- **Caching Support:** Built-in caching libraries and AWS Lambda extensions to improve performance and reduce costs when retrieving secrets
- **Batch Operations:** Ability to retrieve multiple secrets in a single API call for improved efficiency in applications

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/secrets-manager/

FAQ: https://aws.amazon.com/secrets-manager/faqs/

Technical Overview

Backup & Recovery Capabilities

Backup: Versioning mechanism maintains multiple secret versions with staging labels (AWSCURRENT, AWSPREVIOUS, AWSPENDING) automatically during credential rotation or manually via the AWS Management Console, AWS CLI, or SDKs. Cross-Region replication provides geographic backup with synchronised replicas maintaining independent versions while staying current with primary secrets.

Recovery: Secret restoration available within configurable 7–30-day recovery window, version rollback capabilities, and cross-region failover with application access. Multi-AZ deployment with 99.95% SLA provides fault tolerance. Supports low RPO objectives (near real-time cross-region synchronisation) and RTO objectives (10,000 requests/second GetSecretValue capacity with cached retrieval).

Setup Process

Onboarding: Create secrets through AWS Management Console, AWS CLI, or SDKs with built-in tutorials for migrating hard-coded credentials and database rotation setup. Requires SecretsManagerReadWrite IAM policy and proper VPC configuration for private access.

Offboarding: DeleteSecret API operation with configurable 7–30-day recovery window before permanent deletion. Replicated secrets across Regions must be deleted separately with all activities logged through AWS CloudTrail for auditing.

Data Management

- **Storage Options:** Versioned secret storage with staging labels, maximum 64KB secret value size, multi-Availability Zone deployment
- **Data Removal:** Configurable 7–30-day recovery window with DeleteSecret API operation removing all versions and metadata
- **Cross-Region Replication:** Automatic replication across multiple Regions with synchronised copies and Region-specific endpoints for disaster recovery
- **Encryption:** End-to-end encryption at rest using AWS KMS keys and in transit using TLS with customer-managed key support

Optimise Cost and Consumption

Built-in caching libraries and AWS Lambda extensions reduce API call costs for high-frequency secret retrieval operations. Batch operations allow retrieving multiple secret credentials in single API calls, improving application efficiency. Cost allocation tags provide granular usage tracking across business units for precise chargeback management. Secret consolidation strategies reduce per-secret monthly charges while managed rotation capabilities streamline operational overhead.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/secretsmanager/latest/userguide/intro.html

Getting Started: https://docs.aws.amazon.com/secretsmanager/latest/userguide/intro.html

API Reference: https://docs.aws.amazon.com/secretsmanager/latest/apireference/Welcome.html

Service Quotas:
https://docs.aws.amazon.com/secretsmanager/latest/userguide/reference_limits.html

AWS Security Hub

General Information

AWS Security Hub is a unified cloud security solution that correlates and enriches security signals from multiple AWS services including GuardDuty, Inspector, and Macie. It transforms complex security data into actionable insights through automated correlation, exposure findings, and attack path visualization for comprehensive threat response.

Main Benefits

Security Hub eliminates disparate security tools through unified operations, delivering 30-70% cost reduction via automated workflows and consolidated management. Near real-time correlation transforms overwhelming alerts into prioritized insights, while supporting up to 10,000 member accounts with cross-Region aggregation and 1-year historical data retention.

Key Use Cases

- **Centralized Security Operations:** Unify security management across multiple AWS accounts with consolidated visibility
- **Security Posture Management:** Monitor cloud configurations against CIS, PCI DSS, and NIST frameworks
- **Threat Detection and Response:** Correlate findings to identify attack paths and automate workflows
- **Compliance Monitoring:** Automated security checks against industry standards and regulatory requirements

Core Features

- **Exposure Findings:** Correlates findings across services to detect exploitable vulnerabilities and misconfigurations
- **Attack Path Visualization:** Interactive visualization showing how attackers can access and control resources
- **Automated Correlation:** Automatically correlates and enriches security signals across multiple AWS services
- **Multi-Account Support:** Centralized management across AWS Organizations with cross-Region aggregation capabilities
- **Security Standards Compliance:** Automated checks against CIS, PCI DSS, NIST, and AWS best practices
- **Automation Rules:** Automated response workflows with Jira Cloud and ServiceNow ITSM integration
- **OCSF Integration:** Generates findings in Open Cybersecurity Schema Framework for standardized data exchange
- **Historical Trends:** Up to 1 year of data with customizable widgets for comprehensive analysis

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/security-hub

FAQ: https://aws.amazon.com/security-hub/features/

Technical Overview

Backup & Recovery Capabilities

Backup: Security Hub does not provide traditional backup capabilities as it focuses on security findings aggregation rather than data backup. The service maintains findings for 90 days with automatic retention management. Organizations requiring backup of security configurations must implement custom solutions using APIs or third-party integrations for long-term storage.

Recovery: Security Hub supports disaster recovery through cross-Region aggregation, allowing findings from multiple regions to be viewed centrally. Multi-AZ architecture provides inherent resilience with automatic replication across Availability Zones. Supports low RPO through continuous monitoring and RTO objectives via cross-Region aggregation for rapid visibility restoration.

Setup Process

Onboarding: Enable Security Hub through AWS Console with straightforward setup. Prerequisites include AWS Config and proper IAM permissions. Service automatically creates service-linked roles and offers 30-day free trial. Organizations can designate delegated administrators for multi-account management.

Offboarding: Disable service through AWS Console or API, removing dashboard access and stopping finding ingestion. Existing findings retained for 90 days before automatic deletion. Service-linked roles automatically cleaned up upon disabling.

Data Management

- **Storage Options:** Managed service storing security findings with 90-day automatic retention and 1-year historical data
- **Data Removal:** Findings automatically deleted after 90-day retention period or manually archived before offboarding
- **Cross-Region Replication:** Supports cross-Region aggregation allowing findings from multiple regions to be centrally viewed
- **Encryption:** AWS-managed infrastructure with automatic replication across multiple Availability Zones for data durability

Optimize Cost and Consumption

Security Hub's Cost Estimator feature enables organizations to compare deployment scenarios and export estimates for accurate budgeting. Resource-based pricing anchored on EC2 instances provides predictable costs while volume pricing reduces per-resource expenses for larger deployments. Automated correlation workflows minimize manual security analysis time, and cross-Region aggregation optimizes regional infrastructure requirements for comprehensive security operations management.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/securityhub/latest/userguide/index.html

Getting Started: https://docs.aws.amazon.com/securityhub/latest/userguide/securityhub-v2-enable.html

API Reference: https://docs.aws.amazon.com/securityhub/1.0/APIReference/Welcome.html

Service Quotas: https://docs.aws.amazon.com/general/latest/gr/sechub.html

AWS Security Hub – Cloud Security Posture Management (CSPM)

Feature of AWS Security Hub

General Information

AWS Security Hub Cloud Security Posture Management (AWS Security Hub CSPM) is a cloud security posture management service that continuously evaluates AWS resources against proven security practices and compliance frameworks. It centralises security findings across multiple accounts and Regions, providing automated security scoring and remediation capabilities.

Main Benefits

Security Hub CSPM removes complexity by consolidating multiple security tools into a single platform, reducing operational overhead and licensing costs. It supports up to 10,000 member accounts per Region with automated security scoring, cross-Region aggregation, and in-depth integration for streamlined remediation workflows.

Key Use Cases

- **Continuous Security Posture Assessment:** Automatically evaluate AWS resources against proven security practices with real-time monitoring and scoring
- **Multi-Account Security Management:** Centrally manage security across AWS Organizations with delegated administration and cross-Region aggregation
- **Compliance Monitoring and Reporting:** Track compliance status against industry standards like Center for Internet Security (CIS), Payment Card Industry Data Security Standard (PCI-DSS), and National Institute of Standards and Technology (NIST) frameworks
- **Security Finding Consolidation:** Aggregate findings from AWS services and third-party products in standardised AWS Security Finding Format (ASFF) format

Core Features

- **Automated Security Checks:** Performs continuous security practice checks against AWS resources using prepackaged security standards and custom controls
- **Cross-Region Aggregation:** Consolidates findings and provides unified view across multiple AWS Regions from a single aggregation Region
- **Security Score Calculation:** Evaluates security posture with automated scoring based on control status across all allowed standards
- **Central Configuration:** Delegated administrators can configure AWS Security Hub CSPM settings across multiple accounts and organizational units
- **Automation Rules:** Automatically updates findings, triggers responses, and integrates with ticketing, Security Orchestration, Automation, and Response (SOAR) and Security Information and Event Management (SIEM) tools
- **Finding Consolidation:** Aggregates findings from AWS services and partner integrations in standardised AWS Security Finding Format (ASFF)
- **Multi-Framework Compliance:** Aligns with industry and regulatory frameworks including AWS Foundational Security Best Practices, CIS, PCI DSS, and NIST
- **Real-Time Monitoring:** Provides continuous monitoring and trend analysis to identify priority security issues across AWS environments

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/security-hub/
FAQ: https://aws.amazon.com/security-hub/faqs/

Technical Overview

Backup & Recovery Capabilities

Backup: AWS Security Hub CSPM does not provide traditional backup capabilities or integrate with AWS Backup service. The service focuses on security posture assessment rather than data backup. Finding data and configurations are managed internally by AWS with built-in redundancy and high availability across availability zones within each Region. Customers can export findings to Amazon S3 using custom actions with Amazon EventBridge rules for retention purposes.

Recovery: AWS Security Hub CSPM does not directly support disaster recovery scenarios or integrate with AWS Elastic Disaster Recovery. AWS manages service infrastructure with built-in resilience for high availability. The service does not contribute to traditional RPO/RTO objectives but supports business continuity by continuously monitoring security posture and identifying vulnerabilities that could impact recovery processes.

Setup Process

Onboarding: AWS Security Hub CSPM can be set up through AWS Console, API, or CLI with AWS Organizations integration or manual account setup. Essential prerequisites include allowing AWS Config with appropriate resource recording. Default security standards automatically activate unless configured otherwise.

Offboarding: Users can disable service through the AWS Management Console or API to deactivate all standards and controls. Active findings are deleted after 90 days, archived findings are deleted after 30 days. Users can export findings to Amazon S3 using Amazon EventBridge rules before disabling for retention purposes.

Data Management

- **Storage Options:** Finding data and configurations managed internally by AWS with built-in redundancy and high availability
- **Data Removal:** Active findings permanently deleted after 90 days, archived findings deleted after 30 days automatically
- **Cross-Region Replication:** Supports cross-Region aggregation with consolidated findings from multiple AWS Regions into single aggregation Region
- **Encryption:** Data encrypted in transit and at rest using AWS-managed encryption with replication across availability zones

Optimise Cost and Consumption

AWS Security Hub CSPM provides selective control disabling to reduce security check volumes and associated charges. Organisations can disable controls for global resources in non-primary Regions to avoid duplicate charges while maintaining service coverage. The integrated cost estimator tool predicts monthly expenses across security services, while volume pricing tiers reduce per-check costs as usage scales. Strategic AWS Config recording optimisation helps minimise ancillary service charges.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/securityhub/latest/userguide/what-is-securityhub.html
Getting Started: https://docs.aws.amazon.com/securityhub/latest/userguide/securityhub-settingup.html
API Reference: https://docs.aws.amazon.com/securityhub/1.0/APIReference/Welcome.html
Service Quotas: https://docs.aws.amazon.com/general/latest/gr/sechub.html

AWS Serverless Application Repository

General Information

AWS Serverless Application Repository is a managed service allowing developers to discover, deploy, and publish serverless applications in the AWS Cloud. Each application is packaged with an AWS Serverless Application Model (SAM) template that defines the AWS resources used and can be deployed with minimal configuration through AWS Lambda console integration.

Main Benefits

AWS Serverless Application Repository accelerates development by taking away the need to build common functionality from scratch through proven serverless patterns. One-click deployment verified author badges, and 5GB free Amazon Simple Storage Service (Amazon S3) storage per Region provide trusted, cost-efficient application sharing across teams and communities.

Key Use Cases

- **Application Publishing:** Share serverless applications publicly with community or privately within teams and organisations
- **Pre-built Application Deployment:** Discover and deploy applications for web backends, data processing, chatbots, and IoT
- **Nested Applications:** Create modular applications reusing common serverless components across multiple projects reducing duplication
- **Component Sharing:** Share code samples, components, and complete applications accelerating development through proven patterns

Core Features

- **Application Publishing:** Publish serverless applications using AWS Management Console, AWS SAM CLI, or AWS SDKs with templates
- **One-Click Deployment:** Browse and search applications by categories, keywords, publishers, or event sources with instant deployment
- **Access Control:** Control application visibility with private, privately shared, or public access levels including AWS Organization integration
- **Nested Applications:** Create and deploy applications containing other applications as components for code reuse and modularity
- **AWS Lambda Console Integration:** Deep integration with AWS Lambda console for serverless application management and browsing
- **Verified Author Badge:** Author verification system to establish trust and credibility for published applications in the repository
- **AWS SAM Templates:** Applications defined using AWS SAM templates for standardised deployment and configuration
- **CloudFormation Stacks:** Applications deployed as AWS CloudFormation stacks allowing unified resource management and complete application lifecycle control

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/serverless/serverlessrepo/
FAQ: https://aws.amazon.com/serverless/serverlessrepo/faqs/

Technical Overview

Backup & Recovery Capabilities

Backup: AWS Serverless Application Repository does not provide traditional backup capabilities. Published applications are replicated across multiple Regions for public availability with code packages stored in Amazon S3 with standard Amazon S3 durability. The repository maintains application metadata and versions but does not integrate with AWS Backup service.

Recovery: Applications can be quickly redeployed in different Regions to support disaster recovery scenarios through deployment capabilities and regional replication of public applications. The repository benefits from AWS's infrastructure resilience across multiple regions. RPO and RTO objectives depend on the architecture of individual deployed serverless applications rather than the repository service itself.

Setup Process

Onboarding: Access AWS Serverless Application Repository through AWS Management Console or AWS Lambda console. Browse applications, configure required parameters, acknowledge IAM permissions, and deploy with few clicks. Publishing requires AWS SAM templates, source code URLs, and README files via Console, SAM CLI, or SDKs.

Offboarding: Delete published applications using AWS Management Console or CLI to remove from repository. Delete deployed applications through standard AWS CloudFormation stack deletion processes, removing all associated resources automatically.

Data Management

- **Storage Options:** Code packages stored in Amazon S3 with standard S3 durability, 5GB free storage per Region
- **Data Removal:** Delete applications via AWS Management Console or CLI; AWS CloudFormation stack deletion removes all associated resources
- **Cross-Region Replication:** Public applications automatically replicated across multiple regions; private applications remain region-specific only
- **Encryption:** Inherits encryption capabilities from underlying AWS services like Amazon S3, AWS Lambda, and Amazon CloudFormation for deployed applications

Optimise Cost and Consumption

AWS Serverless Application Repository optimises costs through application reuse patterns that reduce development time and resource consumption. Nested applications allow component sharing across multiple projects, avoiding duplication. Organisations can share applications internally through private repositories, minimising redundant development across teams. The repository's integration with pay-per-execution serverless services provides automatic scaling and cost optimisation based on actual demand patterns.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/serverlessrepo/latest/devguide/what-is-serverlessrepo.html
Getting Started: https://docs.aws.amazon.com/serverlessrepo/latest/devguide/serverlessrepo-quick-start.html
API Reference: https://docs.aws.amazon.com/serverlessrepo/latest/devguide/resources.html
Service Quotas: https://docs.aws.amazon.com/serverlessrepo/latest/devguide/quotas.html

AWS Service Catalog

General Information

AWS Service Catalog enables organizations to create and manage catalogues of approved IT services for AWS deployment. It provides centralized governance controls while allowing end users to provision standardized cloud resources through self-service portals without requiring direct AWS permissions.

Main Benefits

AWS Service Catalog delivers reduced operational overhead through self-service provisioning, faster time-to-market for development teams, and improved compliance via pre-approved templates. Organizations achieve cost savings by preventing resource sprawl, eliminating duplicate deployments, and maintaining consistent governance across multi-account environments without complex IAM permissions.

Key Use Cases

- **Self-Service Provisioning:** Enable users to deploy approved AWS resources through personalized portals
- **Governance and Compliance:** Centrally manage organizational policies and standardized configurations for deployments
- **Multi-Account Management:** Share portfolios across multiple AWS accounts with centralized control
- **Application Registry:** Manage applications with associated configurations, resources, and metadata efficiently

Core Features

- **Product Management:** Create and manage IT services using CloudFormation templates, Terraform configurations, or AWS Marketplace offerings
- **Portfolio Organization:** Organize products into portfolios with customizable access controls and configuration information for users
- **Version Control:** Manage multiple product versions and automatically propagate updates across portfolios with centralized control
- **Constraints and Governance:** Apply launch, notification, and template constraints to control and govern resource deployment processes
- **Service Actions:** Define custom operational actions that can be performed on provisioned products for ongoing tasks
- **IAM Integration:** Fine-grained access control through AWS IAM integration without requiring complex permission management structures
- **Multi-Account Support:** Share portfolios and products across multiple AWS accounts while maintaining centralized governance and control

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/servicecatalog/
FAQ: https://aws.amazon.com/servicecatalog/faqs/

Technical Overview

Backup & Recovery Capabilities

Backup: Service Catalog leverages underlying AWS services for backup capabilities, with provisioned products inheriting backup features from deployed resources (EC2, RDS). AWS Backup policies can be implemented for Service Catalog-deployed resources, with service-specific backup features and cross-region replication available. Service metadata and configurations support regional replication.

Recovery: Recovery utilizes AWS Elastic Disaster Recovery integration and cross-region deployment capabilities through CloudFormation StackSets for multi-region provisioning. Pre-defined disaster recovery products enable automated failover processes via Service Actions. Supports customer RPO objectives (automated provisioning and standardized recovery procedures) and RTO objectives (infrastructure-as-code deployment speed minimizing recovery times).

Setup Process

Onboarding: Complete prerequisites through Getting Started Library templates or step-by-step tutorials. Create portfolios, add products using CloudFormation templates or Terraform configurations, set constraints and permissions, then grant end-user access via console, APIs, and CLI.

Offboarding: Terminate provisioned products using TerminateProvisionedProduct API, deleting associated CloudFormation stacks. Delete products, portfolios, and constraints through console or APIs with automated cleanup.

Data Management

- **Storage Options:** Service metadata and configurations stored regionally, inheriting storage from underlying AWS services
- **Data Removal:** Terminate products via TerminateProvisionedProduct API, delete portfolios and constraints through console or APIs
- **Cross-Region Replication:** Service metadata and configurations support replication across regions for disaster recovery strategies
- **Encryption:** Encryption capabilities inherited from underlying AWS services deployed through Service Catalog provisioned products

Optimize Cost and Consumption

Service Catalog prevents resource sprawl through standardized templates that enforce consistent sizing and configurations. Launch constraints enable automated rightsizing by restricting oversized deployments before provisioning occurs. Budget controls integrate directly with Service Catalog products to monitor spending per portfolio. Tagging strategies embedded in templates provide granular cost allocation across departments and projects, while centralized governance reduces administrative overhead through automated workflows and self-service provisioning capabilities.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/servicecatalog/latest/adminguide/what-is_concepts.html
Getting Started: https://docs.aws.amazon.com/servicecatalog/latest/adminguide/getstarted.html
API Reference: https://docs.aws.amazon.com/servicecatalog/latest/dg/API_Reference.html
Service Quotas: https://docs.aws.amazon.com/servicecatalog/latest/adminguide/limits.html

AWS Shield

General Information

AWS Shield is a managed DDoS protection service that safeguards web applications running on AWS. It offers Shield Standard with free automatic Layer 3/4 protection and Shield Advanced with enhanced application-layer mitigation and 24/7 expert support.

Main Benefits

AWS Shield uses AWS's massive global infrastructure for mitigation capacity while offering unique DDoS cost protection through service credits. Shield Advanced includes 24/7 expert support, integrated AWS WAF capabilities, and near-zero RTO objectives, preventing revenue loss from attacks.

Key Use Cases

- **Web Application Protection:** Protecting internet-facing applications from DDoS attacks preventing downtime and service degradation
- **E-commerce and Critical Applications:** Providing high availability for revenue-generating applications that require guaranteed uptime during attacks
- **Gaming and Media Streaming:** Protecting applications with unpredictable traffic patterns from volumetric attacks to help ensure consistent experience
- **API and Microservices Protection:** Safeguarding API gateways and distributed architectures from application layer request floods

Core Features

- **Shield Standard:** Free automatic protection against common Layer 3/4 DDoS attacks with static threshold detection
- **Shield Advanced:** Enhanced DDoS protection with application layer mitigation, health-based detection, and tailored traffic analysis
- **AWS WAF Integration:** Automatic application layer DDoS mitigation using WAF rules and rate-based controls included
- **Shield Response Team:** 24/7 access to AWS DDoS experts for attack response and custom mitigation creation
- **DDoS Cost Protection:** Service credits covering scaling charges resulting from DDoS attacks against protected resources
- **Advanced Event Visibility:** Real-time metrics, detailed attack reports, and Amazon CloudWatch integration for in-depth DDoS monitoring
- **Protection Groups:** Logical groupings of protected resources for enhanced detection, mitigation, and centralised management
- **Network Security Director:** Preview capability providing network topology visibility, security configuration analysis, and remediation recommendations

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/shield/

FAQ: https://aws.amazon.com/shield/faqs/

Technical Overview

Backup & Recovery Capabilities

Backup: AWS Shield does not provide backup capabilities as it operates as a real-time traffic analysis and mitigation service without storing customer data. Service configuration settings are maintained in AWS's control plane, but individual attack data remains transient. Customers must rely on other AWS services like AWS Backup for data protection needs.

Recovery: AWS Shield provides immediate automatic DDoS mitigation without manual intervention, supporting near-zero RTO objectives through real-time protection. The service maintains high availability across AWS's global infrastructure with automatic failover capabilities. Shield's real-time nature helps ensure no data loss during DDoS attacks, though customers require separate backup strategies for data protection.

Setup Process

Onboarding: AWS Shield Standard is automatically initiated for all AWS customers at no cost with no configuration set up required. AWS Shield Advanced requires subscription through AWS Management Console or API, adding resource protections, configuring health-based detection with Amazon Route 53 health checks, and setting up Amazon CloudWatch alarms. Supports AWS CloudFormation and Terraform deployment.

Offboarding: AWS Shield Standard cannot be disabled as it's automatically included. AWS Shield Advanced subscriptions can be cancelled through AWS Management Console after fulfilling one-year commitment. All protections are removed, custom mitigations deleted, and AWS Shield Response Team (SRT) access is terminated automatically.

Data Management

- **Storage Options:** No customer data storage; real-time traffic analysis only with transient attack data
- **Data Removal:** No customer data stored by AWS Shield; Amazon CloudWatch event data follows standard retention policies
- **Cross-Region Replication:** AWS Shield protections should be configured for both primary and secondary Regions for coverage
- **Encryption:** The service operates at network perimeter with AWS infrastructure security; no stored data encryption

Optimise Cost and Consumption

AWS Shield Advanced provides DDoS cost protection through service credits covering scaling charges during attacks, preventing unexpected bills. Consolidated billing allows one Shield Advanced subscription across all AWS Organizations accounts. Fee waivers allow single subscriptions spanning multiple organisations. Automated mitigation reduces operational overhead and manual intervention costs. Protection Groups allow centralised resource management for enhanced efficiency across distributed architectures.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/waf/latest/developerguide/shield-chapter.html

Getting Started: https://docs.aws.amazon.com/waf/latest/developerguide/getting-started-ddos.html

API Reference: https://docs.aws.amazon.com/waf/latest/DDOSAPIReference/Welcome.html

Service Quotas: https://docs.aws.amazon.com/general/latest/gr/shield.html

AWS Step Functions

General Information

AWS Step Functions is a serverless workflow orchestration service that coordinates multiple AWS services into complex application workflows using Amazon States Language. It offers Standard workflows for long-running processes and express workflows for high-volume workloads with built-in error handling.

Main Benefits

Step Functions removes infrastructure management overhead while integrating with more than 220 AWS services and more than 10,000 APIs through its serverless architecture. Express workflows deliver up to 90% cost reduction for high-volume scenarios, while the visual Workflow Studio accelerates development and built-in error handling handles reliable execution.

Key Use Cases

- **Microservice Orchestration:** Combining AWS Lambda functions to build web-based applications with human approval steps
- **Security and IT Automation:** Orchestrating security incident response workflows and automating operational event responses
- **Data Processing and ETL:** Building data processing pipelines with parallel job execution for streaming data
- **Machine Learning Operations:** Running ETL jobs and building, training, deploying ML models using Amazon SageMaker integration

Core Features

- **Visual Workflow Studio:** Drag-and-drop interface with design, code, and config modes for creating workflows visually
- **Two Workflow Types:** Standard workflows for long-running processes and Express workflows for high-volume short-running workloads
- **Service Integration Patterns:** Three patterns supporting more than 220 AWS services: Request Response, Run Job, and Wait for Callback
- **Built-in Error Handling:** Automatic retry mechanisms, dead-letter queues, error notifications, and exception handling with catch states
- **State Management:** Maintains workflow state with execution history, input/output tracking, and resumption after interruptions
- **Parallel Processing:** Map state in distributed mode for large-scale processing and parallel state for concurrent branches
- **Monitoring and Observability:** Integration with Amazon CloudWatch for logging, X-Ray for tracing, and visual execution history tracking
- **Amazon States Language:** JSON-based language for defining workflows as event-driven steps with flow logic and state transitions

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/step-functions/
FAQ: https://aws.amazon.com/step-functions/faqs/

Technical Overview

Backup & Recovery Capabilities

Backup: State machine definitions are automatically backed up within the AWS control plane infrastructure with built-in durability. Customers can export definitions as JSON/YAML manually through the AWS Management Console, AWS CLI, or programmatically via APIs for automated backup processes. Execution history is preserved through Amazon CloudWatch Logs retention policies when logging is selected.

Recovery: Workflows resume from the last successful state after interruptions, with cross-Region redeployment for disaster recovery. Multi-Availability Zone deployment provides automatic failover within Regions, while JSON-based definitions allow recovery by redeploying in alternate Regions. Supports aggressive RPO (real-time state information) and RTO (automatic failover capabilities) objectives through serverless architecture resilience.

Setup Process

Onboarding: Start through AWS Management Console's Hello World tutorial (20-30 minutes) using Workflow Studio's drag-and-drop interface in AWS Step Functions. Configure identity and access management (IAM) roles, policies, logging, tracing, and security settings through design, code, and config modes with visual or Amazon States Language editing.

Offboarding: Delete state machines via AWS Management Console, AWS CLI, or API to remove definitions and stop executions. Remove associated AWS IAM roles, Amazon CloudWatch log groups, AWS Lambda functions, and workflow-specific resources to back complete cleanup and prevent residual charges.

Data Management

- **Storage Options:** State machine definitions stored in AWS control plane with execution history of up to 25,000 entries
- **Data Removal:** Delete state machines via AWS Management Console, AWS CLI, or API removing definitions and stopping new executions
- **Cross-Region Replication:** Export state machine definitions as JSON/YAML and redeploy in alternate Regions for disaster recovery
- **Encryption:** Enterprise-grade security features with integration to IAM for permission management and access control

Optimise Cost and Consumption

Express workflows deliver up to 90% cost reduction for high-volume scenarios through optimised billing structures. Workflow type selection allows nesting express workflows within standard workflows for mixed workload optimisation. Efficient state machine design minimises state transitions, the primary billing unit. Amazon S3 ARNs reduce payload costs while avoiding size limits. Map state in distributed mode provides cost-effective parallel processing. Proper timeout configuration prevents runaway executions and unnecessary charges.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/step-functions/latest/dg/welcome.html
Getting Started: https://docs.aws.amazon.com/step-functions/latest/dg/getting-started.html
API Reference: https://docs.aws.amazon.com/step-functions/latest/apireference/index.html
Service Quotas: https://docs.aws.amazon.com/general/latest/gr/step-functions.html

AWS Storage Gateway

General Information

AWS Storage Gateway is a hybrid cloud storage service that connects on-premises environments with AWS storage infrastructure using standard protocols like iSCSI, SMB, and NFS. It provides file-based, volume-based, and tape-based storage solutions with intelligent local caching for low-latency access.

Main Benefits

Supports hybrid cloud storage without application rewrites, delivering up to 98% cost savings through intelligent data tiering and Amazon Simple Storage Service (Amazon S3) storage classes. Provides low-latency access via local caching, 99.999999999% data durability, and supports RPO objectives with hourly incremental snapshots for backup and disaster recovery.

Key Use Cases

- **Backup and Disaster Recovery:** Move backups to cloud while maintaining on-premises access for recovery
- **Hybrid Cloud File Shares:** Reduce on-premises storage footprint with cloud-backed SMB and NFS file shares
- **Data Lake Ingestion:** Store file data as Amazon S3 objects for analytics and machine learning workflows
- **Virtual Tape Library Replacement:** Replace physical tape infrastructure with virtual tapes in Amazon S3 Glacier

Core Features

- **S3 File Gateway:** Supports file storage as Amazon S3 objects that are accessible via NFS and SMB protocols with local caching
- **Volume Gateway:** Provides cached and stored volume options with Amazon S3 storage and local data access capabilities
- **Tape Gateway:** Virtual tape library interface for backup applications with cloud-based virtual tape storage in Amazon S3
- **FSx File Gateway:** Delivers low-latency access to Amazon FSx for Windows File Server file shares on-premises
- **Local Caching:** Intelligent caching of frequently accessed data on-premises for optimised low-latency performance and access
- **Data Optimisation:** Compresses data and transfers only changed data to optimise bandwidth usage and reduce costs
- **Multiple Deployment Options:** Deploy as VM appliance, hardware appliance, or Amazon Elastic Compute Cloud (Amazon EC2) instance across multiple hypervisors
- **Built-in AWS Integration:** Integrates with Amazon S3, Amazon CloudWatch, AWS Identity and Access Management (AWS IAM), AWS Key Management Service (AWS KMS), and AWS Backup for cloud storage management

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/storagegateway/

FAQ: https://aws.amazon.com/storagegateway/faqs/

Technical Overview

Backup & Recovery Capabilities

Backup: Point-in-time EBS snapshots and virtual tape archives stored in Amazon S3 with 99.999999999% durability can be created manually through the console or automated via AWS Backup with hourly scheduling frequencies. Incremental backups reduce storage costs and support cross-Region replication for disaster recovery.

Recovery: Fast recovery from Amazon S3 snapshots, failover to alternate gateway instances, and cloud-based recovery to Amazon Elastic Compute Cloud (Amazon EC2) instances. Cross-region snapshot replication provides disaster recovery coverage. Supports customer RPO objectives through hourly incremental snapshots and RTO objectives through local caching and snapshot restoration.

Setup Process

Onboarding: Launch through AWS Management Console or SDKs by choosing gateway type (File, Volume, or Tape) and host environment (VM, hardware appliance, or EC2). Configure gateway activation, allocate local cache storage, and create cloud storage resources using setup wizards.

Offboarding: Delete gateway-specific resources (volumes, file shares, virtual tapes) then remove gateway through the AWS Management Console. Manual cleanup required for associated AWS resources like Amazon EBS snapshots and Amazon S3 objects to avoid ongoing charges.

Data Management

- **Storage Options:** File-based, volume-based, and tape-based storage using Amazon S3's Glacier, and Deep Archive storage classes
- **Data Removal:** Delete gateway resources then manually remove associated Amazon EBS snapshots and Amazon S3 objects
- **Cross-Region Replication:** Supports cross-Region snapshot replication for disaster recovery and backup scenarios
- **Encryption:** Use AWS Key Management Service (AWS KMS) for encryption at rest with enterprise-grade security capabilities

Optimise Cost and Consumption

Intelligent data tiering automatically moves infrequently accessed data to lower-cost Amazon S3 storage classes like Glacier and Deep Archive. Local caching optimises bandwidth usage by storing frequently accessed data on-premises. Data compression and incremental synchronisation transfer only changed data, reducing bandwidth charges. Right-sizing local cache storage based on access patterns optimises performance and costs.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/storagegateway/

Getting Started: https://docs.aws.amazon.com/storagegateway/latest/vgw/GettingStarted.html

API Reference: https://docs.aws.amazon.com/storagegateway/latest/APIReference/Welcome.html

Service Quotas: https://docs.aws.amazon.com/general/latest/gr/sg.html

AWS Supply Chain

General Information

AWS Supply Chain is a cloud-based supply chain management application that integrates with existing ERP systems to provide unified data visibility. It uses machine learning and generative AI to enhance demand forecasting, supply planning, and risk management across end-to-end supply chains.

Main Benefits

AWS Supply Chain removes overheads such as replatforming requirements and upfront licensing fees and provides simplified ERP integration. It delivers reduced time-to-market, improved operational efficiency, and automated risk detection with machine learning insights. The service offers pay-as-you-use pricing and global scalability through AWS infrastructure.

Key Use Cases

- **Demand Planning:** Create machine learning-driven forecasts and aid cross-team collaboration for market conditions.
- **Supply Planning:** Automate replenishment and manufacturing plans for raw materials, components, and finished goods.
- **Risk Management:** Generate automated insights for stockouts, excess inventory, and lead time deviations.
- **N-Tier Visibility:** Extend supply chain transparency and insights to external trading partners.

Core Features

- **Data Lake:** Aggregates supply chain data using extensible models supporting Amazon Simple Storage Service (Amazon S3), Electronic Data Interchange (EDI), SAP S/4HANA, and SAP ECC 6.0.
- **Machine Learning Insights:** Automatically generates supply chain risk insights and rebalance recommendations using advanced algorithms.
- **Order Planning and Tracking:** Provides work order status, arrival predictions, delivery risk assessments, and actionable recommendations.
- **Demand Planning:** Creates machine learning forecasts, adjusts for market conditions, and supports cross-team collaboration.
- **Supply Planning:** Supports auto replenishment and manufacturing plans for raw materials, components, and finished goods.
- **N-Tier Visibility:** Extends supply chain visibility and insights beyond organisation boundaries to external trading partners.
- **Sustainability Management:** Supports partner collaboration and sustainability information mapping for environmental tracking.
- **Amazon Quick Integration:** Provides generative AI assistant for natural language querying and data analysis capabilities.

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://docs.aws.amazon.com/aws-supply-chain/latest/adminguide.html

FAQ: https://docs.aws.amazon.com/aws-supply-chain/latest/adminguide/faq.html

Technical Overview

Backup & Recovery Capabilities

Backup: AWS Supply Chain relies on underlying AWS service backup capabilities with Amazon S3 durability and versioning. Centralised backup policies can be created via AWS Backup for Amazon DynamoDB and Amazon S3 supporting Supply Chain operations.

Recovery: Cross-region replication supports disaster recovery through AWS CloudFormation deployment and multi-Region architecture. AWS Elastic Disaster Recovery provides continuous replication capabilities. Supports customer RPO objectives through Amazon S3 cross-Region replication and Amazon DynamoDB global tables, and RTO objectives via multi-Region deployments with automated failover mechanisms.

Setup Process

Onboarding: Setup involves four steps via AWS Management Console: 1) Assign IAM Identity Center user profile 2) Create Supply Chain instance 3) Choose application owner 4) Login to web application. The service supports up to 10 instances per account with test drive feature and role-based access configuration.

Offboarding: Instance deletion uses programmatic APIs including DeleteInstance and DeleteDataLakeDataset for asynchronous cleanup of AWS resources and AWS Key Management Service (AWS KMS) keys. Amazon S3 data requires manual cleanup as it's not automatically deleted during termination.

Data Management

- **Storage Options:** Amazon S3, Amazon DynamoDB, and data lake storage supporting EDI, SAP S/4HANA, and SAP ECC 6.0 formats.
- **Data Removal:** Programmatic APIs including DeleteInstance, DeleteDataLakeDataset, and DeleteDataIntegrationFlow with asynchronous cleanup operations.
- **Cross-Region Replication:** Supported through AWS services like Amazon S3's cross-Region replication and Amazon DynamoDB global tables capabilities.
- **Encryption:** AWS KMS keys automatically managed during instance creation and deletion for data protection.

Optimise Cost and Consumption

AWS Supply Chain provides tiered pricing that reduces per-unit costs at high volumes, with Supply Planning combinations decreasing from higher to lower rates across the first million combinations. N-Tier Visibility capabilities are bundled within Supply Planning and Insights pricing without additional charges. Automated machine learning insights identify inventory inefficiencies and optimise levels, potentially reducing working capital costs through intelligent rebalancing recommendations and demand forecasting accuracy improvements.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/aws-supply-chain/latest/userguide/index.html

Getting Started: https://docs.aws.amazon.com/aws-supply-chain/latest/adminguide/getting-started.html

API Reference: https://docs.aws.amazon.com/aws-supply-chain/latest/APIReference/index.html

Service Quotas: https://docs.aws.amazon.com/aws-supply-chain/latest/adminguide/quotas.html

AWS Systems Manager

General Information

AWS Systems Manager is a management solution that provides centralised operations for compute nodes across AWS, on-premises, and multicloud environments. It offers secure end-to-end infrastructure management through a unified console experience with automated operational capabilities.

Main Benefits

AWS Systems Manager manages thousands of nodes from a single console across hybrid and multicloud environments, removing the need to access bastion hosts and SSH. It reduces the attack surface as inbound ports are not required, offers free tier capabilities, and delivers cost savings through automation and reduced operational overhead.

Key Use Cases

- **Node Management:** Centrally manage Amazon Elastic Compute Cloud (Amazon EC2) instances, on-premises servers, and multicloud environments through a unified console
- **Secure Access and Automation:** Allow secure shell access without SSH/RDP and automate administrative tasks remotely
- **Configuration Management:** Store and manage application configuration data, secrets, and parameters centrally at scale
- **Compliance Monitoring:** Monitor and report on configuration compliance across managed nodes with automated tracking

Core Features

- **Run Command:** Execute commands remotely on managed nodes without SSH access, with rate control and targeting capabilities
- **Session Manager:** Secure browser-based shell access to managed nodes without opening inbound ports or managing bastion hosts
- **Patch Manager:** Automate the patching process for managed nodes with security updates and patches using patch baselines
- **Parameter Store:** Secure hierarchical storage for configuration data, secrets, and parameters with encryption and access control
- **Automation:** Automate common administrative and operational tasks using predefined or custom runbooks for streamlined operations
- **State Manager:** Maintain consistent configuration across managed nodes by defining and enforcing desired state configurations automatically
- **Inventory:** Collect metadata and software inventory information from managed nodes for compliance and asset management purposes
- **OpsCenter:** Central location for operations engineers to view, investigate, and resolve operational work items efficiently

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/systems-manager/
FAQ: https://aws.amazon.com/systems-manager/faq/

Technical Overview

Backup & Recovery Capabilities

Backup: AWS Systems Manager provides limited direct backup capabilities but integrates with AWS Backup for data protection. Parameter Store supports parameter versioning with durable storage in Amazon Simple Storage Service (Amazon S3) and Amazon CloudWatch. Automation documents and runbooks can be exported manually using the AWS Systems Manager Console or APIs for backup purposes.

Recovery: Configuration recovery through cross-Region parameter sharing and Infrastructure as Code practices aids deployment across Regions. Automation runbooks orchestrate disaster recovery procedures, reducing manual intervention. Supports customer RPO objectives through cross-region replication and RTO objectives through automated provisioning processes.

Setup Process

Onboarding: Setup requires SSM Agent installation and Identity and Access Management (IAM) role configuration via the unified AWS Systems Manager console. Default Host Management Configuration simplifies deployment automatically. Setup options configure commonly used features with proven practices and one-click runbooks for connectivity remediation.

Offboarding: Data removal involves deleting Parameter Store data, automation runbooks, and inventory information through AWS Systems Manager console or APIs. Remove SSM Agent registration and clean local directories for complete node deregistration.

Data Management

- **Storage Options:** Parameter Store provides hierarchical storage with durable backing in Amazon S3 and AWS CloudWatch services
- **Data Removal:** Delete parameters, runbooks, inventory data, and compliance information through AWS Systems Manager console or APIs with cleanup automation
- **Cross-Region Replication:** Parameter Store supports cross-region parameter sharing and automation runbooks deployable across multiple regions
- **Encryption:** AWS Key Management Service integration provides encryption for sensitive parameters with comprehensive access control capabilities

Optimise Cost and Consumption

The AWS Systems Manager's adaptive concurrency feature automatically scales automation executions up to 500 concurrent operations, optimising resource usage during peak workloads. Run Command provides rate control and targeting capabilities to throttle command execution across managed nodes. Automation runbooks reduce manual labour costs through predefined workflows, while Parameter Store versioning minimises storage overhead by efficiently managing configuration data history and reducing duplicate parameter storage requirements.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/systems-manager/

Getting Started: https://docs.aws.amazon.com/systems-manager/latest/userguide/what-is-systems-manager.html

API Reference: https://docs.aws.amazon.com/systems-manager/latest/APIReference/

Service Quotas: https://docs.aws.amazon.com/general/latest/gr/ssm.html

AWS Transfer Family

General Information

AWS Transfer Family is a fully managed file transfer service for secure file transfers—into and out of—AWS storage services. It supports multiple protocols including SFTP, FTPS, FTP, AS2, and web-based transfers with automatic scaling and high availability.

Main Benefits

AWS Transfer Family takes away the operational burden of managing file transfer infrastructure while allowing migration of existing file transfer workflows. Existing applications and workflows can continue using the same protocols without any modifications. Features 24/7 availability with multi-availability zone (multi-AZ) redundancy, automatic scaling, pay-as-you-use pricing without any upfront costs, and fast deployment. Supports up to 5TB individual file transfers.

Key Use Cases

- **Data Lakes:** Ingesting external data into Amazon Simple Storage Service (Amazon S3) for analytics and machine learning
- **B2B File Exchanges:** Secure business-to-business document exchange using AS2 protocol with trading partners
- **Compliance Workflows:** Helps to meet regulatory requirements for secure file transfers with in-depth audit trails
- **Supply Chain Management:** Automated file processing and distribution across multiple business units

Core Features

- **Multi-Protocol Support:** Supports SFTP, FTPS, FTP, AS2, and web browser-based file transfers
- **Managed Workflows:** Automates file processing tasks like decryption, tagging, and AWS Lambda function execution
- **Multiple Authentication Methods:** Service managed, AWS Directory Service, or custom identity providers via AWS Lambda/Amazon API Gateway
- **Storage Integration:** Built-in integration with Amazon S3 for object storage and Amazon Elastic File Service (Amazon EFS) file systems
- **High Availability:** Multi-AZ deployment with automatic failover and elastic scaling capabilities
- **SFTP Connectors:** Bidirectional file transfers between AWS and remote SFTP servers without intermediate storage
- **Web Applications:** Browser-based file transfer interface for end users with secure session management
- **Security Features:** Encryption in transit and at rest, Pretty Good Privacy (PGP) encryption, and malware protection

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/aws-transfer-family/
FAQ: https://aws.amazon.com/aws-transfer-family/faqs/

Technical Overview

Backup & Recovery Capabilities

Backup: AWS Transfer Family uses underlying storage services with Amazon S3 versioning, cross-region replication, and AWS Backup integration for automated scheduling. Amazon EFS provides automatic backups and point-in-time recovery. Server configurations can be exported through AWS CloudFormation or Terraform templates for version control.

Recovery: Multi-AZ deployment allows automated failover with minimal downtime, cross-region disaster recovery through Infrastructure as Code for server recreation, and AWS Application Recovery Controller integration. Supports customer RPO objectives through real-time data replication with near-zero RPO and RTO objectives via elastic infrastructure, taking away capacity constraints.

Setup Process

Onboarding: Create servers through AWS Management Console or APIs after configuring Amazon S3/Amazon EFS storage and AWS Identity and Access Management (IAM) roles. Select protocols, authentication methods, and user directories. Optional managed workflows that allow automation. No infrastructure provisioning required, aiding fast deployment.

Offboarding: Delete servers, users, and resources via the AWS Management Console or APIs with immediate endpoint decommissioning. Remove workflows, connectors, and certificates separately. Delete IAM roles and custom identity configurations for a complete cleanup.

Data Management

- **Storage Options:** Built-in integration with Amazon S3 for object storage and Amazon EFS for file systems
- **Data Removal:** Files in Amazon S3/Amazon EFS remain until explicitly deleted; servers decommissioned immediately upon deletion via AWS Management Console/APIs
- **Cross-Region Replication:** Supported through underlying Amazon S3 cross-region replication and Amazon EFS backup capabilities for disaster recovery
- **Encryption:** Encryption in transit and at rest, PGP encryption, and AWS Key Management Service integration for key management

Optimise Cost and Consumption

Protocol-based pricing allows selecting only required protocols to minimise hourly charges. SFTP connectors reduce data transfer costs through direct server-to-server transfers bypassing intermediate storage. Managed workflows automate file processing without custom infrastructure. Multi-protocol support on single servers reduces endpoint proliferation. Data compression and S3 Intelligent Tiering integration optimise transfer patterns and long-term storage costs.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/transfer/latest/userguide/what-is-aws-transfer-family.html
Getting Started: https://docs.aws.amazon.com/transfer/latest/userguide/tutorials-transfer.html
API Reference: https://docs.aws.amazon.com/transfer/latest/APIReference/api-welcome.html
Service Quotas: https://docs.aws.amazon.com/general/latest/gr/transfer-service.html

AWS Transfer Family – AWS Transfer for SFTP

Feature of AWS Transfer Family

General Information

AWS Transfer for Secure Shell [SSH] File Transfer Protocol (SFTP) is a fully managed file transfer service, part of the AWS Transfer Family, that supports secure file transfers using SFTP protocol into and out of AWS storage services. The service supports transferring data to Amazon Simple Storage Service (Amazon S3) and Amazon Elastic File Service (Amazon EFS) without requiring server infrastructure management.

Main Benefits

Takes away the infrastructure management overhead of SFTP servers and maintains compatibility with existing client applications and workflows. Provides automatic scaling, high availability across multiple availability zones, and 99.999999999% data durability. AWS Transfer Family is compliant with the GDPR and uses the pay-as-you-go pricing model.

Key Use Cases

- **B2B File Exchanges:** Securely exchange files with trading partners using standardised SFTP protocol
- **Data Lake Ingestion:** Upload files from vendors and partners to AWS data lakes for analytics
- **Internal Data Transfers:** Move files between departments using familiar SFTP workflows within organisations
- **Compliance Workflows:** Help meet regulatory requirements for secure transfers in financial services and healthcare

Core Features

- **SFTP Protocol Support:** Full support for SFTP version 3 protocol with SSH authentication and encryption
- **Multiple Storage Backends:** Built-in integration with Amazon S3 and Amazon EFS for data storage
- **Identity Provider Options:** Support for service-managed, AWS Directory Service, custom AWS Lambda, and API Gateway authentication
- **SFTP Connectors:** Automated file transfers between Amazon S3 and remote SFTP servers with bidirectional capabilities
- **Virtual Private Connection (VPC) Endpoint Support:** Private connectivity through Amazon VPC endpoints for secure, internal network access
- **Managed Workflows:** Post-upload file processing including encryption, decryption, tagging, and custom Lambda functions
- **Logical Directories:** Simplified directory structures mapped to Amazon S3 bucket paths
- **Multi-Protocol Support:** Single server can support multiple protocols including SFTP, FTPS, FTP, and AS2

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/transfer-family/

FAQ: https://aws.amazon.com/transfer-family/faqs/

Technical Overview

Backup & Recovery Capabilities

Backup: AWS Transfer for SFTP uses the underlying storage service capabilities with Amazon S3 providing 99.999999999% durability and cross-region replication, while Amazon EFS offers automatic daily backups and point-in-time recovery. Backups can be automated via AWS Backup integration with configurable scheduling frequencies for centralised management across services.

Recovery: Multi-Availability Zone architecture automatically recovers from individual zone failures within minutes, cross-region disaster recovery through Amazon S3's Cross-Region Replication and multiple AWS Transfer Family servers. Supports customer RPO objectives (near-zero with S3 cross-region replication) and RTO objectives (minutes for zone-level failures).

Setup Process

Onboarding: Create SFTP-enabled server through AWS Transfer Family console or APIs/CLI with step-by-step wizards. Configure identity provider options, endpoint settings, Amazon S3/Amazon EFS storage, IAM roles, and SSH keys for user authentication. Test using standard SFTP clients.

Offboarding: Stop server to prevent new connections, backup required data from storage locations, delete user accounts and SSH keys, remove IAM roles and policies, delete AWS Transfer Family server and associated resources.

Data Management

- **Storage Options:** Built-in integration with Amazon S3 for object storage and Amazon EFS for file system storage
- **Data Removal:** Data stored in Amazon S3 or Amazon EFS must be explicitly deleted as it persists independently
- **Cross-Region Replication:** Supported through Amazon S3 Cross-Region Replication and EFS replication for comprehensive disaster recovery strategies
- **Encryption:** Built-in encryption capabilities with AWS Key Management Service integration for encryption key management and secure file transfers

Optimise Cost and Consumption

VPC endpoints avoid data transfer charges for internal traffic while SFTP connectors automate file transfers to reduce manual operations. Rightsizing server configurations based on usage patterns optimises server uptime costs. Amazon S3 storage classes support long-term cost optimisation of transferred data. Automated workflows reduce manual processing overhead and pay-per-use pricing aligns costs with actual consumption patterns.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/transfer/latest/userguide/what-is-aws-transfer-family.html

Getting Started: https://docs.aws.amazon.com/transfer/latest/userguide/getting-started.html

API Reference: https://docs.aws.amazon.com/transfer/latest/APIReference/

Service Quotas: https://docs.aws.amazon.com/general/latest/gr/transfer-service.html

AWS Transform

General Information

AWS Transform accelerates enterprise workload modernisation through agentic AI automation. It streamlines migration of mainframes, VMware workloads, and.NET applications to AWS cloud infrastructure. The service automates complex transformation tasks while preserving business logic.

Main Benefits

AWS Transform accelerates enterprise modernisation up to 4x faster than traditional methods through agentic AI automation. The service delivers up to 40% licensing cost savings for.NET modernisation, automates network conversion 80x faster for VMware migration, and transforms mainframe applications from months to minutes.

Key Use Cases

- **Mainframe Modernisation:** Migrate IBM z/OS and Fujitsu GS21 applications with automated COBOL to Java conversion
- **VMware Migration:** Migrate VMware workloads to Amazon Elastic Compute Cloud (Amazon EC2) with automated discovery and network conversion
- **.NET Application Modernisation:** Modernise Windows.NET applications to Linux-ready cross-platform.NET with framework upgrades
- **Custom Transformations:** Large-scale code modernisation including API migrations, language upgrades, and framework refactoring

Core Features

- **Agentic AI Experience:** Uses specialised AI agents with expertise in languages, frameworks, and infrastructure to automate complex transformations
- **Conversational Interface:** Natural language processing interface for describing current architecture and transformation goals in everyday language
- **Mainframe Modernisation:** Automated code analysis, business logic extraction, and COBOL to Java refactoring for IBM z/OS applications
- **VMware Migration:** AI-driven network architecture conversion, migration planning, and automated server rehosting to Amazon EC2
- **.NET Modernisation:** Automated porting from.NET Framework to cross-platform.NET with Linux compatibility and package modernisation
- **Wave Planning:** Intelligent dependency analysis and migration wave planning using graph neural networks for optimised sequences
- **Human-in-the-Loop Collaboration:** Unified web experience for cross-functional teams with collaboration features and approval workflows
- **Infrastructure as Code Generation:** Automated generation of AWS CloudFormation templates and deployment configurations for target environments
- AWS Application Migration Service (AWS MGN): Automates lift-and-shift migrations by converting source servers to run natively on AWS with minimal downtime.

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/transform/

FAQ: https://aws.amazon.com/transform/faq/

Technical Overview

Backup & Recovery Capabilities

Backup: AWS Transform does not provide dedicated backup capabilities as a core service feature. Transformation artifacts and job history are stored using underlying AWS infrastructure with standard durability guarantees in Amazon Simple Storage Service (Amazon S3). Customers are responsible for backing up source code repositories, transformation results, and infrastructure templates using their own backup policies.

Recovery: AWS Transform does not offer built-in disaster recovery capabilities or integration with AWS Elastic Disaster Recovery. The service operates within standard AWS regional infrastructure and inherits underlying AWS service resilience. Customers must implement their own disaster recovery strategies using standard AWS patterns. The service does not directly provide RPO and RTO capabilities for transformation workflows.

Setup Process

Onboarding: Initiate AWS Transform in the AWS Management Console, configure IAM Identity Center or third-party identity providers, and set encryption preferences. Create workspaces, configure connectors for external systems, establish account connections, and define user roles with appropriate permissions.

Offboarding: Delete workspaces, connectors, and jobs through service console or APIs. Remove transformation artifacts from Amazon S3 buckets, clean up AWS CloudFormation stacks, and revoke external account connections for complete data removal.

Data Management

- **Storage Options:** Transformation artifacts and code repositories stored in Amazon S3 with built-in durability and versioning
- **Data Removal:** Users can delete workspaces, connectors, and jobs through service console or APIs with programmatic deletion capabilities
- **Cross-Region Replication:** No built-in cross-region replication; operates within standard AWS Regional infrastructure with inherited resilience
- **Encryption:** Uses AWS Key Management Service for encryption of data, conversion history, and transformation artifacts

Optimise Cost and Consumption

AWS Transform uses graph neural networks for wave planning to optimise transformation sequences and reduce resource consumption. It provides automated EC2 instance optimisation for VMware migrations to right-size infrastructure. Assessment and transformation are zero-cost, while automated modernisation reduces consultant needs by accelerating timelines up to 4x faster than traditional methods.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/transform/latest/userguide/what-is-service.html

Getting Started: https://docs.aws.amazon.com/transform/latest/userguide/getting-started.html

API Reference:
https://docs.aws.amazon.com/transform/latest/userguide/security_iam_permissions.html

Service Quotas: https://docs.aws.amazon.com/transform/latest/userguide/transform-limits.html

AWS Transit Gateway

General Information

AWS Transit Gateway is a fully managed network transit hub that connects multiple virtual private connections (VPCs) and on-premises networks through a central gateway using hub-and-spoke architecture. It simplifies VPC peering and supports centralised connections routing and traffic inspection capabilities.

Main Benefits

AWS Transit Gateway dramatically reduces operational complexity by replacing hundreds of VPC peering connections with a single managed hub, supporting thousands of connections with up to 100 Gbps bandwidth per attachment. It provides enterprise-grade scalability, built-in encryption controls, and flexible cost allocation, removing networking operational overhead.

Key Use Cases

- **Centralised Connectivity:** Connect thousands of VPCs and on-premises networks through single hub
- **Multi-Account Architecture:** Allow secure communication between VPCs across multiple AWS accounts centrally
- **Hybrid Cloud Integration:** Integrate on-premises networks using VPNs, direct connections, and SD-WAN solutions
- **Network Segmentation:** Implement isolated routing domains with centralised security appliances and inspection

Core Features

- **Hub-and-Spoke Architecture:** Central gateway connecting thousands of VPCs and on-premises networks with simplified routing
- **Multiple Attachment Types:** Support for VPC, VPN, direct connect, transit gateway connect, peering, and network functions
- **Route Table Management:** Default and custom route tables with dynamic routing, route propagation, and traffic segmentation
- **High Bandwidth Performance:** Up to 100 Gbps per VPC attachment per AZ and 20 Gbps per Connect attachment
- **Cross-Region Peering:** Inter-region connectivity keeping traffic on AWS backbone with automatic encryption at physical layer
- **Multicast Support:** Integrated multicast routing with IGMP snooping for broadcasting applications like financial data feeds
- **Encryption Control:** Enforce encryption-in-transit for all traffic on attached VPCs with configurable encryption policies
- **Resource Sharing:** AWS RAM integration for sharing transit gateways across multiple accounts with granular permissions

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/transit-gateway/

FAQ: https://aws.amazon.com/transit-gateway/faqs/

Technical Overview

Backup & Recovery Capabilities

Backup: Transit Gateway configurations use Infrastructure as Code templates through AWS CloudFormation, AWS Cloud Development Kit (CDK), or Terraform for version-controlled backup. Route tables and attachment configurations are exported to templates rather than traditional snapshots. Flow logs and monitoring data can be archived to Amazon Simple Storage Service (Amazon S3) for compliance requirements.

Recovery: Network configurations restore through IaC deployment within 5-15 minutes, with cross-region peering allowing automatic traffic rerouting during failures. Multi-Availability Zone architecture provides sub-second failover. Supports customer RPO of near-zero for connectivity and RTO objectives of minutes for DNS-based failover scenarios.

Setup Process

Onboarding: Create Transit Gateway through AWS Management Console, CLI, or API, specifying ASN and allowing DNS support and multicast features. Create VPC attachments by selecting subnets, then configure route tables for traffic flow control and establish hybrid connectivity through VPN or Direct Connect attachments.

Offboarding: Systematic detachment requires deleting route table associations first, then detaching all VPCs, VPN connections, and peering connections until reaching 'deleted' state. Finally delete the transit gateway, to automatically remove all configurations and metadata permanently.

Data Management

- **Storage Options:** Infrastructure-as-Code templates store configurations; flow logs and monitoring data archived to Amazon S3
- **Data Removal:** Systematic detachment of attachments required before deletion; all configurations and metadata permanently removed
- **Cross-Region Replication:** Multi-region peering supports cross-region connectivity with automatic failover through routing updates
- **Encryption:** Enforces encryption-in-transit for VPC traffic; cross-region peering includes automatic encryption at physical layer

Optimise Cost and Consumption

Flexible Cost Allocation (FCA) supports custom allocation of data processing charges to source, destination, or central accounts based on defined policies. VPN Concentrator attachments consolidate up to 100 connections under single attachments for low-bandwidth sites. Transit Gateway Connect attachments reduce data processing charges through strategic deployment. Cost allocation tags provide granular charge tracking across teams and projects while careful route table design minimises unnecessary traffic traversal.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/vpc/latest/tgw/

Getting Started: https://docs.aws.amazon.com/vpc/latest/tgw/tgw-getting-started.html

API Reference: https://docs.aws.amazon.com/AWSEC2/latest/APIReference/OperationList-query-ec2.html

Service Quotas: https://docs.aws.amazon.com/vpc/latest/tgw/transit-gateway-quotas.html

AWS Trusted Advisor

General Information

AWS Trusted Advisor is a proactive guidance service that inspects AWS environments and provides real-time recommendations to optimise costs, improve performance, and enhance security. It continuously analyses resource configurations and usage patterns, delivering actionable insights based on AWS's proven practices.

Main Benefits

AWS Trusted Advisor delivers automated, proactive guidance based on proven practices from hundreds of thousands of AWS customers. It provides over 482 checks across six critical areas—cost optimisation, security, performance, resilience, operational excellence, and service limits—with more than 16 specialised cost optimisation checks offering accurate savings estimates and continuous security monitoring to help optimise your AWS environment.

Key Use Cases

- **Cost Optimisation:** Identify less used resources, rightsizing recommendations, and reserved instance opportunities to reduce spending
- **Security Enhancement:** Automated security gap detection, access permission analysis, and compliance monitoring across resources
- **Performance Optimisation:** Analyse resource configurations, identify bottlenecks, and provide recommendations for improved availability
- **Service Limits Monitoring:** Real-time tracking of AWS service quotas with alerts to prevent disruptions

Core Features

- **Cost Optimisation Checks:** Over 16 integrated checks providing customised recommendations for Amazon EC2, Amazon RDS, AWS Lambda with accurate savings estimates
- **Security Monitoring:** Automated security gap detection including exposed access keys, security group configurations, and 111+ CSPM controls
- **Performance Analysis:** Resource usage analysis and performance recommendations—throughput optimisation and availability zone redundancy checks
- **Service Limits Monitoring:** Real-time tracking of AWS service quotas and limits with alerts when approaching thresholds
- **Organisational View:** Centralised visibility across multiple AWS accounts through integration with AWS Organizations for enterprise-wide governance
- **Trusted Advisor Priority:** Prioritised and context-driven recommendations with collaboration features for Enterprise Support customers including multi-account views
- **CloudWatch Integration:** Automated notifications and workflow triggers through Amazon CloudWatch Events and Amazon EventBridge for operational automation
- **Well-Architected Framework:** Actionable recommendations to optimise cloud infrastructure according to AWS Well-Architected Framework principles

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/premiumsupport/

FAQ: https://aws.amazon.com/premiumsupport/faqs/

Technical Overview

Backup & Recovery Capabilities

Backup: AWS Trusted Advisor does not provide direct backup capabilities as it is an advisory service analysing configurations rather than storing data. However, it provides recommendations for backup strategies, identifies resources without proper backup configurations, and validates backup readiness through integration with AWS Resilience Hub and other AWS services.

Recovery: Trusted Advisor provides critical guidance for disaster recovery planning by identifying single points of failure, recommending availability zone redundancy, and providing fault tolerance checks through AWS Resilience Hub integration. Supports customer RPO (backup strategy recommendations and multi-AZ deployments) and RTO (redundancy configurations and disaster recovery architectures) objectives.

Setup Process

Onboarding: Customers access AWS Trusted Advisor through AWS Support plans via AWS Management Console. No installation is required. The service automatically analyses environments and provides recommendations. Enterprise customers need AWS Organizations setup for accessing AWS Trusted Advisor Priority features.

Offboarding: The service automatically stops generating recommendations when support plans are discontinued or accounts are closed. Historical check results are automatically removed without specific data removal.

Data Management

- **Storage Options:** Stores only recommendation metadata and check results, not actual customer data or configurations
- **Data Removal:** Automatically removes access to historical check results when support plans are discontinued or accounts closed
- **Cross-Region Replication:** Not applicable as service analyses configurations without storing customer data across regions
- **Encryption:** Standard AWS encryption applies to recommendation metadata stored within the advisory service infrastructure

Optimise Cost and Consumption

AWS Trusted Advisor delivers more than 16 specialised cost optimisation checks through AWS Cost Optimization Hub integration, providing accurate savings estimates for Amazon Elastic Compute Cloud, Amazon Relational Database Service, and AWS Lambda services. The service identifies underused resources and provides rightsizing recommendations based on actual usage patterns. Reserved instance and savings plan analysis helps maximise purchase decisions. Organisational visibility across multiple accounts allow cost optimisation strategies with prioritised recommendations for Enterprise Support customers.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/awssupport/latest/user/trusted-advisor.html

Getting Started: https://docs.aws.amazon.com/awssupport/latest/user/get-started-with-aws-trusted-advisor.html

API Reference: https://docs.aws.amazon.com/awssupport/latest/user/get-started-with-aws-trusted-advisor-api.html

Service Quotas: https://docs.aws.amazon.com/general/latest/gr/aws_service_limits.html

AWS Verified Access

General Information

AWS Verified Access is a Zero Trust security service that provides secure access to corporate applications without requiring VPN infrastructure. It evaluates each request based on user identity and device posture before granting access.

Main Benefits

AWS Verified Access takes away the complexity around VPN infrastructure and provides Zero Trust security through application-specific access controls. The service reduces attack surface, prevents lateral movement, supports both HTTP(S) and non-HTTP protocols, and offers detailed logging for compliance with pay-per-use cost optimisation.

Key Use Cases

- **VPN Replacement:** Remove traditional VPN infrastructure while maintaining secure remote access to applications
- **Zero Trust Access:** Implement identity-aware, contextual access controls based on user and device trust
- **Secure Database Access:** Provide secure access to databases, SAP systems, and TCP-based applications
- **Remote Workforce Enablement:** Support hybrid workforces with secure, device-independent access to corporate resources

Core Features

- **Zero Trust Access Control:** Evaluates user identity and device posture before granting access with continuous monitoring
- **Multi-Protocol Support:** Supports HTTP(S) and non-HTTP(S) protocols including TCP, SSH, and RDP for comprehensive coverage
- **Device Posture Assessment:** Integrates with CrowdStrike, Jamf, and JumpCloud for device security validation and trust verification
- **Policy-Based Access:** Uses Cedar policy language for fine-grained, conditional access control based on user attributes
- **Identity Provider Integration:** Supports IAM Identity Center, Active Directory, SAML 2.0, and OpenID Connect providers seamlessly
- **Detailed Logging:** Logs every access attempt with detailed visibility for security incidents and compliance audits
- **Connectivity Client:** Desktop application provides secure access to non-HTTP resources with encrypted tunnelling capabilities
- **Real-Time Evaluation:** Evaluates each application request in real-time based on contextual policies and security requirements

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/verified-access/

FAQ: https://aws.amazon.com/verified-access/faqs/

Technical Overview

Backup & Recovery Capabilities

Backup: AWS Verified Access does not have traditional backup capabilities as a managed access control service. Service configurations and Cedar policies should be backed up through Infrastructure as Code practices using AWS CloudFormation or Terraform for redeployment. Access logs are stored in AWS CloudTrail for historical audit purposes.

Recovery: Recovery involves recreating service configurations rather than data recovery through multi-Region deployment capabilities. Organisations can achieve near-zero RTO by maintaining Infrastructure as Code templates and using Amazon Route 53 health checks for automatic failover. Supports low RTO objectives (near-zero with proper templates) and RPO is not applicable as no customer data is stored.

Setup Process

Onboarding: Setup requires using IAM Identity Center, creating instances, trust providers, groups, and endpoints through AWS console or AWS CLI. Configure DNS CNAME records, Cedar policies, and install Connectivity Client for non-HTTP applications. Step-by-step tutorials provide guided deployment.

Offboarding: Systematic deletion in reverse order: endpoints, groups, instances, then trust providers through console or CLI with confirmation prompts. Remove DNS configurations and client applications with automatic service-linked role cleanup.

Data Management

- **Storage Options:** Managed service stores configurations and policies; access logs retained in AWS CloudTrail for audits
- **Data Removal:** Systematic deletion in reverse order: endpoints, groups, instances, trust providers via console/CLI confirmation
- **Cross-Region Replication:** Multi-region deployment supported; configurations recreated using Infrastructure as Code across regions for disaster recovery
- **Encryption:** TLS/SSL certificate integration via AWS Certificate Manager; encrypted tunnelling through Connectivity Client for non-HTTP resources

Optimise Cost and Consumption

Organisations can consolidate similar applications into groups to share Cedar policies, reducing administrative overhead and endpoint costs. Rightsizing the number of endpoints based on actual usage patterns optimises resource allocation. Non-HTTP applications benefit from 100 free connections per endpoint hour, allowing strategic deployment of database and TCP-based access points. The serverless architecture provides automatic scaling without capacity planning requirements.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/verified-access/latest/ug/what-is-verified-access.html

Getting Started: https://docs.aws.amazon.com/verified-access/latest/ug/getting-started.html

API Reference: https://docs.aws.amazon.com/AWSEC2/latest/APIReference/operation-list-verified-access.html

Service Quotas: https://docs.aws.amazon.com/verified-access/latest/ug/verified-access-quotas.html

AWS VPN

General Information

AWS VPN provides secure connectivity through two services: AWS Site-to-Site VPN which creates IPsec tunnels between on-premises networks and AWS, and AWS Client VPN which offers remote workforce access to cloud and on-premises resources.

Main Benefits

AWS VPN takes away the management overhead of VPN infrastructure while delivering high-performance connectivity of up to 5 Gbps per tunnel. Features automatic failover under 30 seconds, dual tunnel redundancy, and easy integration with AWS services, reducing operational complexity and costs.

Key Use Cases

- **Hybrid Connectivity:** Securely extend on-premises networks to AWS cloud resources using IPsec tunnels
- **Remote Workforce Access:** Allow secure employee access to AWS and on-premises resources via AWS Client VPN
- **Multi-Site Networking:** Connect distributed branch offices through Amazon VPN CloudHub for hub-and-spoke topology
- **Application Migration:** Support migration and multi-site communication with encrypted VPN connections

Core Features

- **Site-to-Site VPN:** IPsec connections with dual tunnels supporting up to 5 Gbps per tunnel
- **Client VPN:** OpenVPN-based remote access with SAML/Active Directory integration and multi-device support
- **VPN Concentrator:** Cost-effective multi-site connectivity supporting up to 100 remote sites per concentrator
- **High Availability:** Dual tunnel redundancy with automatic failover typically under 30 seconds
- **Advanced Security:** Certificate-based authentication, NAT traversal, and integration with AWS security services
- **Global Acceleration:** AWS Global Accelerator integration for Site-to-Site VPN performance enhancement
- **Amazon CloudWatch Integration:** Detailed monitoring with connection logging and metrics for visibility and troubleshooting
- **Multi-Platform Support:** Client compatibility across Windows, macOS, Linux, iOS, and Android operating systems

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/vpn/

FAQ: https://aws.amazon.com/vpn/faqs/

Technical Overview

Backup & Recovery Capabilities

Backup: AWS VPN does not integrate with AWS Backup as it's a connectivity service. Configuration data is automatically maintained by AWS with built-in redundancy. Customers must manually backup customer gateway configurations and Client VPN certificates outside AWS using external storage systems.

Recovery: Automatic tunnel failover typically under 30 seconds, multi-region deployments with Transit Gateway ECMP routing, and cross-region disaster recovery capabilities. Dual tunnel redundancy provides continuous connectivity without data loss. Supports customer RPO (continuous connectivity) and RTO (sub-minute recovery through proper routing configurations) objectives.

Setup Process

Onboarding: Setup via AWS Management Console, AWS CLI, or APIs. AWS Site-to-Site VPN requires creating customer gateways, virtual private gateways, and configuring routing. AWS Client VPN needs certificate generation, importing to Certificate Manager, and creating endpoints with authorisation rules.

Offboarding: Delete VPN connections, endpoints, and associated resources through AWS Management Console, AWS CLI, or API. Connection logs and configuration data automatically removed during resource deletion, supporting complete cleanup.

Data Management

- **Storage Options:** Configuration data stored in AWS control plane with built-in redundancy for connectivity service
- **Data Removal:** Connection logs and configuration data automatically removed when VPN resources are deleted
- **Cross-Region Replication:** Multi-region VPN deployments supported with Transit Gateway ECMP routing for high availability
- **Encryption:** Enterprise-grade IPsec encryption for AWS Site-to-Site VPN and OpenVPN protocol for AWS Client VPN connections

Optimise Cost and Consumption

VPN Concentrator provides shared bandwidth across up to 100 remote sites, reducing multi-site deployment costs by up to 50%. Rightsizing tunnel bandwidth between standard 1.25 Gbps and large 5 Gbps options optimises throughput costs based on actual requirements. Transit Gateway termination allows Equal-Cost Multi-Path (ECMP) routing for enhanced bandwidth usage, while unaccelerated VPN configurations reduce Global Accelerator charges for non-critical workloads.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/vpn/

Getting Started: https://docs.aws.amazon.com/vpn/latest/s2svpn/SetUpVPNConnections.html

API Reference: https://docs.aws.amazon.com/AWSEC2/latest/APIReference/OperationList-query-vpc.html

Service Quotas: https://docs.aws.amazon.com/vpn/latest/s2svpn/vpn-limits.html

AWS Web Application Firewall (AWS WAF)

General Information

AWS WAF is a web application firewall service that protects applications from common attacks like SQL injection and cross-site scripting by monitoring HTTP(S) requests. It uses customisable rules and AWS managed rule groups to allow, block, or count traffic based on defined conditions across multiple AWS services.

Main Benefits

AWS WAF reduces setup time by up to 80% through pre-configured protection packs and removes the need to install separate security appliances with native AWS integration. It provides automatic scaling without capacity planning, pay-as-you-use pricing with no upfront costs, and sub-second latency for real-time request processing across multiple services.

Key Use Cases

- **Web Application Protection:** Shield applications from SQL injection, cross-site scripting, and common exploits
- **Bot Control and Mitigation:** Block malicious bots allowing legitimate search engines using machine learning
- **API Security:** Protect REST and GraphQL APIs from abuse and malicious requests
- **DDoS Protection:** Provide application-layer protection with rate-based rules and automatic baseline detection

Core Features

- **Web Traffic Filtering:** Create rules based on IP addresses, HTTP headers, body, URI strings, and query parameters
- **AWS Managed Rules:** Pre-configured rule groups maintained by AWS for common threats like OWASP Top 10
- **Bot Control:** Advanced bot detection using machine learning to distinguish legitimate from malicious bots
- **Rate-based Rules:** Limit the number of requests from specific IP addresses within defined time windows
- **Fraud Control:** Account takeover prevention and account creation fraud detection with behavioural analysis
- **Real-time Visibility:** Real-time metrics, sampled web requests, and comprehensive logging with Amazon CloudWatch integration
- **Custom Response:** Configure custom response codes, headers, and error pages when blocking requests
- **Geographic Blocking:** Allow or block requests based on country of origin using GeoIP database

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/waf

FAQ: https://aws.amazon.com/waf/faqs

Technical Overview

Backup & Recovery Capabilities

Backup: AWS WAF does not provide integrated backup capabilities for configuration data. Customers must use Infrastructure as Code practices with AWS CloudFormation, AWS CDK, or direct API calls to maintain configuration backups. Rule sets and web ACL configurations should be version controlled externally for backup purposes.

Recovery: Configuration recovery is achieved through multi-region deployment strategies and programmatic recreation of rule configurations, IP sets, and web ACL settings. AWS WAF integrates with Application Recovery Controller for coordinated failover. Supports customer RPO objectives with near-zero targets for active configurations and RTO objectives of minutes to hours depending on automation sophistication.

Setup Process

Onboarding: Get started through the simplified console by creating protection packs tailored to application types. Configure using guided workflows, AWS CloudFormation, or AWS CDK for programmatic deployment. One-click integration with supported AWS services and count mode for testing rules.

Offboarding: Disassociate web ACLs from protected resources, then delete web ACLs, rules, rule groups, IP sets, and regex patterns. Configuration data removes immediately upon deletion. Log data requires separate deletion from respective services.

Data Management

- **Storage Options:** Logs stored in Amazon Simple Storage Service (Amazon S3), Amazon CloudWatch Logs, or Amazon Data Firehose with respective service capabilities
- **Data Removal:** Configuration data removed immediately upon deletion; log data requires separate deletion from storage services
- **Cross-Region Replication:** No built-in cross-region backup; customers must implement strategies using Infrastructure as Code tools
- **Encryption:** Inherits encryption capabilities from integrated storage services like Amazon S3, Amazon CloudWatch Logs, and Amazon Data Firehose

Optimize Cost and Consumption

AWS WAF provides scope-down statements to limit advanced rule evaluation to specific pages, reducing Bot Control costs. Protection packs bundle rules to avoid individual rule charges. Rate-based rules offer cost-effective DDoS protection through automated thresholds. Apply lower-cost rules before expensive managed rules to optimise processing order and reduce evaluation costs for filtered traffic.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/waf/latest/developerguide/waf-chapter.html
Getting Started: https://docs.aws.amazon.com/waf/latest/developerguide/setup-iap-console.html
API Reference: https://docs.aws.amazon.com/waf/latest/APIReference/Welcome.html
Service Quotas: https://docs.aws.amazon.com/waf/latest/developerguide/limits.html

AWS Wavelength

General Information

AWS Wavelength deploys standard AWS compute and storage services to the edge of 5G networks, supporting ultra-low latency applications. It extends virtual private connections (VPCs) to Wavelength Zones for real-time gaming, Augmented Reality and Virtual Reality (AR/VR), and connected vehicles while maintaining consistent regional service connectivity.

Main Benefits

AWS Wavelength delivers ultra-low latency by positioning compute resources within 5G network infrastructure, removing the need for expensive on-premises edge hardware. It provides AWS access with consistent APIs, automatic scaling, and data sovereignty compliance and reduces Motion to Photon latencies for critical applications.

Key Use Cases

- **Real-time Gaming:** Game streaming with ultra-low latency for devices with limited processing power
- **AR/VR:** Reducing Motion to Photon latencies for realistic AR/VR experiences
- **Connected Vehicles:** Facilitating C-V2X functionality for intelligent driving and autonomous vehicle capabilities
- **Smart Factories:** Industrial automation using ML inference for quality control without expensive GPU servers

Core Features

- **EC2 Compute Services:** Multiple Amazon Elastic Compute Cloud (Amazon EC2) instance families (T3, R5, M5, C5, G4) including g4dn.2xlarge with NVIDIA T4 GPUs
- **EBS Storage:** Amazon Elastic Block Store (Amazon EBS) gp2 volumes for persistent block storage with snapshot and restore capabilities
- **Container Support:** Amazon Elastic Kubernetes Service (Amazon EKS) and Amazon Elastic Container Service (Amazon ECS) clusters for Kubernetes and containerised workloads in Wavelength Zones
- **Load Balancing:** Application Load Balancer available in select Wavelength Zones for distributing traffic across targets
- **Carrier Gateway:** Allows connectivity from Wavelength subnet to CSP network, internet, and AWS Region
- **Data Sovereignty:** Built with AWS sovereign-by-design principles providing verifiable control over data location
- **Regional Service Access:** Coherent connection to AWS Regional services like Amazon Simple Storage Service (Amazon S3) and Amazon DynamoDB through APIs
- **Management Tools:** Support for AWS CloudFormation, Amazon CloudWatch, AWS CloudTrail, AWS Systems Manager, and AWS Cost Explorer tools

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/wavelength

FAQ: https://aws.amazon.com/wavelength/faqs

Technical Overview

Backup & Recovery Capabilities

Backup: Amazon EBS volumes support snapshots stored in the parent AWS Region for point-in-time recovery capabilities. AMIs are stored regionally for instance recovery and replication. AWS Backup provides centralised backup policy management where applicable.

Recovery: Disaster recovery requires architecting solutions using multiple Wavelength Zones, Regional failover patterns, or hub-and-spoke models. Customers can implement DNS-based failover using Route 53 health checks. Supports customer RPO (EBS snapshots for data protection) and RTO (easy AMI deployment and automated scaling) objectives.

Setup Process

Onboarding: Choose Wavelength Zones through Amazon EC2 Dashboard, create VPC, configure carrier gateway, and establish subnets. Launch instances with Carrier IP addresses, configure security groups and route tables. Uses AWS Management Console, AWS CLI, and SDKs.

Offboarding: Terminate Amazon EC2 instances, delete Amazon EBS volumes and snapshots, remove Amazon VPC components including carrier gateways. Delete container clusters and Regional service data separately through standard AWS APIs.

Data Management

- **Storage Options:** Amazon EBS gp2 volumes for persistent block storage with snapshot and restore capabilities
- **Data Removal:** Standard AWS deletion process through APIs and console interfaces for consistent resource cleanup
- **Cross-Region Replication:** AMIs stored regionally offer instance recovery and replication across zones and regions
- **Encryption:** AWS Nitro System architecture provides hardware-based security and isolation with data sovereignty controls

Optimise Cost and Consumption

Amazon EC2 Instance Savings Plans reduce compute costs for Wavelength deployments. Right-sizing instance families (T3, R5, M5, C5, G4) optimises workload requirements. Processing data locally in Wavelength Zones before sending results to regional services minimises data transfer costs. Hub-and-spoke architectural patterns allow cost-effective resource sharing and helps to maintain edge performance benefits.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/wavelength/latest/developerguide/what-is-wavelength.html

Getting Started: https://docs.aws.amazon.com/wavelength/latest/developerguide/get-started-wavelength.html

API Reference: https://docs.aws.amazon.com/AWSEC2/latest/APIReference/

Service Quotas: https://docs.aws.amazon.com/wavelength/latest/developerguide/wavelength-quotas.html

AWS Well-Architected Tool

General Information

AWS Well-Architected Tool is a zero-cost service that helps to measure and document workloads against AWS's proven architectural practices—across five core pillars, reliability, security, performance efficiency, cost optimisation, and sustainability. It provides actionable improvement recommendations across these pillars through in-depth architecture reviews.

Main Benefits

Delivers architecture assessment backed by AWS Solutions Architect expertise at zero-cost. Provides prioritised improvement plans with actionable recommendations, supports up to 1,000 workloads per region, and has built-in integration with AWS Trusted Advisor for contextual insights, allowing systematic risk reduction across technology portfolios.

Key Use Cases

- **Architecture Review:** Conduct detailed workload reviews against proven practices to identify risks and improvements
- **Product Launch Governance:** Document architectures and provide governance for new launches helping to ensure organisational alignment
- **Portfolio Risk Management:** Systematically evaluate multiple workloads and track improvements across technology portfolios
- **Compliance and Audit:** Support regulatory requirements with built-in compliance checks and standardised assessment frameworks

Core Features

- **Workload Reviews:** Conduct detailed assessments against AWS Well-Architected Framework pillars with customisable lenses and question sets
- **Custom Lenses:** Create industry-specific or custom evaluation criteria to supplement standard AWS proven practices and requirements
- **Trusted Advisor Integration:** Access real-time infrastructure recommendations and checks integrated within workload reviews for comprehensive insights
- **Milestone Tracking:** Save workload states at specific points in time to track improvements and changes over time
- **Improvement Plans:** Generate prioritised improvement plans with step-by-step guidance for addressing identified risks and optimisation opportunities
- **Collaboration Features:** Share workloads and reviews with team members, AWS accounts, and AWS Organizations for collaborative assessment
- **Cost Explorer Integration:** Connect with AWS Cost Explorer through cost optimisation recommendations for actionable cost reduction insights

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/well-architected-tool

FAQ: https://aws.amazon.com/well-architected-tool/faqs

Technical Overview

Backup & Recovery Capabilities

Backup: AWS Well-Architected Tool does not provide traditional backup capabilities as it operates as a managed assessment platform. Data persistence is handled through AWS's managed infrastructure across regions. Customers can export reports and assessments manually through the console for external record-keeping purposes.

Recovery: The service operates with built-in resilience as part of AWS's managed service architecture with regional availability and service continuity. Recovery capabilities are inherent to the managed service design rather than customer-configurable solutions. Supports customer RPO objectives through reliability pillar guidance and RTO objectives through disaster recovery proven practices recommendations.

Setup Process

Onboarding: Access through AWS Management Console in all AWS Regions without additional setup requirements. Configure IAM permissions for users and activate service integrations like AWS Trusted Advisor. Define workloads by specifying name, description, owner, environment, regions, and accounts.

Offboarding: Delete individual workloads, custom lenses, and review templates through console or API. All associated data including answers, milestones, and improvement plans are permanently removed. Revoke sharing permissions and disable service integrations.

Data Management

- **Storage Options:** Stores workload definitions, review responses, custom lenses, and milestone data within regional service infrastructure
- **Data Removal:** Delete workloads, custom lenses, and templates through console or API with permanent data removal
- **Cross-Region Replication:** Built-in resilience across AWS regions with workloads and assessments accessible from supported regions
- **Encryption:** Data persistence handled through managed service architecture with AWS's underlying infrastructure security standards

Optimise Cost and Consumption

The Cost Optimisation pillar systematically identifies rightsizing opportunities, Savings Plans implementation, and Spot Instance usage to reduce expenditure. Integration with Cost Explorer provides actionable cost reduction insights while AWS Trusted Advisor checks highlight over-provisioned resources and idle assets. Milestone tracking quantifies optimisation improvements over time, helping organisations to measure business value from architecture changes and establish cost-aware operational cultures.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/wellarchitected/latest/userguide/intro.html
Getting Started: https://docs.aws.amazon.com/wellarchitected/latest/userguide/getting-started.html
API Reference: https://docs.aws.amazon.com/wellarchitected/latest/APIReference/Welcome.html
Service Quotas: https://docs.aws.amazon.com/general/latest/gr/wellarchitected.html

AWS X-Ray

General Information

AWS X-Ray is a distributed tracing service that analyses and debugs production applications by tracking requests across microservices architectures. It generates visual service maps, identifies performance bottlenecks, and provides end-to-end visibility into application behaviour and dependencies.

Main Benefits

AWS X-Ray significantly reduces mean time to resolution through faster root cause analysis compared to traditional logging approaches. It offers seamless AWS integration with automatic instrumentation, eliminating complex setup overhead. The generous free tier provides 100,000 traces recorded and 1,000,000 traces retrieved monthly at no cost.

Key Use Cases

- **Performance Optimization:** Identifying latency bottlenecks and performance issues across distributed applications and microservices
- **Error Debugging:** Root cause analysis of application errors and exceptions in production environments
- **Application Monitoring:** Understanding request flow and dependencies between services for operational visibility
- **Microservices Analysis:** Visualizing complex service relationships and interactions in distributed architectures

Core Features

- **End-to-End Tracing:** Tracks requests across multiple services and components for complete application behaviour visibility
- **Service Map Generation:** Creates visual representations of application architecture showing service dependencies and relationships
- **Latency Detection:** Measures server-side and client-side latency to identify performance bottlenecks across distributed systems
- **Error Analysis:** Automatically highlights bugs and errors by analysing response codes and exceptions in code
- **Data Annotation and Filtering:** Allows adding annotations and metadata to traces for easier analysis and filtering
- **Sampling Configuration:** Customizable sampling rules to control trace volume and costs without requiring code modifications
- **Multi-Language SDK Support:** SDKs available for Java, Node.js, Python, Ruby, Go, and .NET development platforms
- **Query and Analytics APIs:** Programmatic access to trace data for building custom analysis and visualization applications

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/xray/

FAQ: https://aws.amazon.com/xray/faqs/

Technical Overview

Backup & Recovery Capabilities

Backup: AWS X-Ray does not provide traditional backup capabilities as it is designed for operational monitoring with automatic 30-day data retention and deletion. Trace data is considered ephemeral operational telemetry rather than persistent business data. Customers needing long-term analysis must export data using X-Ray APIs before the retention period expires.

Recovery: X-Ray operates as a highly available service across multiple Availability Zones within each region and supports multi-region deployment for disaster recovery scenarios. The service indirectly supports RTO objectives by enabling rapid root cause analysis and reducing mean time to resolution, minimizing application downtime through faster problem diagnosis.

Setup Process

Onboarding: Setup involves instrumenting applications using X-Ray SDKs or AWS Distro for OpenTelemetry, configuring the X-Ray daemon or CloudWatch agent for data collection, and deploying instrumented applications. AWS Lambda enables tracing with simple configuration changes.

Offboarding: Disable tracing by removing SDK instrumentation, stopping the X-Ray daemon or CloudWatch agent, and deleting custom sampling rules through console or API. Trace data automatically purges after 30-day retention period.

Data Management

- **Storage Options:** Managed storage with automatic 30-day retention for trace data considered ephemeral operational telemetry
- **Data Removal:** Automatic deletion after 30 days with no manual intervention required for data removal
- **Cross-Region Replication:** Multi-region tracing supported with cross-account observability features for centralized monitoring across regions
- **Encryption:** Standard AWS encryption capabilities for trace data in transit and at rest

Optimize Cost and Consumption

Customizable sampling rules control trace volume without code modifications, defaulting to 1 request per second plus 5% of additional requests to manage costs effectively. Strategic annotation usage reduces indexed data volume while the built-in sampling algorithm automatically maintains statistical significance. Intelligent sampling strategies capture representative traces based on application criticality, optimizing resource consumption while preserving analytical value for performance monitoring and debugging activities.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/xray/latest/devguide/index.html
Getting Started: https://docs.aws.amazon.com/xray/latest/devguide/xray-gettingstarted.html
API Reference: https://docs.aws.amazon.com/xray/latest/api/
Service Quotas: https://docs.aws.amazon.com/general/latest/gr/xray.html

FreeRTOS

General Information

FreeRTOS is an open-source real-time operating system for microcontrollers and small microprocessors. It provides IoT connectivity libraries, security features, and built-in AWS IoT integration for edge devices. The service allows secure deployment and management of small, low-power connected devices.

Main Benefits

FreeRTOS delivers enterprise-grade reliability with open-source flexibility, supporting more than 40 architectures with 18 years of proven stability. It offers the only RTOS with native AWS IoT integration, simplifying integration while reducing development time and costs through comprehensive security libraries and simplified connectivity.

Key Use Cases

- **Smart Meters:** Monitor electricity usage in real time and allow efficient load balancing for utilities
- **Oil Pump Sensors:** Monitor well output and process data locally via AWS IoT Greengrass for real-time analytics
- **Industrial IoT Devices:** Build connected security systems, fitness trackers, and sensor networks for commercial applications
- **Consumer Appliances:** Facilitate home automation, voice control, and connected home product development at scale

Core Features

- **FreeRTOS Kernel:** Market-leading real-time operating system kernel supporting more than 40 architectures including RISC-V and ARMv8-M
- **IoT Connectivity Libraries:** Libraries for MQTT, HTTP, TCP/IP, Wi-Fi, Bluetooth Low Energy, and cellular connectivity
- **Security Libraries:** TLS v1.2, data encryption, key management, code signing, and secure communications
- **AWS IoT Integration:** Native integration with AWS IoT Core, Device Shadows, Device Defender, and Device Management
- **Over-the-Air Updates:** Secure firmware updates and device management capabilities with code signing
- **Long Term Support (LTS):** LTS releases with security updates and critical bug fixes for at least two years
- **Extended Maintenance Plan:** Up to 10 additional years of security patches and critical bug fixes beyond LTS
- **Edge Computing Capabilities:** Process data locally on microcontrollers while maintaining cloud connectivity for management and analytics

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://aws.amazon.com/freertos/

FAQ: https://aws.amazon.com/freertos/faqs/

Technical Overview

Backup & Recovery Capabilities

Backup: FreeRTOS relies on AWS IoT Device Shadows for cloud-based device state storage and firmware images in device flash memory for OTA updates. Customers implement application-level backup strategies using Amazon S3 or Amazon DynamoDB through AWS service APIs and console interfaces.

Recovery: Device recovery occurs through OTA updates restoring firmware within minutes to hours, Device Shadow synchronisation providing near real-time state recovery, and automatic reconnection to AWS IoT services. The distributed edge architecture provides natural resilience against single points of failure. Supports customer RPO objectives with seconds-level Device Shadow sync and RTO objectives through OTA recovery capabilities.

Setup Process

Onboarding: Download FreeRTOS from FreeRTOS.org, select qualified hardware from AWS Partner Device Catalog, and setup development environment with supported IDEs (Eclipse, Keil µVision). Configure AWS IoT credentials, X.509 certificates, build demo applications, and validate connectivity using AWS IoT Device Tester.

Offboarding: Remove device certificates from AWS IoT Device Registry, delete device shadows and metadata, deprovision devices from AWS IoT Device Management. Control device-side data removal by reflashing firmware or resetting devices through standard AWS service APIs.

Data Management

- **Storage Options:** Device flash memory for firmware images, AWS IoT Device Shadows for cloud-based state storage
- **Data Removal:** Remove device certificates from AWS IoT Device Registry, delete device shadows and metadata via APIs
- **Cross-Region Replication:** Uses standard AWS IoT services cross-region capabilities for cloud-side data replication and backup
- **Encryption:** TLS v1.2, data encryption, key management, code signing, and secure communications through security libraries

Optimise Cost and Consumption

FreeRTOS's efficient resource usage minimises hardware requirements, allowing deployment on low-cost microcontrollers with just 16KB RAM. Edge computing capabilities process data locally, reducing cloud processing and data transfer costs. Over-the-air update mechanisms reduce field service expenses through remote firmware management. The Extended Maintenance Plan offers predictable annual costs for long-term budget planning.

Technical Resources

Developer Guide: https://docs.aws.amazon.com/freertos/latest/userguide/what-is-freertos.html
Getting Started: https://docs.aws.amazon.com/freertos/latest/userguide/freertos-getting-started.html
API Reference: https://docs.aws.amazon.com/freertos/latest/lib-ref/html1/annotated.html
Service Quotas: https://docs.aws.amazon.com/general/latest/gr/amazon-freertos.html

Kiro

General Information

Kiro is AWS's native AI-powered development environment that enables developers to build applications using natural language specifications and AI assistance. It provides comprehensive development capabilities with integrated AWS services access, contextual code analysis, and automated infrastructure code generation.

Main Benefits

Kiro dramatically reduces development time from weeks to hours through natural language programming and automated code generation. It enables citizen development for non-technical users, accelerates cloud modernization assessment from weeks to days, and automatically applies AWS best practices for secure, optimized deployments.

Key Use Cases

- **Cloud Modernization:** Automate assessment, planning, and implementation for migrating legacy applications to AWS
- **Infrastructure Optimization:** Find cost optimization opportunities and automatically apply recommendations to Infrastructure as Code
- **AI Agent Development:** Build and deploy production-ready AI agents with integrated Amazon Bedrock AgentCore support
- **Citizen Development:** Enable business users to build applications using natural language without coding skills

Core Features

- **AI-Powered Development:** Natural language code generation and contextual analysis with real-time AI assistance
- **Model Context Protocol Integration:** Connects with specialized MCP servers for extended domain-specific capabilities
- **Infrastructure Code Generation:** Automated generation of Infrastructure as Code with AWS best practices
- **Visual Project Management:** Comprehensive project visualization and management capabilities for development teams
- **Cost Optimization:** Automated identification and application of cost optimization recommendations to infrastructure
- **Collaborative Workflows:** Multi-user collaboration features enabling seamless teamwork across development projects
- **Spec-Driven Development:** Requirements gathering, design, implementation planning, and execution workflow methodology
- **Contextual Code Analysis:** Advanced code analysis capabilities that understand project context and AWS patterns

Pricing

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Key Resources

Service Page: https://kiro.dev/
FAQ: https://kiro.dev/faq/

Technical Overview

Backup & Recovery Capabilities

Backup: Kiro relies on standard development environment backup practices focusing on project workspaces, configuration files, and integration settings. Project preservation is handled through integration with version control systems and standard code versioning practices via CLI and IDE interfaces.

Recovery: Recovery involves standard development tool restoration practices and reconfiguration of MCP integrations and AWS connections. Disaster recovery focuses on restoring development environment configurations and reconnecting to AWS services. Recovery objectives depend on underlying infrastructure hosting the development environment and standard development workflow practices.

Setup Process

Onboarding: Download Kiro CLI or access Kiro IDE through AWS with multiple entry points including direct download, existing IDE integration, and web-based access. Configure AWS credentials, install Python 3.10+ dependencies, and set up MCP servers. Initial setup takes just a few hours with guided tutorials.

Offboarding: Clear project workspaces, disconnect MCP server integrations, remove AWS credential configurations, and delete locally stored project artifacts following standard development environment cleanup practices.

Data Management

- **Storage Options:** Local project workspaces, configuration files, and integration settings with Amazon S3 integration for storage management
- **Data Removal:** Clear project workspaces, disconnect MCP integrations, remove AWS credentials, and delete locally stored project artifacts
- **Cross-Region Replication:** Cross-region replication capabilities not explicitly documented in available Kiro documentation and technical specifications
- **Encryption:** Encryption capabilities follow AWS best practices automatically but specific encryption details not documented in available information

Optimize Cost and Consumption

Kiro's automated cost optimization engine identifies opportunities across entire infrastructure files and applies recommendations directly to Infrastructure as Code. The platform suggests storage tier improvements, networking enhancements, compute right-sizing, and auto-scaling configurations. Embedded cost optimization rules ensure every new resource is optimized from deployment, while the context system allows organizations to integrate their cost optimization best practices into development workflows.

Technical Resources

Getting Started: https://kiro.dev/docs/enterprise/getting-started/

Kiro IDE Reference: https://kiro.dev/docs/

Kiro CLI Reference: https://kiro.dev/cli/