G-Cloud 13 Amazon Web Services EMEA SARL, UK Branch (AWS) – Cloud Compute Infrastructure Services Service Definition Catalogue

May 2022

G-Cloud 13 Amazon Web Services EMEA SARL, UK Branch (AWS) – Cloud Compute Infrastructure Services Service Definition Catalogue

Table of Contents

1. Introduction

This document provides service definitions for the Amazon Web Services EMEA SARL, UK Branch (AWS) Service Offerings included in the G-Cloud 13 framework catalogue. We have broken out service definitions in accordance with Invitation to Tender (ITT) requirements.

1.1. 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 descriptions 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:

- Cloud Compute Infrastructure Services (Lot 1 & 2)
- VMWare Cloud on AWS (Lot 1)
- Professional Services (Lot 3)
- Support Services (Lot 3)
- Training Services (Lot 3)
- AWS Managed Services (Lot 3)

This AWS Cloud Compute Infrastructure Services Service Definition document describes the key features for each of the different Cloud Compute Services available to Customers on G-Cloud 13 in Lots 1 & 2.

Notwithstanding that AWS has combined its service descriptions 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 13 Call-Off Contract. AWS recommends that Buyers list all of the Digital Marketplace Service ID's for every service described in this document 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 through aws-cloud@amazon.com.

Please note that we have consolidated common elements of each Service Offering (e.g., on-boarding and off-boarding) and have provided descriptions for these common elements that apply equally to each Service Offering. To find out more about AWS on G-Cloud and AWS Cloud services, visit us at AWS on G-Cloud UK.

The AWS Free Tier enables you to gain free, hands-on experience with AWS products and services. It is designed to enable you to get hands-on experience with AWS at no charge for 12 months after you sign up. After creating your AWS account, you can use products and services listed at http://aws.amazon.com/free/ for free within certain usage limits.

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 willing to provide additional information about our services upon request.

2. AWS Security Assurance

Moving IT infrastructure to AWS means that both the customer and AWS have important roles in the operation and management of security in their areas of responsibility. AWS operates, manages, and controls the components from the host operating system and virtualisation layer down to the physical security of the facilities in which the services operate. The customer is responsible for management of the guest operating system (including updates and security patches to the guest operating system) and associated application software, as well as the configuration of the AWS provided security group firewall and other security-related features. The customer will generally connect to the AWS environment through services provided by third parties (e.g., internet service providers). AWS does not provide these connections, and the customer should consider the security of such connections and the security responsibilities of such third parties in relation to their systems. This is really no different from working with a network service provider who brings connectivity to on-premises data centres.

We are vigilant about the security of our underlying cloud environment and have implemented sophisticated technical and organisational measures against unauthorised access. Customers can validate the security controls in place within the AWS environment through AWS certifications and reports, including the AWS System and Organisation Controls (SOC) 1, 2, and 3 reports, International Organisation for Standardization (ISO) 27001 certification, and Payment Card Industry Data Security Standard (PCI DSS) compliance reports. These reports and certifications are produced by independent third-party auditors and attest to the design and operating effectiveness of AWS security controls. The applicable AWS compliance certifications and reports can be requested at https://aws.amazon.com/compliance/contact. More information on AWS compliance certifications, reports, and alignment with best practices and standards can be found at https://aws.amazon.com/compliance and https://aws.amazon.com/compliance/programs/.

British Standard 7858:2019

Buyers selecting AWS Services and expressly requiring AWS conformity to BS7858:2019 acknowledge that AWS scopes BS7858:2019 compliance to those AWS employees with 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.

2.1. Information Assurance

The following subsections provide information relating to information assurance.

2.1.1. ISO 27001 Certification

AWS is certified under the ISO 27001 standard. ISO 27001 is a widely adopted global security standard that outlines the requirements for information security management systems. It provides a systematic approach to managing company and customer information that is based on periodic risk assessments. In order to achieve the certification, a company must show it has a systematic and ongoing approach to managing information security risks that affect the confidentiality, integrity, and availability of company and customer information.

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

2.1.2. 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. In order to provide you with more information regarding NCSC UK's Cloud Security Principles and to make an informed decision when performing risk assessments, we have published a whitepaper called Using AWS in the Context of NCSC UK's Cloud Security Principles.

This whitepaper provides insights into implementation and assurance approaches within AWS based on the published guidance for each of the 14 Cloud Security Principles and provides an in-depth view into the AWS implementation approach in relation to the Cloud Security Principles. Based on this information, UK public sector organisations and their information security functions can conduct informed risk assessments and select the appropriate AWS Cloud services for their cloud environment.

2.2. 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/de/compliance/gdpr-center/
- AWS EU Data Protection website: https://aws.amazon.com/compliance/eudata-protection/

3. Amazon API Gateway

3.1. Service Overview

Amazon API Gateway is a fully managed service that makes it easy for developers to create, publish, maintain, monitor, and secure APIs at any scale. APIs act as the "front door" for applications to access data, business logic, or functionality from your backend services. Using API Gateway, you can create RESTful APIs and WebSocket APIs that enable real-time two-way communication applications. API Gateway supports containerized and serverless workloads, as well as web applications.

API Gateway handles all the tasks involved in accepting and processing up to hundreds of thousands of concurrent API calls, including traffic management, CORS support, authorization and access control, throttling, monitoring, and API version management. API Gateway has no minimum fees or start-up costs. You pay for the API calls you receive and the amount of data transferred out and, with the API Gateway tiered pricing model, you can reduce your cost as your API usage scales.

3.1.1. Features

- **Support for RESTful APIs and WebSocket APIs:** With API Gateway, you can create RESTful APIs using either HTTP APIs or REST APIs. HTTP APIs are the best way to build APIs that do not require API management features. HTTP APIs are optimized for serverless workloads and HTTP backends—they offer up to 71% cost savings and 60% latency reduction compared to REST APIs from API Gateway.
- **Private integrations with AWS ELB & AWS Cloud Map:** With API Gateway, you can route requests to private resources in your VPC. Using HTTP APIs, you can build APIs for services behind private ALBs, private NLBs, and IP-based services registered in AWS Cloud Map, such as ECS tasks.
- **Easy API Creation and Deployment:** With API Gateway, you can quickly and easily create a custom API to your code running in AWS Lambda and then call the Lambda code from your API.
- **API Operations Monitoring:** After an API is deployed and in use, API Gateway provides you with a dashboard to visually monitor calls to the services. The API Gateway console is integrated with Amazon CloudWatch, so you get backend performance metrics such as API calls, latency, and error rates.
- **AWS Authorization:** To authorize and verify API requests to AWS services, API Gateway can help you leverage signature version 4 for REST APIs and WebSocket APIs.
- **API Keys for Third-Party Developers:** If you're using REST APIs, API Gateway helps you manage the ecosystem of third-party developers accessing your APIs.
- **SDK Generation:** If you're using REST APIs, API Gateway can generate client SDKs for a number of platforms which you can use to quickly test new APIs from your applications and distribute SDKs to third-party developers.
- **API Lifecycle Management:** If you're using REST APIs, API Gateway lets you run multiple versions of the same API simultaneously so that applications can continue to call previous API versions even after the latest versions are published.

3.1.2. Benefits

- **Efficient API development:** Run multiple versions of the same API simultaneously with API Gateway, allowing you to quickly iterate, test, and release new versions. You pay for calls made to your APIs and data transfer out and there are no minimum fees or upfront commitments.
- **Performance at any scale:** Provide end users with the lowest possible latency for API requests and responses by taking advantage of our global network of edge locations using Amazon CloudFront. Throttle traffic and authorize API calls to ensure that backend operations withstand traffic spikes and backend systems are not unnecessarily called.
- **Cost savings at scale:** API Gateway provides a tiered pricing model for API requests. With an API Requests price as low as $0.90 per million requests at the highest tier, you can decrease your costs as your API usage increases per region across your AWS accounts.
- **Easy monitoring:** Monitor performance metrics and information on API calls, data latency, and error rates from the API Gateway dashboard, which allows you to visually monitor calls to your services using Amazon CloudWatch.
- **Flexible security controls:** Authorize access to your APIs with AWS Identity and Access Management (IAM) and Amazon Cognito. If you use OAuth tokens, API Gateway offers native OIDC and OAuth2 support. To support custom authorization requirements, you can execute a Lambda authorizer from AWS Lambda.
- **RESTful API options:** Create RESTful APIs using HTTP APIs or REST APIs. HTTP APIs are the best way to build APIs for a majority of use cases—they're up to 71% cheaper than REST APIs. If your use case requires API proxy functionality and management features in a single solution, you can use REST APIs.

3.2. Backup/Restore and Disaster Recovery

This service allows for backup and recovery. The service can back up swagger exports. Users control this via a manual or scheduled API call. Users schedule and recover backups through a web interface.

3.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

3.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/apigateway/
- **Service quotas:** https://docs.aws.amazon.com/apigateway/latest/developerguide/limits.html
- **Service FAQs:** https://aws.amazon.com/api-gateway/faqs/

3.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/apigateway/ and the following links for comprehensive technical documentation regarding this service.

- **Developer Guide:** Provides a conceptual overview of Amazon API Gateway and includes detailed instructions for using the service.
- **V1 API Reference:** Describes the API Gateway Version 1 API operations for creating and deploying REST APIs.
- **V2 API Reference:** Describes the API Gateway Version 2 API operations for creating and deploying APIs.

4. Amazon AppFlow

4.1. Service Overview

Amazon AppFlow is a fully managed integration service that enables you to securely transfer data between Software-as-a-Service (SaaS) applications like Salesforce, SAP, Zendesk, Slack, and ServiceNow, and AWS services like Amazon S3 and Amazon Redshift, in just a few clicks. With AppFlow, you can run data flows at enterprise scale at the frequency you choose - on a schedule, in response to a business event, or on demand. You can configure data transformation capabilities like filtering and validation to generate rich, ready-to-use data as part of the flow itself, without additional steps. AppFlow automatically encrypts data in motion, and allows users to restrict data from flowing over the public Internet for SaaS applications that are integrated with AWS PrivateLink, reducing exposure to security threats.

4.1.1. Features

- **Point and click user interface:** You can use AppFlow to set up data flows in minutes - no coding required. A point and click user interface enables you to select your data sources and destinations, configure optional transformations and validations, and run your flow without creating dependencies on technical teams.
- **Flexible data flow triggers:** AppFlow enables you to run data flows on demand to do bulk transfers or tests, set up a routine schedule to keep data in sync, or run flows in response to business events like the creation of a sales opportunity, the status change of a support ticket, or the completion of a registration form.
- **Native SaaS integrations:** AppFlow include native integration with the Software-as-a-Service (SaaS) applications used daily for business operations, including Salesforce, Marketo, Slack and more - and more integrations are planned. With AppFlow, you can easily transfer data from any supported SaaS application in just a few clicks.
- **Easy to use field mapping:** You use the AppFlow interface to map source and destination fields together all at once through bulk mapping, or map each field one at a time. For data flows with a large number of fields, you can also upload a csv file to map many fields quickly.
- **Pay as you go:** Amazon AppFlow offers significant cost-savings advantage compared to building connectors in-house or using other application integration services. There are no upfront charges or licensing fees to use AppFlow, and customers only pay for the number of flows they run and the volume of data processed.
- **High scale data transfer:** Amazon AppFlow can run up to 100 GB of data per flow, which enables you to easily transfer millions of Salesforce records, Marketo leads or Zendesk tickets - all while running a single flow.
- **Enterprise grade data transformations:** AppFlow enables you to perform data transformations like mapping, merging, masking, filtering, and validation as part of the flow itself, so there's no need for additional steps. For example, you can validate that

data is in the right numeric format, merge first and last names, or mask credit card details.

- **Data privacy defaults through PrivateLink:** AWS PrivateLink simplifies the security of data shared with cloud-based applications by eliminating the exposure of data to the public Internet. For SaaS applications that have PrivateLink enabled, AppFlow automatically creates and configures private endpoints so your data remains private by default.
- **Custom encryption keys:** All data flowing through AppFlow is encrypted at rest and in transit, and you can encrypt data with AWS keys, or bring your own custom keys.

4.1.2. Benefits

- **Integrate with a few clicks:** Anyone can use AppFlow to integrate applications in a few minutes – no more waiting days or weeks to code custom connectors. Features like data pagination, error logging, and network connection retries are included by default so there's no coding or management. With Appflow, data flow quality is built in, and you can enrich the flow of data through mapping, merging, masking, filtering, and validation as part of the flow itself.
- **Transfer data at massive scale:** AppFlow easily scales up without the need to plan or provision resources, so you can move large volumes of data without breaking it down into multiple batches. AppFlow can run up to 100 GB per flow, which enables you to easily transfer millions of Salesforce records or Zendesk events or Marketo responses or other data - all while running a single flow.
- **Automate data security:** All data flowing through AppFlow is encrypted at rest and in transit, and you can encrypt data with AWS keys, or bring your own custom keys. With AppFlow, you can use your existing Identity and Access Management (IAM) policies to enforce fine-grained permissions, rather than creating new policies. For SaaS integrations with AWS PrivateLink enabled, data is secured from the public internet by default.

4.2. Backup/Restore and Disaster Recovery

This requirement is not applicable for this service. For additional information beyond what is described herein, please refer to http://aws.amazon.com/documentation/.

4.3. Pricing Overview

"Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace."

4.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/appflow/index.html
- **Service quotas:** https://docs.aws.amazon.com/appflow/latest/userguide/service-quotas.html
- **Service FAQs:** https://aws.amazon.com/appflow/faqs/

4.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/appflow/index.html and the following links for comprehensive technical documentation regarding this service.

- **User Guide:** Describes key concepts of Amazon AppFlow and provides instructions for using the features of Amazon AppFlow.
- **API Reference Guide:** Describes all the API operations for Amazon AppFlow in detail. Also provides sample requests, responses, and errors for the supported web service protocols.
- **AWS CloudFormation User Guide for Amazon AppFlow:** Documents the reference information for all Amazon AppFlow resource and property types that are supported by AWS CloudFormation.
- **Amazon AppFlow SDK for Java API Reference:** Describes all the API operations for the AWS SDK for JavaScript in detail. Also provides sample requests, responses, and errors for the supported web service protocols.
- **Amazon AppFlow SDK for Python (Boto3):** Describes all of the classes included in the Amazon AppFlow SDK for Python.
- **AWS CLI Reference for Amazon AppFlow:** Documents the Amazon AppFlow commands available in the AWS Command Line Interface (AWS CLI).
- **Amazon AppFlow Custom Connector SDK (Python):** Use the Python Custom Connector SDK to build custom source and destination connectors for Amazon AppFlow. With custom connectors, you can transfer data between private APIs, on-premise systems, other cloud services, and AWS.
- **Amazon AppFlow Custom Connector SDK (Java):** Use the Java Custom Connector SDK to build custom source and destination connectors for Amazon AppFlow. With custom connectors, you can transfer data between private APIs, on-premise systems, other cloud services, and AWS.

5. Amazon AppStream 2.0

5.1. Service Overview

Amazon AppStream 2.0 lets you move your desktop applications to AWS, without rewriting them. It's easy to install your applications on AppStream 2.0, set launch configurations, and make your applications available to users. AppStream 2.0 offers a wide selection of virtual machine options so that you can select the instance type that best matches your application requirements, and set the auto-scale parameters so that you can easily meet the needs of your end users. AppStream 2.0 allows you to launch applications in your own network, which means your applications can interact with your existing AWS resources.

5.1.1. Features

- **Simple application management:** Amazon AppStream 2.0 enables you to quickly and easily install, test, and update your applications using the image builder. Any application that runs on Microsoft Windows Server 2012 R2, Windows Server 2016, Windows Server 2019, or Amazon Linux 2 is supported and you don't need to make any modifications.

- **Multiple virtual machine types:** Amazon AppStream 2.0 runs your applications on virtual machines (VMs) called streaming instances, which provide the GPU, CPU, memory, storage, and networking capacity you need.
- **GPU optimized virtual machines:** Amazon AppStream 2.0 offers Graphics Design, Graphics Pro and Graphics G4 instance families.
- **Amazon VPC support:** With Amazon AppStream 2.0, your desktop applications can be launched inside an Amazon Virtual Private Cloud (VPC). You can use VPC security groups to provide granular access control to streaming instances, and to manage users' access to the resources in your VPC, such as your databases, file shares, license servers, and application servers.
- **Identity federation:** Amazon AppStream 2.0 supports federated sign-in using SAML 2.0. Users can sign in to AppStream 2.0 using their existing credentials, and start streaming applications.
- **Microsoft Active Directory domain support:** Your Amazon AppStream 2.0 Always-On and On-Demand fleet streaming instances and image builders can connect to your Microsoft Active Directory (AD) domain.
- **Smart card support:** Your users can use their Personal Identity Verification (PIV) and Common Access Card (CAC) smart card, and other types of smart cards, to sign in to a Windows OS based AppStream 2.0 streaming instance that is joined to a Microsoft Active Directory domain.
- **Monitoring:** Amazon AppStream 2.0 allows you to monitor the utilization of your AppStream 2.0 fleet resources using Amazon CloudWatch metrics. With Elastic fleets, you can track the number of instances that are being used over time. With Always-On and On-Demand fleets, you can see the size of your fleet, the number of instances you have running, and the available capacity for new connections.
- **Elastic fleets:** Elastic fleets are a serverless fleet type that allows you to deliver your applications to end users without needing to predict concurrency, create or manage auto scaling policies, or create any images. Your applications are stored within virtual hard disks that are downloaded to streaming instances on user request simplifying how you deliver streaming applications to users.
- **Fleet auto scaling:** With Always-On and On-Demand fleets, you can use auto scaling policies to adjust the number of instances that are running to reduce your streaming costs. With Elastic fleets, AppStream 2.0 manages the size of the fleet for you without needing to use autoscaling policies.
- **Programmatic control:** Amazon AppStream 2.0 includes APIs that you can use to easily integrate and extend the service. The APIs enable you to create, update, and delete AppStream 2.0 resources, and provide detailed information about resource states. You can create URLs for administrators to connect to their image builders to install applications, and create URLs for users to access their AppStream 2.0 applications.
- **Browser and client access:** Amazon AppStream 2.0 allows you to access your desktop applications from HTML5-capable browsers such as Google Chrome, Mozilla Firefox, Microsoft Internet Explorer, and Microsoft Edge. No plug-ins are needed.

- **Persistent storage:** Amazon AppStream 2.0 offers multiple options for persistent file storage to allow users to store and retrieve files between their application streaming sessions. You can use a home folder backed by Amazon S3, Google Drive for G Suite, or Microsoft OneDrive for Business.
- **Simple user interface:** Amazon AppStream 2.0 offers an intuitive user interface, making it easy to control your experience.
- **NICE DCV protocol:** Amazon AppStream 2.0 uses NICE DCV technology to provide secure, high-performance access to your applications.
- **HTTPS access:** With Amazon AppStream 2.0, your application streams and user input flows through a secure streaming gateway on AWS over HTTPS. Streaming instances are not directly accessible from the internet, and users can only access their applications through the streaming gateway after being authenticated.
- **Globally available:** Amazon AppStream 2.0 is available in multiple AWS regions globally. You can host your AppStream 2.0 resources in multiple AWS regions, and direct users to the closest AWS region for the best end-user experience.

5.1.2. Benefits

- **Empower your remote workforce:** React quickly to changing conditions with access to applications and desktops from anywhere.
- **Strengthen security:** Store data on AWS instead of vulnerable endpoint devices.
- **Optimize costs:** Benefit from on-demand cloud scalability with a range of compute, memory, and storage options.
- **Reduce downtime:** Fully managed application delivery and reliable AWS infrastructure offering 99.9% uptime.

5.2. Backup/Restore and Disaster Recovery

This service allows for backup and recovery. The service can back up files. Users control this Via the "home folder", which is Appstream 2.0's native persistent storage option. More information is available at https://docs.aws.amazon.com/appstream2/latest/developerguide/home-folders.html. Users schedule and recover backups through a web interface.

5.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

5.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/appstream2/
- **Service quotas:** https://docs.aws.amazon.com/appstream2/latest/developerguide/limits.html
- **Service FAQs:** https://aws.amazon.com/appstream2/faqs/

5.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/appstream2/ and the following links for comprehensive technical documentation regarding this service.

- **Administration Guide:** Provides a conceptual overview of Amazon AppStream 2.0 and includes detailed instructions for using the features available in this service.
- **API Reference:** Describes all the API operations for Amazon AppStream 2.0 in detail. Also provides sample requests, responses, and errors for the supported web services protocols.

6. Amazon Athena

6.1. Service Overview

Amazon Athena is an interactive query service that makes it easy to analyse data in Amazon S3 using standard SQL. Athena is serverless, so there is no infrastructure to manage, and you pay only for the queries that you run.

Athena is easy to use. Simply point to your data in Amazon S3, define the schema, and start querying using standard SQL. Most results are delivered within seconds. With Athena, there's no need for complex ETL jobs to prepare your data for analysis. This makes it easy for anyone with SQL skills to quickly analyse large-scale datasets.

Athena is out-of-the-box integrated with AWS Glue Data Catalog, allowing you to create a unified metadata repository across various services, crawl data sources to discover schemas and populate your Catalog with new and modified table and partition definitions, and maintain schema versioning.

6.1.1. Features

- **Serverless. Zero infrastructure. Zero administration:** Amazon Athena is serverless, so there is no infrastructure to manage. You don't need to worry about configuration, software updates, failures or scaling your infrastructure as your datasets and number of users grow. Athena automatically takes care of all of this for you, so you can focus on the data, not the infrastructure.
- **Easy to get started:** To get started, log into the Athena console, define your schema using the console wizard or by entering DDL statements, and immediately start querying using the built-in query editor. You can also use AWS Glue to automatically crawl data sources to discover data and populate your Data Catalog with new and modified table and partition definitions.
- **Easy to query, just use standard SQL:** Amazon Athena uses Presto, an open source, distributed SQL query engine optimized for low latency, ad hoc analysis of data. This means you can run queries against large datasets in Amazon S3 using ANSI SQL, with full support for large joins, window functions, and arrays.
- **Pay per query:** With Amazon Athena, you pay only for the queries that you run. You are charged based on the amount of data scanned by each query. You can get significant cost savings and performance gains by compressing, partitioning, or converting your data to a columnar format, because each of those operations reduces the amount of data that Athena needs to scan to execute a query.

- **Fast performance:** With Amazon Athena, you don't have to worry about managing or tuning clusters to get fast performance. Athena is optimized for fast performance with Amazon S3. Athena automatically executes queries in parallel, so that you get query results in seconds, even on large datasets.
- **Highly available & durable:** Amazon Athena is highly available and executes queries using compute resources across multiple facilities, automatically routing queries appropriately if a particular facility is unreachable. Athena uses Amazon S3 as its underlying data store, making your data highly available and durable.
- **Secure:** Amazon Athena allows you to control access to your data by using AWS Identity and Access Management (IAM) policies, access control lists (ACLs), and Amazon S3 bucket policies. With IAM policies, you can grant IAM users fine-grained control to your S3 buckets.
- **Integrated:** Amazon Athena integrates out-of-the-box with AWS Glue. With Glue Data Catalog, you will be able to create a unified metadata repository across various services, crawl data sources to discover data and populate your Data Catalog with new and modified table and partition definitions, and maintain schema versioning. You can also use Glue's fully-managed ETL capabilities to transform data or convert it into columnar formats to optimize query performance and reduce costs. Learn more about AWS Glue.
- **Federated query:** Athena enables you to run SQL queries across data stored in relational, non-relational, object, and custom data sources. You can use familiar SQL constructs to JOIN data across multiple data sources for quick analysis, and store results in Amazon S3 for subsequent use.
- **Machine learning:** You can invoke your SageMaker Machine Learning models in an Athena SQL query to run inference. The ability to use ML models in SQL queries makes complex tasks such anomaly detection, customer cohort analysis and sales predictions as simple as writing a SQL query. Athena makes it easy for anyone with SQL experience to run ML models deployed on Amazon SageMaker.

6.1.2. Benefits

- **Start querying instantly:** Athena is serverless. You can quickly query your data without having to setup and manage any servers or data warehouses. Just point to your data in Amazon S3, define the schema, and start querying using the built-in query editor. Amazon Athena allows you to tap into all your data in S3 without the need to set up complex processes to extract, transform, and load the data (ETL).
- **Pay per query:** With Amazon Athena, you pay only for the queries that you run. You are charged $5 per terabyte scanned by your queries. You can save from 30% to 90% on your per-query costs and get better performance by compressing, partitioning, and converting your data into columnar formats. Athena queries data directly in Amazon S3. There are no additional storage charges beyond S3.
- **Open, powerful, standard:** Amazon Athena uses Presto with ANSI SQL support and works with a variety of standard data formats, including CSV, JSON, ORC, Avro, and Parquet. Athena is ideal for quick, ad-hoc querying but it can also handle complex analysis, including large joins, window functions, and arrays. Amazon Athena is highly available; and executes queries using compute resources across multiple facilities and multiple devices in each facility. Amazon Athena uses Amazon S3 as its underlying data store, making your data highly available and durable.

- **Interactive performance even for large datasets:** With Amazon Athena, you don't have to worry about having enough compute resources to get fast, interactive query performance. Amazon Athena automatically executes queries in parallel, so most results come back within seconds.

6.2. Backup/Restore and Disaster Recovery

This requirement is not applicable for this service. For additional information beyond what is described herein, please refer to http://aws.amazon.com/documentation/.

6.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

6.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/athena/
- **Service quotas:** https://docs.aws.amazon.com/athena/latest/ug/service-limits.html
- **Service FAQs:** https://aws.amazon.com/athena/faqs/

6.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/athena/ and the following links for comprehensive technical documentation regarding this service.

- **User Guide:** Learn how to use Athena to query data stored in Amazon S3.
- **API Reference:** Describes the Athena API operations in detail.

7. Amazon Augmented AI (A2I)

7.1. Service Overview

Some machine learning applications need human oversight to ensure accuracy with sensitive data, to help provide continuous improvements and, retrain models with updated predictions. However, in these situations, you're often forced to choose between a machine learning only or human only system. Companies are looking for the best of both worlds -- integrating machine learning systems into your workflow while keeping a human eye on the results to guarantee a needed precision.

Amazon Augmented AI is a machine learning service which makes it easy to build the workflows required for human review. Amazon A2I brings human review to all developers, removing the undifferentiated heavy lifting associated with building human review systems or managing large numbers of human reviewers whether it runs on AWS or not.

Please note that procurement of services via AWS Marketplace or Mechanical Turk is not possible via this framework.

7.1.1. Features

- **Easy integration:** Amazon Augmented AI (Amazon A2I) is integrated with Amazon Textract for document processing and Amazon Rekognition for content moderation, so you can implement human review workflows for these use cases with just a few clicks in

the Amazon A2I console or a few API parameters. The Amazon A2I API also allows you to integrate your workflows into custom models that you've built with Amazon SageMaker or other machine learning tools.

- **Flexibility to work with reviewers inside and outside of your organization:** Amazon A2I supports multiple choices for human reviewers. You can use your private team of reviewers for in-house review jobs, especially when handling sensitive data that needs to stay within your organization.
- **Easy instructions for reviewers:** With Amazon A2I, you provide instructional guidance to human reviewers to help ensure consistency. These detailed instructions are available to reviewers within their review interface. You can update these instructions at any time, which makes it easy to add more detail to tasks where reviewers often commit mistakes or to adjust instructions based on evolving needs.
- **Workflows to simplify the human review process:** Amazon A2I provides built-in workflows that route predictions to reviewers and take the reviewers step by step through their tasks. The conditions under which workflows route predictions to reviewers can be either a confidence threshold or a random sampling percentage. You can also build custom workflows by providing an AWS Lambda function that you write to tell Amazon A2I when to trigger human reviews, and a web interface that you create using one of the over 60 available HTML templates or from scratch.
- **Content moderation:** The reviewer interface for content moderation allows you to specify clear instructions to help reviewers complete their tasks.
- **Form extraction:** The reviewer interface for form extraction enables you to extract key-value pairs from document images or online forms. The interface allows you to specify clear instructions to help reviewers complete their tasks.
- **Improve results with multiple reviews:** You can use multiple workers in reviews to increase the confidence level of the results. When defining an Amazon A2I workflow, you can specify the number of workers per review, and Amazon A2I routes each review to that many reviewers.

7.1.2. Benefits

- **Easily implement human review of ML predictions:** Amazon A2I gives you the flexibility to incorporate human review into ML applications based on your specific requirements. Low-confidence predictions are sent to humans to review and take action. If needed, you can also require multiple reviewers to review a prediction to achieve consensus. Additionally, to audit models, you can randomly sample predictions for human review so that you can regularly evaluate if the model is still performing well. Amazon A2I helps people and machines do what they do best.
- **Integrate human oversight with any application:** Amazon A2I provides you an easy way to integrate human oversight into your machine learning workflows, with no machine learning experience required. No need to go with an all human review system vs. machine learning only, Amazon A2I brings together machine learning and humans to provide you with automation while keeping a human eye on the results to provide needed precision. Amazon A2I makes it easy to integrate human judgement and AI into any ML application, regardless of whether it's run on AWS or on another platform.
- **Get to market quicker:** Deciding between machine learning vs. humans doing manual processes can be the decision to get to market today vs. months from now. Integrating

Amazon A2I into your workflow not only aids in getting to market with your machine learning quicker but you can also update and retrain your models over time. As your business needs change, so can your workflows and Amazon A2I can help provide you continuously improve your models at whatever stage you are in your machine learning journey.

7.2. Backup/Restore and Disaster Recovery

This requirement is not applicable for this service. For additional information beyond what is described herein, please refer to http://aws.amazon.com/documentation/.

7.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

7.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/sagemaker/latest/dg/a2i-getting-started.html
- **Service quotas:** https://docs.aws.amazon.com/general/latest/gr/a2i.html
- **Service FAQs:** https://aws.amazon.com/augmented-ai/faqs/

7.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/sagemaker/latest/dg/a2i-getting-started.html for comprehensive technical documentation regarding this service.

8. Amazon Aurora

8.1. Service Overview

Amazon Aurora is a MySQL and PostgreSQL-compatible relational database built for the cloud that combines the performance and availability of traditional enterprise databases with the simplicity and cost-effectiveness of open source databases.

Amazon Aurora is up to five times faster than standard MySQL databases and three times faster than standard PostgreSQL databases. It provides the security, availability, and reliability of commercial databases at 1/10th the cost. Amazon Aurora is fully managed by Amazon Relational Database Service (RDS), which automates time-consuming administration tasks like hardware provisioning, database setup, patching, and backups.

Amazon Aurora features a distributed, fault-tolerant, self-healing storage system that auto-scales up to 128TB per database instance. It delivers high performance and availability with up to 15 low-latency read replicas, point-in-time recovery, continuous backup to Amazon S3, and replication across three Availability Zones.

Visit the Amazon RDS Management Console to create your first Aurora database instance and start migrating your MySQL and PostgreSQL databases.

8.1.1. Features

- **High Performance and Scalability:** Get 5X the throughput of standard MySQL and 3X the throughput of standard PostgreSQL. This performance is on par with commercial

databases, at 1/10th the cost. You can easily scale your database deployment up and down from smaller to larger instance types as your needs change, or let Aurora Serverless handle scaling automatically for you. To scale read capacity and performance, you can add up to 15 low latency read replicas across three Availability Zones. Amazon Aurora automatically grows storage as needed, up to 128TB per database instance.

- **High Availability and Durability:** Amazon Aurora is designed to offer 99.99% availability, replicating 6 copies of your data across 3 Availability Zones and backing up your data continuously to Amazon S3. It transparently recovers from physical storage failures; instance failover typically takes less than 30 seconds. You can also backtrack within seconds to a previous point in time to recover from user errors. With Global Database, a single Aurora database can span multiple AWS Regions to enable fast local reads and quick disaster recovery.
- **Highly Secure:** Amazon Aurora provides multiple levels of security for your database. These include network isolation using Amazon VPC, encryption at rest using keys you create and control through AWS Key Management Service (KMS) and encryption of data in transit using SSL. On an encrypted Amazon Aurora instance, data in the underlying storage is encrypted, as are the automated backups, snapshots, and replicas in the same cluster.

8.1.2. Benefits

- **MySQL and PostgreSQL Compatible:** The Amazon Aurora database engine is fully compatible with existing MySQL and PostgreSQL open source databases, and adds support for new releases regularly. This means you can easily migrate MySQL or PostgreSQL databases to Aurora using standard MySQL or PostgreSQL import/export tools or snapshots. It also means the code, applications, drivers, and tools you already use with your existing databases can be used with Amazon Aurora with little or no change.
- **Fully Managed:** Amazon Aurora is fully managed by Amazon Relational Database Service (RDS). You no longer need to worry about database management tasks such as hardware provisioning, software patching, setup, configuration, or backups. Aurora automatically and continuously monitors and backs up your database to Amazon S3, enabling granular point-in-time recovery. You can monitor database performance using Amazon CloudWatch, Enhanced Monitoring, or Performance Insights, an easy-to-use tool that helps you quickly detect performance problems.
- **Migration Support:** MySQL and PostgreSQL compatibility make Amazon Aurora a compelling target for database migrations to the cloud. If you're migrating from MySQL or PostgreSQL, see our migration documentation for a list of tools and options. To migrate from commercial database engines, you can use the AWS Database Migration Service for a secure migration with minimal downtime.

8.2. Backup/Restore and Disaster Recovery

This service allows for backup and recovery. The service can back up block volumes. Users control this via a manual or scheduled API call. Users schedule and recover backups through a web interface.

8.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

8.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/rds/index.html
- **Service quotas:** https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/CHAP_Limits.html
- **Service FAQs:** https://aws.amazon.com/rds/aurora/faqs/

8.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/rds/index.html and the following link for comprehensive technical documentation regarding this service.

- Amazon Aurora User Guide: Describes all Amazon Aurora concepts and provides instructions on using the various features with both the console and the command line interface.

9. Amazon Braket

9.1. Service Overview

Amazon Braket is a fully managed quantum computing service that helps you get started with quantum computing and accelerate your research. Amazon Braket provides everything you need to build, test, and run quantum algorithms on AWS. It includes access to different types of quantum computers, a unified development environment, a choice of classical circuit simulators, and fully managed execution of hybrid quantum-classical algorithms.

9.1.1. Features

- **Hardware-agnostic developer framework**: To simplify the process of designing and executing quantum algorithms, you can use the Amazon Braket SDK. It has been designed to be technology agnostic, removing the need to code against different quantum programming tools for each type of quantum hardware. The SDK supports a unified developer framework that enables you to build quantum algorithms and run them on any compatible quantum hardware or circuit simulator provided through the Amazon Braket service. As new quantum technologies emerge and are added to the Amazon Braket service, you can be confident that your development experience will remain consistent and that your existing designs and quantum algorithms can be tested on these new systems.
- **Fully managed executions of quantum-classical algorithms with Hybrid Jobs**: Amazon Braket Hybrid Jobs simplifies the process of setting up, monitoring, and running hybrid quantum-classical algorithms. After you provide your algorithm script and select the QPU or simulator to run on, Amazon Braket spins up the classical compute, executes the algorithm, and releases the resources once the job is completed. You can define custom metrics for your algorithms, which are automatically logged by Amazon CloudWatch and displayed in real-time in the Amazon Braket console as the algorithm runs. This gives you live insights into how your algorithm is progressing, so you can make adjustments to your algorithm if needed. Most importantly, Hybrid Jobs provides prioritized access to your chosen QPU to help your algorithm execute quickly and predictably, enabling you to improve the quality and reproducibility of results.

- **Develop variational quantum algorithms with PennyLane**: Amazon Braket natively supports PennyLane, an open source software framework built around the concept of quantum differentiable programming, to help you build and run hybrid quantum-classical, or variational, algorithms. This approach enables you to train quantum circuits in the same way that you would train a machine learning neural network to find solutions to computational problems in quantum chemistry, quantum machine learning, and optimization. PennyLane is performance-optimized for Amazon Braket and provides interfaces to familiar machine learning tools, including PyTorch and TensorFlow, to make training quantum circuits fast, easy, and intuitive.
- **Fully managed Jupyter notebooks**: You have the choice of using your own development environment or fully managed Jupyter notebooks in Amazon Braket to build your quantum algorithms and manage experiments. Amazon Braket makes it easy to create notebooks with a single click. You can select the notebook instance type to match your performance requirements and configure security settings such as encryption for stored data. Amazon Braket notebooks come pre-configured with a suite of quantum computing developer tools, including the Amazon Braket SDK, PennyLane, and Ocean, to help you get started quickly.
- **Pre-built algorithms and tutorials**: Amazon Braket notebooks come pre-installed with the Amazon Braket SDK, tutorials and a selection of pre-built algorithms to give you everything you need to get started in a single place. Use them to familiarize yourself with the recommended steps to build and execute quantum algorithms using Amazon Braket. Learn more in the Amazon Braket documentation.
- **Choice of simulation tools**: With Amazon Braket, you have a choice of four circuit simulators to run and test quantum algorithms. These include the local simulator that is included in the Amazon Braket SDK and three fully managed simulators. The local simulator can run on a laptop or within an Amazon Braket managed notebook and supports simulation of quantum circuits with and without noise. The fully managed simulators are SV1, a general-purpose state vector simulator; DM1, a density matrix simulator that supports noise modeling; and TN1, a tensor network simulator that specializes in certain larger scale structured quantum circuits. Learn more »
- **Consistent experience**: You can run a circuit on Amazon Braket simulators with a single API call. A request to run your algorithm on a simulator works in the same way as a request to run on quantum hardware; by changing a single line of code, you can change from running on a simulator to an actual quantum computer.
- **Choice of result types**: You can choose different result types for your simulation tasks, including individual samples, custom observables, individual amplitudes, or the full state vector. Amazon Braket simulators can calculate exact results, or return measurement samples emulating the behavior of quantum computers.

9.1.2. Benefits

- **Simplified access to quantum computers**: Amazon Braket provides secure access to a variety of quantum computing technologies. There is no upfront commitment or contract to sign, and you pay only for what you use through your AWS bill.
- **Choice of quantum processing units (QPUs)**: Amazon Braket provides access to both annealing and gate-based quantum computers. Following the gate-based quantum computing paradigm, you can access trapped-ion technology from IonQ and superconducting quantum processors from Oxford Quantum Circuits and Rigetti. Alternatively, you can solve quantum annealing problems using the latest QPUs from D-

Wave. This helps you test different technologies, compare the compute performance of different machines for the problem that you are trying to solve, and choose the hardware system that is best suited to your application. Please visit the Hardware Providers page to learn more about the QPUs that are available on Amazon Braket.

- **Amazon Quantum Solutions Lab**: Today, quantum computing is still in its infancy, and there are still many unknowns and challenges to solve. The Amazon Quantum Solutions Lab can help. It is a collaborative research and professional services program staffed with quantum computing experts who can assist you to more effectively explore quantum computing and assess the current performance of this nascent technology. Additionally, you can work with our qualified technology and consulting partners in the AWS Partner Network (APN) that specialize in applications for quantum computing and can help you address your specific requirements. Please visit the Quantum Solutions lab webpage to get started.
- **Management Console**: As a native AWS service, Amazon Braket is accessible through the AWS Management Console, a centralized and easy to use web interface for Amazon Web Services, which provides you with a secure login using your AWS account or AWS Identity and Access Management (IAM) credentials. You can use the console to manage and monitor your Amazon Braket resources, such as notebooks and tasks, and access detailed information about quantum circuit simulators and QPUs.
- **User access management, security, and monitoring**: Amazon Braket is integrated with AWS services such as Amazon CloudWatch, AWS CloudTrail, Amazon EventBridge, and AWS IAM to enable the monitoring of workloads, generate notifications when your tasks are completed, and manage access controls and permissions. Your simulation and quantum task results are delivered to your preferred Amazon Simple Storage Service (S3) bucket for storage and analysis, giving you full control over your data.

9.2. Backup/Restore and Disaster Recovery

This requirement is not applicable for this service. For additional information beyond what is described herein, please refer to http://aws.amazon.com/documentation/.

9.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

9.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/braket/index.html
- **Service quotas:** https://docs.aws.amazon.com/braket/latest/developerguide/braket-quotas.html
- **Service FAQs:** https://aws.amazon.com/braket/faqs/

9.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/braket/index.html and the following links for comprehensive technical documentation regarding this service.

- Developer Guide : Provides a conceptual overview of Amazon Braket, and detailed information about how to design and create quantum tasks, test them on quantum simulators, and then run them on a quantum computer.
- API Reference : The Amazon Braket API Reference provides information about the operations and structures supported in Amazon Braket.
- Amazon Braket Boto 3 SDK : Allows you to interact with Amazon Braket using the AWS SDK for Python (Boto3).
- Braket Python SDK : The Amazon Braket Python SDK is an open source library to design and build quantum circuits, submit them to Amazon Braket devices as quantum tasks, and monitor their execution.

10. Amazon Chime

10.1. Service Overview

Amazon Chime is a communications service that lets you meet, chat, and place business calls inside and outside your organization, all using a single application.

With Amazon Chime, customers can:

- Conduct and attend online meetings with HD video, audio, screen sharing, meeting chat, dial-in numbers, and in-room video conference support;
- Use chat and chat rooms for persistent communications across desktop and mobile devices;
- Administer enterprise users, manage policies, and setup SSO or other advanced features in minutes using the Amazon Chime management console.

Amazon Chime offers an easy-to-use app available for Windows, Mac, web, IOS, Android devices.

For a set of real-time communications components to quickly add messaging, audio, video, and screen sharing capabilities to your web or mobile applications, please visit Amazon Chime SDK.

10.1.1. Features

- **Online Meetings:** Easily schedule, join, and participate in online meetings
- **Video Conferencing:** You can use high-quality wideband audio and high-definition video conferencing for up to 16 people on your desktop, or 8 people on mobile devices. Conference room video systems: Amazon Chime supports most Session Initiation Protocol (SIP) and H.323 video conferencing systems, and meeting participants can join meetings by simply entering the meeting ID into the device console.
- **Team Collaboration:** You can use Amazon Chime to chat with colleagues directly, in a group, or in a chat room. You can share attachments up to 50 megabytes, search contacts, conversations, and chat rooms, and read conversations across all devices. Chat Rooms can be created to enable multiple users to collaborate on projects, share files with colleagues, and use @mentions to direct messages to specific participants. There is no limit to the number of users that can be invited to a chat room and individual users can join as many chat rooms as needed.
- **Business Calling:** You can place and receive phone calls and send and receive text messages directly from the Amazon Chime application. Calls to your phone number ring will ring the Amazon Chime app on all of your devices -- across desktop, mobile, or web.

Just answer on one device and the others will stop ringing. ou can add additional participants to your 1:1 phone call to turn it into a Amazon Chime meeting, regardless of whether they are using the Amazon Chime application or their telephones.

- **Security and Administration:** Amazon Chime is an AWS service, which means you benefit from a data center and network architecture built to meet the requirements of the most security-sensitive organizations. In addition, Amazon Chime features security capabilities built directly into the service. Messages, voice, video, and content are encrypted using AES 256-bit encryption. The Amazon Chime desktop, mobile, and web clients contain logic to keep users on the latest versions providing the most recent features and security patches. You can log, continuously monitor, and record account activity related to actions taken in the Amazon Chime console, using AWS CloudTrail.

10.1.2. Benefits

- **Choose how you communicate:** Amazon Chime lets you choose the communication options that are best suited for your business. You have the option to choose from meetings, chat, and business calling. With Amazon Chime, you have the flexibility to choose the communication option that fits with your business needs, and the freedom to scale up or down as needed.
- **Use one app for all your communication:** Amazon Chime lets you meet, chat, and place business phone calls with a single, secure application. You don't need to switch between applications to collaborate and can instantly go from a chat to a call, share your screen, and even invite more people to join your meeting. When it's time for your meeting, Amazon Chime will call you on all your devices to help ensure you are never late, and that your meetings start on time.
- **Pay only for what you use:** Amazon Chime offers pay-per-use pricing which lets you pay for features you use, on the days that you use them. With pay-per-use pricing there's no upfront investment or long-term contracts. You can switch between Basic features that don't include a charge, and Pro features that do include a charge. You can use the right features for your business without worrying about overspending.

10.2. Backup/Restore and Disaster Recovery

This requirement is not applicable for this service. For additional information beyond what is described herein, please refer to http://aws.amazon.com/documentation/.

10.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

10.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/chime/
- **Service quotas:** https://docs.aws.amazon.com/general/latest/gr/chime.html
- **Service FAQs:** https://aws.amazon.com/chime/faq/

10.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/chime/ and the following links for comprehensive technical documentation regarding this service.

- Administrator Guide: Helps you use Amazon Chime to perform several administrative tasks, such as creating an Amazon Chime account, inviting users, and managing licenses.
- User Guide: Helps you use Amazon Chime, including joining, attending, and scheduling meetings for your organization.
- Amazon Chime section of AWS CLI Reference: Documents the Amazon Chime commands available in the AWS Command Line Interface (AWS CLI).

11. Amazon CloudFront

11.1. Service Overview

Amazon CloudFront is a fast content delivery network (CDN) service that securely delivers data, videos, applications, and APIs to customers globally with low latency, high transfer speeds, all within a developer-friendly environment. Amazon CloudFront is integrated with AWS—both with physical locations that are directly connected to the AWS global infrastructure and with other AWS Cloud services. Amazon CloudFront works seamlessly with services like AWS Shield for distributed denial of service (DDoS) mitigation, Amazon S3, Elastic Load Balancing or Amazon EC2 as origins for your applications, and Lambda@Edge to run custom code closer to your users and to customize the user experience.

11.1.1. Features

- **Network Connectivity and Backbone:** Amazon CloudFront peers with thousands of Tier 1/2/3 telecom carriers globally, is well connected with all major access networks for optimal performance, and has hundreds of terabits of deployed capacity. CloudFront Edge locations are connected to the AWS Regions through the AWS network backbone - fully redundant, multiple 100GbE parallel fiber that circles the globe and links with tens of thousands of networks for improved origin fetches and dynamic content acceleration.
- **Protection against Network and Application Layer Attacks:** With CloudFront as the “front door” to an application and infrastructure, the primary attack surface is moved away from critical content, data, code and infrastructure.
- **SSL/TLS Encryptions and HTTPS:** With Amazon CloudFront, content, APIs or applications can be delivered over HTTPS using the latest version Transport Layer Security (TLSv1.3) to encrypt and secure communication between viewer clients and CloudFront. AWS Certificate Manager (ACM) can be used to easily create a custom SSL certificate and deploy to an CloudFront distribution for free.
- **Real-time Metrics:** Amazon CloudFront is integrated with Amazon CloudWatch, and automatically publishes six operational metrics per distribution, which are displayed in a set of graphs in the CloudFront console. Additional, granular metrics are available with simple click on the console or via API. It also includes standard and real time logginh.
- **Fast Change Propagation and Invalidations:** CloudFront offers fast change propagation and invalidations, within a matter of minutes. Typically, changes are propagated to the edge in a matter of a few minutes, and invalidation times are under two minutes.
- **CloudFront Functions:** Amazon CloudFront offers programmable and secure edge CDN computing capabilities through CloudFront Functions and AWS Lambda@Edge. CloudFront Functions is ideal for high scale and latency sensitive operations like HTTP header manipulations, URL rewrites/redirects, and cache-key normalizations.

- **Origin Shield:** Amazon CloudFront automatically reduces the volume of application origin requests. Content is stored in CloudFront's edge and regional caches and only fetched from origins when needed. The load on application origins can be further reduced by using Origin Shield to enable a centralized caching layer. Origin Shield optimizes cache hit ratios and collapses requests across regions leading to as few as one origin request per object.
- **Enabling redundancy for origins:** CloudFront supports multiple origins for backend architecture redundancy. CloudFront's native origin failover capability automatically serves content from a backup origin when the primary origin is unavailable.

11.1.2. Benefits

- **Latency:** Reduce latency by delivering data through 310+ globally dispersed Points of Presence (PoPs) with automated network mapping and intelligent routing.
- **Security:** Improve security with traffic encryption and access controls, and use AWS Shield Standard to defend against DDoS attacks at no additional charge.
- **Price:** Cut costs with consolidated requests, customizable pricing options, and zero fees for data transfer out from AWS origins.
- **Customization:** Customize the code you run at the AWS content delivery network (CDN) edge using serverless compute features to balance cost, performance, and security.
- **Deliver fast, secure websites:** Reach viewers across the globe in milliseconds with built-in data compression, edge compute capabilities, and field-level encryption.
- **Accelerate dynamic content delivery and APIs:** Optimize dynamic web content delivery with the purpose-built and feature-rich AWS global network infrastructure supporting edge termination and WebSockets.
- **Stream live and on-demand video:** Start streams quickly, play them with consistency, and deliver high-quality video to any device with AWS Media Service and AWS Elemental integration.
- **Distribute patches and updates:** Scale automatically to deliver software, game patches, and IoT over-the-air (OTA) updates at scale with high transfer rates.

11.2. Backup/Restore and Disaster Recovery

This requirement is not applicable for this service. For additional information beyond what is described herein, please refer to http://aws.amazon.com/documentation/.

11.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace."

11.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/cloudfront/
- **Service quotas:** https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/cloudfront-limits.html

- **Service FAQs:** https://aws.amazon.com/cloudfront/faqs/

11.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/cloudfront/ and the following links for comprehensive technical documentation regarding this service.

- **API Reference:** Describes all the API operations for Amazon CloudFront in detail. Also provides sample requests, responses, and errors for the supported web services protocols.
- **Developer Guide:** Provides a conceptual overview of Amazon CloudFront and includes detailed instructions for using the service.

12. Amazon CloudSearch

12.1. Service Overview

Amazon CloudSearch is a fully managed service in the cloud that makes it easy to set up, manage, and scale a search solution for your website. Amazon CloudSearch enables you to search large collections of data such as web pages, document files, forum posts, or product information. With Amazon CloudSearch, you can quickly add search capabilities to your website without having to become a search expert or worry about hardware provisioning, setup, and maintenance. As your volume of data and traffic fluctuates, Amazon CloudSearch automatically scales to meet your needs.

Amazon CloudSearch supports 34 languages and popular search features such as highlighting, autocomplete, and geospatial search.

12.1.1. Features

- **Index and Search:** You can use Amazon CloudSearch to index and search both structured data and plain text.
- **Wide range of search options:** With Amazon CloudSearch you can perform: Full text search with language-specific text processing; Boolean search; Prefix searches; Range searches; Term boosting; Faceting; Highlighting; Autocomplete Suggestions
- **Search result options:** You can get search results in JSON or XML, sort and filter results based on field values, and sort results alphabetically, numerically, or according to custom expressions.

12.1.2. Benefits

- **Simple:** You can configure and manage an Amazon CloudSearch domain through the AWS Management Console, AWS CLI, and AWS SDKs. Simply point to a sample of your data and Amazon CloudSearch will automatically recommend how to configure your domain's indexing options. You can easily add or delete index fields and customize search options such as faceting and highlighting. Configuration changes do not require you to re-upload your data.
- **Scalable:** Amazon CloudSearch offers powerful autoscaling for all search domains. As your data or query volume changes, Amazon CloudSearch can scale your search domain's resources up or down as needed. You can control scaling if you know that you need more capacity for bulk uploads or are expecting a surge in search traffic.

- **Reliable:** Amazon CloudSearch provides automatic monitoring and recovery for your search domains. When Multi-AZ is enabled, Amazon CloudSearch provisions and maintains resources for a search domain in two Availability Zones to ensure high availability. Updates are automatically applied to the search instances in both Availability Zones. Search traffic is distributed across both Availability Zones and the instances in either zone are capable of handling the full load in the event of a failure.
- **High Performance:** Amazon CloudSearch ensures low latency and high throughput performance, even at large scale through automatic sharding and horizontal and vertical autoscaling.
- **Fully Managed:** Amazon CloudSearch is a fully managed custom search service. Hardware and software provisioning, setup and configuration, software patching, data partitioning, node monitoring, scaling, and data durability are handled for you.

12.2. Backup/Restore and Disaster Recovery

This requirement is not applicable for this service. For additional information beyond what is described herein, please refer to http://aws.amazon.com/documentation/.

12.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

12.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/cloudsearch/
- **Service quotas:** https://docs.aws.amazon.com/general/latest/gr/cloudsearch.html
- **Service FAQs:** https://aws.amazon.com/cloudsearch/faqs/

12.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/cloudsearch/ for comprehensive technical documentation regarding this service.

13. Amazon CloudWatch

13.1. Service Overview

Amazon CloudWatch is a monitoring and observability service built for DevOps engineers, developers, site reliability engineers (SREs), IT managers, and product owners. CloudWatch provides you with data and actionable insights to monitor your applications, respond to system-wide performance changes, and optimize resource utilization. CloudWatch collects monitoring and operational data in the form of logs, metrics, and events. You get a unified view of operational health and gain complete visibility of your AWS resources, applications, and services running on AWS and on-premises. You can use CloudWatch to detect anomalous behavior in your environments, set alarms, visualize logs and metrics side by side, take automated actions, troubleshoot issues, and discover insights to keep your applications running smoothly.

13.1.1. Features

- **Easily collect and store logs:** The Amazon CloudWatch Logs service allows you to collect and store logs from your resources, applications, and services in near real time.
- **Collect and aggregate infrastructure and application metrics:** Amazon CloudWatch allows you to collect infrastructure metrics from more than 70 AWS services, such as Amazon Elastic Compute Cloud (Amazon EC2), Amazon DynamoDB, Amazon Simple Storage Service (Amazon S3), Amazon ECS, AWS Lambda, and Amazon API Gateway, with no action on your part.
- **Unified operational view with dashboards:** Amazon CloudWatch dashboards enable you to create reusable graphs and visualize your cloud resources and applications in a unified view. You can graph metrics and logs data side by side in a single dashboard to quickly get the context and move from diagnosing the problem to understanding the root cause.
- **Composite alarms:** With Amazon CloudWatch composite alarms, you can combine multiple alarms and reduce alarm noise. If an issue affects several resources in an application, you will receive a single alarm notification for the entire application instead of one for each affected resource. This helps you focus on finding the root cause of operational issues to reduce application downtime.
- **Logs and metrics correlation:** Applications and infrastructure resources generate large amounts of operational and monitoring data in the form of logs and metrics. In addition to letting you access and visualize these datasets in a single platform, Amazon CloudWatch also makes it easy to correlate them. This helps you quickly move from diagnosing the problem to understanding the root cause.
- **Application Insights:** Amazon CloudWatch Application Insights provides automated setup of observability for your enterprise applications so you can get visibility into their health. It helps you identify and set up key metrics and logs across your application resources and technology stack, such as database, web (IIS) and application servers, operating system, load balancers, and queues. It constantly monitors this telemetry data to detect and correlate anomalies and errors to notify you of any problems in your application.
- **ServiceLens:** You can use Amazon CloudWatch ServiceLens to visualize and analyze the health, performance, and availability of your applications in a single place. It ties together CloudWatch metrics and logs as well as traces from AWS X-Ray to give you a complete view of your applications and their dependencies. Quickly pinpoint performance bottlenecks, isolate root causes of application issues, and determine the impact on users.
- **Automate response to operational changes with CloudWatch Events:** CloudWatch Events provides a near real-time stream of system events that describe changes to your AWS resources. It allows you to respond quickly to operational changes and take corrective action. You simply write rules to indicate which events are of interest to your application and what automated actions to take when a rule matches an event.
- **Granular data and extended retention:** Amazon CloudWatch allows you to monitor trends and seasonality with 15 months of metric data (storage and retention). This lets you perform historical analysis to fine-tune resource utilization. With CloudWatch, you can also collect up to one second of health metrics, including custom metrics (such as those coming from your on-premises applications). Granular real-time data enables

better visualization and the ability to spot and monitor trends to optimize application performance and operational health.

- **Custom operations on metrics:** Amazon CloudWatch Metric Math enables you to perform calculations across multiple metrics for real-time analysis so you can easily derive insights from your existing CloudWatch metrics and better understand the operational health and performance of your infrastructure. You can visualize these computed metrics in the AWS Management Console, add them to CloudWatch dashboards, or retrieve them using the GetMetricData API action. Metric Math supports arithmetic operations (such as +, -, /, and *) and mathematical functions (such as Sum, Average, Min, Max, and Standard Deviation).

13.1.2. Benefits

- **Use a single platform for observability:** Modern applications, such as those running on microservices architectures, generate large volumes of data in the form of metrics, logs, and events. Amazon CloudWatch allows you to collect, access, and correlate this data on a single platform from across all your AWS resources, applications, and services running on AWS and on-premises, helping you break down data silos to gain system-wide visibility and quickly resolve issues.
- **Collect metrics on AWS and on premises:** Monitoring your AWS resources and applications is easy with CloudWatch. It natively integrates with more than 70 AWS services, such as Amazon EC2, Amazon DynamoDB, Amazon S3, Amazon ECS, Amazon EKS, and AWS Lambda. It automatically publishes detailed one-minute metrics and custom metrics with up to one-second granularity so you can dive deep into your logs for additional context. You can also use CloudWatch in hybrid environments by using the CloudWatch Agent or API to monitor your on-premises resources.
- **Improve operational performance and resource optimization:** Set alarms and automate actions based on predefined thresholds or on machine learning (ML) algorithms that identify anomalous behavior in your metrics. For example, you can start Amazon EC2 Auto Scaling automatically or stop an instance to reduce billing overages. You can also use CloudWatch Events for serverless to trigger workflows with services like AWS Lambda, Amazon SNS, and AWS CloudFormation.
- **Get operational visibility and insight:** To optimize performance and resource utilization, you need a unified operational view, real-time granular data, and historical reference. CloudWatch provides automatic dashboards, data with one-second granularity, and up to 15 months of metrics storage and retention. You can also perform metric math on your data to derive operational and utilization insights; for example, you can aggregate usage across an entire fleet of EC2 instances.
- **Derive actionable insights from logs:** Explore, analyze, and visualize your logs so you can troubleshoot operational problems with ease. With CloudWatch Logs Insights, you pay only for the queries you run. It scales with your log volume and query complexity, giving you answers in seconds. In addition, you can publish log-based metrics, create alarms, and correlate logs and metrics together in CloudWatch Dashboards for complete operational visibility.

13.2. Backup/Restore and Disaster Recovery

This service allows for backup and recovery. The service can back up logs to S3. Users control this via a manual or scheduled API call. Users schedule and recover backups through a web interface.

13.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

13.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/cloudwatch/index.html
- **Service quotas:** https://docs.aws.amazon.com/AmazonCloudWatch/latest/monitoring/cloudwatch_limits.html
- **Service FAQs:** https://aws.amazon.com/cloudwatch/faqs/

13.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/cloudwatch/index.html and the following links for comprehensive technical documentation regarding this service.

- CloudWatch User Guide: Provides a conceptual overview of CloudWatch and includes detailed development instructions for using the various features.
- CloudWatch API Reference: Describes all the API operations for CloudWatch in detail. Also provides sample requests, responses, and errors for the supported web services protocols.
- CloudWatch Logs User Guide: Provides a conceptual overview of CloudWatch Logs and includes detailed development instructions for using the various features.

14. Amazon CodeGuru

14.1. Service Overview

Amazon CodeGuru is a developer tool that provides intelligent recommendations to improve code quality and identify an application’s most expensive lines of code. Integrate CodeGuru into your existing software development workflow to automate code reviews during application development and continuously monitor application's performance in production and provide recommendations and visual clues on how to improve code quality, application performance, and reduce overall cost.

CodeGuru Reviewer uses machine learning and automated reasoning to identify critical issues, security vulnerabilities, and hard-to-find bugs during application development and provides recommendations to improve code quality. Learn more »

CodeGuru Profiler helps developers find an application’s most expensive lines of code by helping them understand the runtime behaviour of their applications, identify and remove code inefficiencies, improve performance, and significantly decrease compute costs. Learn more »

14.1.1. Features

Copy and paste features from service features page

- **Amazon CodeGuru Reviewer:** Amazon CodeGuru Reviewer finds issues in your Java and Python code and provides recommendations to improve your code. For example, CodeGuru Reviewer detects security vulnerabilities, secrets, resource leaks, concurrency issues, incorrect input validation, and deviation from best practices for using

AWS APIs and SDKs. To begin reviewing code, you can associate your existing code repositories on GitHub, GitHub Enterprise, Bitbucket, or AWS CodeCommit with CodeGuru.

- **Security detection:** CodeGuru Reviewer helps you improve code security and provides recommendations based on common vulnerabilities (OWASP Top 10) and AWS internal security best practices. It uses automated reasoning to analyze data flow from source to sink and across multiple functions to detect hard-to-find security vulnerabilities.
- **Secrets detection:** CodeGuru Reviewer Secrets Detector uses machine learning-based analysis to help you detect secrets that are hardcoded in your repository or configuration files, including passwords, API keys, SSH keys, access tokens, database connection strings and JSON Web Tokens. Part of CodeGuru Reviewer, Secrets Detector is an automated mechanism that checks code for these secrets and provides point-and-click steps to secure them using AWS Secrets Manager. It can also identify specific keys generated by the most common API providers, including AWS, Atlassian, GitHub, Salesforce, HubSpot, and Stripe.
- **Code quality:** CodeGuru Reviewer identifies code quality issues and equips your development team to maintain a high bar of coding standards in the software development process
- **CI/CD integration with GitHub Actions:** CI/CD experience for CodeGuru Reviewer allow you to invoke code quality and security analysis as a step within your CI workflow using GitHub Actions. You can configure it to run and provide recommendations on a pull, push, or scheduled run of your pipeline. After you invoke a CodeGuru Reviewer scan via CI/CD, you can view your code quality and security recommendations within the CodeGuru Reviewer Console or within the GitHub's user interface. With CI/CD integration, you can continuously monitor the quality and security of your code to help ensure that you do not miss a recommendation.
- **Amazon CodeGuru Profiler:** Amazon CodeGuru Profiler is always searching for application performance optimizations, identifying your most "expensive" lines of code and recommending ways to fix them to reduce CPU utilization, cut compute costs, and improve application performance. For example, CodeGuru Profiler can identify when your application is consuming excessive CPU capacity on a logging routine instead of executing on core business logic.
- **Always-on profiling of applications in production:** CodeGuru Profiler is designed to continuously run on production with minimal overhead which means you can leave it on all the time with minimal impact on application performance. It enables you to profile and troubleshoot your application using real customer traffic patterns and easily discover performance issues. With the profiler data and ML-powered recommendations, you can identify and fix performance issues for your applications in production. CodeGuru Profiler also provides a heap summary, so you can identify what objects are using up memory at any given time.
- **Understand the runtime behavior of applications:** CodeGuru Profiler continuously analyzes application CPU utilization, heap usage, and latency characteristics to show you where you are spending the most cycles or time in your application. The CPU and latency analysis is presented in an interactive flame graph that helps you easily understand which code paths consume the most resources, verify that your application is performing as expected, and uncover areas that can be optimized further.
- **Intelligent recommendations:** CodeGuru Profiler automatically identifies performance issues in your application and provides ML-powered recommendations on how to

remediate them. These recommendations help you identify and optimize the most expensive or resource intensive methods within your code without requiring you to be a performance engineering expert. These optimizations help you reduce the cost of your infrastructure, reduce latency, and improve your overall end user experience.

- **Anomaly detection:** Amazon CodeGuru Profiler continuously analyzes your application profiles in real-time and detects anomalies in the behavior of your application and its methods. Each anomaly is tracked in the Recommendation report of the CodeGuru Profiler console and you can see time series of how the method's latency behaves over time with anomalies clearly highlighted. If configured, an Amazon SNS notification will also be sent when a new anomaly is detected.

14.1.2. Benefits

- **Catch code problems before they hit production:** For code reviews, developers commit their code to GitHub, GitHub Enterprise, Bitbucket Cloud, and AWS CodeCommit and add CodeGuru Reviewer as one of the code reviewers, with no other changes to their development process. CodeGuru Reviewer analyzes existing code bases in the repository, identifies hard-to-find bugs and critical issues with high accuracy, provides intelligent suggestions on how to remediate them, and creates a baseline for successive code reviews.
- **Fix security vulnerabilities:** CodeGuru Reviewer Security Detector leverages automated reasoning and AWS's years of security experience to improve your code security. It enables you to incorporate security reviews directly into your application development CI/CD processes via a GitHub Action and ensures that your code follows best practices for AWS Key Management Service (AWS KMS), Amazon Elastic Cloud Compute (Amazon EC2), application programming interfaces (APIs), common Java or Python crypto, and TLS (Transport Layer Security)/SSL (Secure Socket Layer) libraries. When the security detector discovers an issue, a recommendation for remediation is provided along with an explanation for why the code improvement is suggested, thereby enabling Security Engineers to focus on architectural and application-specific security best practices.
- **Proactively improve code quality with continuous monitoring:** For every pull request initiated, CodeGuru Reviewer automatically analyzes the incremental code changes and posts recommendations directly on the pull request. Additionally, it supports full repository or code base scan for periodic code maintainability, and code due diligence initiatives to ensure that your code quality is consistent. CodeGuru Reviewer can also be integrated into your CI/CD pipelines. You can configure it to run on a pull, push, or scheduled run of your pipeline and view your code quality and security recommendations within the AWS Console or within your CI/CD provider's user interface.

14.2. Backup/Restore and Disaster Recovery

This requirement is not applicable for this service. For additional information beyond what is described herein, please refer to http://aws.amazon.com/documentation/.

14.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

14.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/codeguru/index.html
- **Service quotas:** https://docs.aws.amazon.com/codeguru/latest/profiler-ug/quotas.html for CodeGuru Profiler and https://docs.aws.amazon.com/codeguru/latest/reviewer-ug/quotas.html for CodeGuru Reviewer
- **Service FAQs:** https://aws.amazon.com/codeguru/faqs/

14.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/codeguru/index.html and the following links for comprehensive technical documentation regarding this service.

- User Guide for CodeGuru: Provides a conceptual overview of Amazon CodeGuru Reviewer, instructions for getting started, product and service integrations, and review recommendations.
- API Reference for CodeGuru: Describes all the Amazon CodeGuru Reviewer API operations in detail. Also provides sample requests, responses, and errors for the supported web services protocols.
- CodeGuru Reviewer section of the AWS CLI Reference: Describes the AWS CLI commands that you can use to work with CodeGuru Reviewer repository associations and code reviews.
- User Guide for Profiler: Provides a conceptual overview of Amazon CodeGuru Profiler, instructions for creating profile groups, using graphs and filters, reviewing recommendations, and exploring code.
- API Reference for Profiler: Describes all the Amazon CodeGuru Profiler API operations in detail. Also provides sample requests, responses, and errors for the supported web services protocols.
- CodeGuru Profiler section of the AWS CLI Reference: Describes the AWS CLI commands that you can use to work with profiling groups.

15. Amazon Cognito

15.1. Service Overview

Amazon Cognito lets you add user sign-up, sign-in, and access control to your web and mobile apps quickly and easily. Amazon Cognito scales to millions of users and supports sign-in with social identity providers, such as Apple, Facebook, Google, and Amazon, and enterprise identity providers via SAML 2.0 and OpenID Connect.

Please note, Amazon Cognito Sync is not available for procurement via this framework.

15.1.1. Features

- **An identity store for all your apps and users:** Amazon Cognito User Pools provide a secure identiy store that scales to millions of users. As a fully managed service, User Pools are easy to set up without provisioning any infrastructure. User Pools store user profiles and support authentication for users who sign up directly and for federated users who sign in with social and enterprise identity providers.

- **Built-in customizable UI to sign in users:** Amazon Cognito provides a built-in and customizable UI for user sign-up and sign-in. You can use Android, iOS, and JavaScript SDKs for Amazon Cognito to add user sign-up and sign-in pages to your apps.
- **Advanced security features to protect your users:** Using advanced security features for Amazon Cognito helps you protect access to user accounts in your applications. These advanced security features provide risk-based adaptive authentication and protection from the use of compromised credentials. With just a few clicks, you can enable these advanced security features for your Amazon Cognito User Pools.
- **Social and enterprise identity federation**: With Amazon Cognito, your users can sign-in through social identity providers such as Google, Facebook, and Amazon, and through enterprise identity providers such as Microsoft Active Directory using SAML.
- **Access control for AWS resources**: Amazon Cognito provides solutions to control access to AWS resources from your app. You can define roles and map users to different roles so your app can access only the resources that are authorized for each user. Alternatively, you can use attributes from identity providers in AWS Identity and Access Management permission policies, so you can control access to resources to users who meet specific attribute conditions.
- **Standards-based authentication**: Amazon Cognito uses common identity management standards including OpenID Connect, OAuth 2.0, and SAML 2.0.
- **Adaptive authentication**: Using advanced security features for Amazon Cognito to add adaptive authentication to your applications helps protect your applications' user accounts and user experience. When Amazon Cognito detects unusual sign-in activity, such as sign-in attempts from new locations and devices, it assigns a risk score to the activity and lets you choose to either prompt users for additional verification or block the sign-in request. Users can verify their identities using SMS or a Time-based One-time Password (TOTP) generator, such as Google Authenticator.
- **Protection from compromised credentials**: Advanced security features for Amazon Cognito helps protect your application users from unauthorized access to their accounts using compromised credentials. When Amazon Cognito detects users have entered credentials that have been compromised elsewhere, it prompts them to change their password.
- **Supports Multiple Compliance Programs**: Amazon Cognito helps you meet multiple security and compliance requirements, including those for highly regulated organizations such as healthcare companies and merchants. Amazon Cognito is HIPAA eligible and PCI DSS, SOC, and ISO/IEC 27001, ISO/IEC 27017, ISO/IEC 27018, and ISO 9001 compliant.

15.1.2. Benefits

- **Social and enterprise identity federation**: With Amazon Cognito, your users can sign in through social identity providers such as Apple, Google, Facebook, and Amazon, and through enterprise identity providers such as SAML and OpenID Connect.
- **Secure and scalable identity store**: Amazon Cognito User Pools provide a secure identity store that scales to millions of users. Cognito User Pools can be more easily set up without provisioning any infrastructure, and all members of the user pool have a directory profile that you can manage through a Software Development Kit (SDK).
- **Security for your apps and users**: Amazon Cognito supports multi-factor authentication and encryption of data-at-rest and in-transit. Amazon Cognito is HIPAA

eligible and PCI DSS, SOC, ISO/IEC 27001, ISO/IEC 27017, ISO/IEC 27018, and ISO 9001 compliant.

- **Access control for AWS resources**: Amazon Cognito provides solutions to control access to AWS resources from your app. You can define roles and map users to different roles so your app can access only the resources that are authorized for each user. Alternatively, you can use attributes from identity providers in AWS Identity and Access Management permission policies, so you can control access to resources to users who meet specific attribute conditions.
- **Easy integration with your app**: With a built-in UI and easy configuration for federating identity providers, you can integrate Amazon Cognito to add user sign-in, sign-up, and access control to your app in minutes. You can customize the UI to put your company branding front and center for all user interactions.

15.2. Backup/Restore and Disaster Recovery

This requirement is not applicable for this service. For additional information beyond what is described herein, please refer to http://aws.amazon.com/documentation/.

15.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

15.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/cognito/
- **Service quotas:** https://docs.aws.amazon.com/cognito/latest/developerguide/limits.html
- **Service FAQs:** https://aws.amazon.com/cognito/faqs/

15.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/cognito/ and the following links for comprehensive technical documentation regarding this service.

- **Developer Guide:** Provides a conceptual overview of Amazon Cognito user pools and includes instructions that show you how to use its features.
- **API Reference:** Describes the REST API for user pools.
- **AWS CLI Reference**: Describes the AWS CLI commands for user pools.

16. Amazon Comprehend

16.1. Service Overview

Amazon Comprehend is a natural language processing (NLP) service that uses machine learning to discover insights from text. Amazon Comprehend provides Keyphrase Extraction, Sentiment Analysis, Entity Recognition, Topic Modelling, and Language Detection APIs so you can easily integrate natural language processing into your applications. You simply call the Amazon Comprehend APIs in your application and provide the location of the source document or text. The APIs will output entities, key phrases, sentiment, and language in a JSON format, which you can use in your application.

16.1.1. Features

- **Keyphrase Extraction:** The Keyphrase Extraction API returns the key phrases or talking points and a confidence score to support that this is a key phrase.
- **Sentiment Analysis:** The Sentiment Analysis API returns the overall sentiment of a text (Positive, Negative, Neutral, or Mixed).
- **Syntax Analysis:** The Amazon Comprehend Syntax API enables customers to analyse text using tokenization and Parts of Speech (PoS), and identify word boundaries and labels like nouns and adjectives within the text.
- **Entity Recognition:** The Entity Recognition API returns the named entities ("People," "Places," "Locations," etc.) that are automatically categorized based on the provided text.
- **Custom Entities:** Custom Entity Recognition allows you to customize Amazon Comprehend to identify terms that are specific to your domain. Using AutoML, Comprehend will learn from a small set of examples (for example, a list of policy numbers, claim numbers, or SSN), and then train a private, custom model to recognize these terms such as claim numbers in any other block of text within PDFs, plain text, or Microsoft Word documents – no machine learning required.
- **Custom Classification:** The Custom Classification API enables you to easily build custom text classification models using your business-specific labels without learning ML. For example, your customer support organization can use Custom Classification to automatically categorize inbound requests by problem type based on how the customer has described the issue.
- **Topic Modelling:** Topic Modelling identifies relevant terms or topics from a collection of documents stored in Amazon S3. It will identify the most common topics in the collection and organize them in groups and then map which documents belong to which topic.
- **Multiple language support:** Amazon Comprehend can perform text analysis on English, French, German, Italian, Portuguese, and Spanish texts. This lets you build applications that can detect text in multiple languages, convert the text to English, French, German, Italian, Portuguese, and Spanish with Amazon Translate, and then use Amazon Comprehend to perform text analysis.

16.1.2. Benefits

- **Uncover valuable insights:** from text in documents, customer support tickets, product reviews, emails, social media feeds, and more.
- **Simplify document processing workflows:** extract text, key phrases, topics, sentiment, and more from documents such as insurance claims.
- **Differentiate your business:** train a model to classify documents and identify terms, with no machine learning experience required.
- **Protect and control who has access to your sensitive data:** identify and redact Personally Identifiable Information (PII) from documents.

16.2. Backup/Restore and Disaster Recovery

This requirement is not applicable for this service. For additional information beyond what is described herein, please refer to http://aws.amazon.com/documentation/.

16.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

16.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/comprehend/
- **Service quotas:** https://docs.aws.amazon.com/comprehend/latest/dg/guidelines-and-limits.html
- **Service FAQs:** https://aws.amazon.com/comprehend/faqs/

16.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/comprehend/ and the following links for comprehensive technical documentation regarding this service.

- Developer Guide: Provides a conceptual overview of Amazon Comprehend and detailed instructions for using its features.
- API Reference: Describes all the API operations for Amazon Comprehend in detail.

17. Amazon Comprehend Medical

17.1. Service Overview

Amazon Comprehend Medical is a HIPAA-eligible natural language processing (NLP) service that uses machine learning to extract health data from medical text–no machine learning experience is required.

With a simple API call to Amazon Comprehend Medical you can quickly and accurately extract information such as medical conditions, medications, dosages, tests, treatments and procedures, and protected health information while retaining the context of the information. Amazon Comprehend Medical can identify the relationships among the extracted information to help you build applications for use cases like population health analytics, clinical trial management, pharmacovigilance, and summarization. You can also use Amazon Comprehend Medical to link the extracted information to medical ontologies such as ICD10-CM or RxNorm to help you build applications for use cases like revenue cycle management (medical coding), claim validation and processing, and electronic health record creation.

Amazon Comprehend Medical is fully managed, so there are no servers to provision, and no machine learning models to build, train, or deploy. You pay only for what you use, and there are no minimum fees and no upfront commitments.

17.1.1. Features

- **Medical Named Entity and Relationship Extraction (NERe):** The Medical NERe API returns the medical information such as medication, medical condition, test, treatment and procedures (TTP), anatomy, and Protected Health Information (PHI). It also identifies relationships between extracted sub-types associated to Medications and TTP. There is also contextual information provided as entity “traits” (negation, or if a diagnosis is a sign or symptom). The table below shows the extracted information with relevant

sub-types and entity traits. To only extract PHI, you can use the Protected Health Information Data Identification (PHId) API.

- **Medical Ontology Linking:** The Medical Ontology Linking APIs identifies medical information and links them to codes and concepts in standard medical ontologies. Medical conditions are linked to ICD-10-CM codes (e.g. “headache” is linked to the “R51” code) with the InferICD10CM API, while medications are linked to RxNorm codes (“Acetaminophine / Codeine” is linked to the “C2341132” cui). The Medical Ontology Linking APIs also detects contextual information as entity traits (e.g. negation).
- **Keyphrase Extraction:** The Keyphrase Extraction API returns the key phrases or talking points and a confidence score to support that this is a key phrase.
- **Sentiment Analysis:** The Sentiment Analysis API returns the overall sentiment of a text (Positive, Negative, Neutral, or Mixed).
- **Syntax Analysis:** The Amazon Comprehend Syntax API enables customers to analyse text using tokenization and Parts of Speech (PoS), and identify word boundaries and labels like nouns and adjectives within the text.
- **Entity Recognition:** The Entity Recognition API returns the named entities ("People," "Places," "Locations," etc.) that are automatically categorized based on the provided text.
- **Custom Entities:** Custom Entity Recognition allows you to customize Amazon Comprehend to identify terms that are specific to your domain. Using AutoML, Comprehend will learn from a small set of examples (for example, a list of policy numbers, claim numbers, or SSN), and then train a private, custom model to recognize these terms such as claim numbers in any other block of text within PDFs, plain text, or Microsoft Word documents – no machine learning required.
- **Custom Classification:** The Custom Classification API enables you to easily build custom text classification models using your business-specific labels without learning ML. For example, your customer support organization can use Custom Classification to automatically categorize inbound requests by problem type based on how the customer has described the issue.
- **Topic Modelling:** Topic Modelling identifies relevant terms or topics from a collection of documents stored in Amazon S3. It will identify the most common topics in the collection and organize them in groups and then map which documents belong to which topic.
- **Multiple language support:** Amazon Comprehend can perform text analysis on English, French, German, Italian, Portuguese, and Spanish texts. This lets you build applications that can detect text in multiple languages, convert the text to English, French, German, Italian, Portuguese, and Spanish with Amazon Translate, and then use Amazon Comprehend to perform text analysis.

17.1.2. Benefits

- **Extract medical information quickly and accurately:** Powered by state-of-the-art machine learning models, Amazon Comprehend Medical understands and identifies complex medical information quickly and more accurately. For example, Amazon Comprehend Medical can extract "methicillin-resistant Staphylococcus aureus" (often input as "MRSA") link it to the "J15.212" ICD-10-CM code, and provide context, such as

whether a patient has tested positive or negative, to make the extracted term meaningful.

- **Protect patient information:** Amazon Comprehend Medical provides a number of capabilities to help healthcare providers stay compliant and protect patient data. The service is HIPAA eligible and can identify protected health information (PHI) stored in medical record systems while adhering to the standards for General Data Protection Regulation (GDPR). Amazon Comprehend Medical allows developers to implement data privacy and security solutions by extracting and then identifying relevant patient identifiers as described in HIPAA's Safe Harbor method of de-identification. Finally, the service does not store or save any customer data.
- **Lower medical document processing costs:** Amazon Comprehend Medical makes it easy to automate and lower the cost of processing and coding unstructured medical text from patient records, billing, and clinical indexing. It offers 2 APIs that developers can integrate into existing workflows and applications with only a few lines of code, costing a penny or less for every 100 characters of analysed text. You pay only for what you use, and there are no minimum fees.

17.2. Backup/Restore and Disaster Recovery

This requirement is not applicable for this service. For additional information beyond what is described herein, please refer to http://aws.amazon.com/documentation/.

17.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

17.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/comprehend-medical/latest/dev/comprehendmedical-welcome.html
- **Service quotas:** https://docs.aws.amazon.com/comprehend-medical/latest/dev/comprehendmedical-quotas.html
- **Service FAQs:** https://aws.amazon.com/comprehend/faqs/

17.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/comprehend-medical/latest/dev/comprehendmedical-welcome.html

18. Amazon Connect

18.1. Service Overview

Amazon Connect is a self-service, cloud-based contact centre service that makes it easy for any business to deliver better customer service at lower cost. Amazon Connect is based on the same contact centre technology used by Amazon customer service associates around the world to power millions of customer conversations. The self-service graphical interface in Amazon Connect makes it easy for non-technical users to design contact flows, manage agents, and track performance metrics—no specialised skills required. There are no up-front payments or

long-term commitments and no infrastructure to manage with Amazon Connect; customers pay by the minute for Amazon Connect usage plus any associated telephony services.

18.1.1. Features

- **Telephony:** Amazon Connect manages a network of telephony providers from around the world, removing the need for you to manage multiple vendors, negotiate complex multi-year contracts, or commit to peak call volumes. Its telephony service includes direct inward dial (DID) and toll-free phone numbers in 20+ countries worldwide. There are also 200+ outbound calling destinations available. The telephony-as-a service model also offers proactive monitoring from telephony experts, which can scale up and down at a moment's notice. Best of all, you only pay for what you use with pay-as-you-go pricing.
- **High-quality audio:** Sound quality in a call impacts productivity. When your customers can't hear you clearly, it can lead to wasted time and frustration. With Amazon Connect, calls are made over the internet from a computing device like a PC, using the Amazon Connect softphone. The Amazon Connect softphone delivers high-quality 16kHz audio and is resistant to packet loss to ensure a high-quality call experience.
- **Web and mobile chat:** With Amazon Lex natively integrated within Amazon Connect, no coding is required for adding chatbots that have natural language understanding (NLU) and for the context of conversations to be passed automatically when escalated to a human agent. Amazon Connect Chat supports asynchronous messaging, so your customers and agents have the ability to message without being available at the same time. Chats are secure and encrypted and support all existing Amazon Connect compliance certifications. Amazon Connect also offers native integration with Apple Business Chat, enabling your agents to support customers through the same popular iOS Messaging application that they use every day. Management for Apple Business Chat is easy because it uses Amazon Connect contact flows, configuration, and routing.
- **Omnichannel routing:** Amazon Connect has a single user interface (UI) across voice, chat, and tasks for contact routing, queuing, analytics, and management. This omnichannel experience means your call centre agents don't have to learn and work across multiple tools. The same automated interactions and chatbots can be used across both channels, increasing operational efficiency because you don't have to rebuild interaction flows. Customers can interact with your agents on voice or chat based on factors such as personal preferences and wait times. Customers can keep working with the same agent across channels, but if it's a different agent, their interaction history is preserved, so they don't have to repeat themselves. The omnichannel contact centre improves customer experience while reducing resolution time. With Amazon Connect, you can build call flows, rules, and reports once and enable them across channels.
- **Task management:** Follow-up items to resolve customer issues and requests are essential to maintaining high customer satisfaction. Amazon Connect Tasks makes it easy for you to prioritize, assign, and track agent tasks to completion, including work in external applications to ensure customer issues are quickly resolved. Today, agents who keep track of tasks and follow-up items for customers manually know that it's difficult and error prone, especially when a task spans multiple systems. Amazon Connect Tasks allows your agents to create and complete tasks just like they do a call or chat. You can also use workflows to automate tasks that don't require agent interaction. This results in improved agent productivity leading to increased customer satisfaction.

- **Contact centre automation:** Amazon Connect works on behalf of your supervisors and agents to save time and money while providing your customers with the best possible contact centre experience. Amazon Connect has self-service tools and intelligent automation, such as natural language chatbots, interactive voice response (IVR), and automated customer voice authentication. Amazon Connect provides a seamless omnichannel experience for agents and managers for voice, chat, and task management to ensure customers are routed with their conversation's full context or work across applications when switching channels. Once Amazon Connect has identified the customer's needs, skills-based routing matches them with the best available agent most likely to resolve their issue. Amazon Connect has all the automation capabilities you need to improve efficiency and reduce agent time performing repetitive functions.
- **Rules engine:** Automatically define filers based on information such as specific keywords and phrases extracted from a customer-agent conversation using real-time or post-call rules in Contact Lens for Amazon Connect. For example, you can set up a rule to alert QA managers when a VIP customer mentions "cancel my subscription," or to inform the sales team when a new customer mentions "I'd like to upgrade my account," assigning and routing an Amazon Connect task to the next available agent.
- **Agent application:** Amazon Connect's agent application consolidates all agent features into one easy to use experience, helping to save your agents valuable time and increasing their productivity. It combines the Contact Control Panel (CCP) with other Amazon Connect agent capabilities, such as task management, customer information, knowledge assist, and workforce schedules into a single UI.
- **Contact Control Panel:** The Contact Control Panel (CCP) provides a single, easy to use communication interface for agents to receive calls, chat with contacts, transfer them to other agents, put them on hold, and perform other tasks. It is also easy to customize your agent experience allowing you to integrate your external applications like CRM or marketing automation.
- **Skills-based routing:** Amazon Connect has a single UI and routing engine for calls and chat, increasing efficiency among agents. Efficient routing is important to minimize wait times and ensure an end customer gets the answer they need. With skills-based routing, Amazon Connect ensures contacts are sent to the right agent at the right time based on variables such as availability, skillset, customer sentiment, and past history. This helps agents quickly and efficiently resolve issues.
- **Unified customer profiles:** Amazon Connect Customer Profiles brings together information from multiple applications into a unified customer profile to empower automated interactions and help your agents improve customer service. It aggregates customer data with built-in connectors for third-party applications like Salesforce, ServiceNow, Zendesk, and Marketo. When a customer contacts the contact centre, Amazon Connect Customer Profiles scans and matches phone numbers or customer IDs to customer information located in connected applications. It also combines contact history information from Amazon Connect—for example, the number of holds, transcripts, and customer information from CRMs.
- **Agent assist:** Amazon Connect Wisdom, a feature of Amazon Connect, delivers your agents the information they need to help reduce time spent searching for answers and improve customer satisfaction. They can search across connected data repositories to find answers and quickly resolve customer issues during a conversation in real time. Amazon Connect Wisdom links relevant knowledge repositories with built-in connectors

for third-party applications like Salesforce and ServiceNow and with internal wikis, FAQ stores, and file shares. Also, it uses ML-powered speech analytics in Contact Lens for Amazon Connect to automatically detect customer issues during calls and recommend content to your agents in real time, helping them resolve the issue without having to manually search.

- **Caller authentication and fraud risk detection:** Amazon Connect Voice ID uses ML to provide real-time caller authentication and fraud risk detection to make voice interactions faster and more secure. Amazon Connect Voice ID analyses caller's unique voice characteristics and carrier network metadata to provide your agents and self-service interactive voice response (IVR) systems with a real-time decision on a caller's identity for faster and more accurate verification. Amazon Connect Voice ID also screens for fraudulent actors in real time, based on your contact centre's custom watchlist, reducing potential losses from fraudulent attacks.
- **Real-time speech and sentiment analysis:** Contact Lens for Amazon Connect enables you to better understand the sentiment, trends, and compliance of customer conversations in your contact centre. This helps supervisors train agents, replicate successful interactions, and identify crucial company feedback. Supervisors can conduct fast full-text search on all transcripts to quickly troubleshoot customer issues. Using real-time analytics powered by ML, you can also get alerted to issues during live customer calls and deliver coaching to agents while calls are in progress, improving customer satisfaction.
- **Call summarization:** With call summarization, a capability of Contact Lens, important aspects of each customer call are automatically summarized, such as the outcome of the agent's actions and any follow-up items (for example, issuing a refund) to complete the request for the customer. Your contact centre agents and supervisors can easily access the call summary with just a few clicks in Amazon Connect without requiring any technical expertise.
- **Data redaction:** With data redaction (that is, data masking), a feature of Contact Lens, sensitive data such as name, address, and social security number are automatically detected and redacted from call recordings and transcripts. In addition, businesses can protect sensitive customer information by controlling access to the redacted and non-redacted data through user-defined permission groups.
- **Natural language chatbots with automated design:** You can easily build natural language chatbot contact flows using Amazon Lex, an AWS artificial intelligence (AI) service that is natively integrated within Amazon Connect and has the same automatic speech recognition (ASR) technology and natural language understanding (NLU) that powers Amazon Alexa.
- **Simple self-service and contact flow builder:** An Amazon Connect contact flow defines the customer experience with your contact centre from start to finish, including setting logging behaviour, setting text-to-speech language and voice, capturing customer inputs (spoken or by pressing 0-9 on the phone keypad), playing prompts, and transferring to appropriate queue. With the contact flow builder's graphical user interface in Amazon Connect, contact centre managers can easily create dynamic, personal, and automated customer experiences without needing to write a single line of code. Amazon Connect makes it possible to design automated contact flows that dynamically adapt to the caller experience in real time. With Amazon Connect, you also have the flexibility to use other AWS services. Using AWS Lambda, you can create targeted and personal

experiences by accessing virtually any back-end system and easily pulling in information, such as past purchases, contact history, and customer tendencies, that can be used to anticipate end-customer needs and deliver answers to questions before they are even asked.

You can also design contact flows to change based on information retrieved by Amazon Connect from AWS services (for example, Amazon DynamoDB, Amazon Redshift, or Amazon Aurora) or third-party systems (for example, CRM or analytics solutions). For example, an airline could design a contact flow to recognize a caller's phone number, look up their travel schedule in a booking database, and present options like "rebook," or "cancel" if the caller just missed a flight. Customers can also build contact flows that understand natural language using Amazon Lex, an AI service that has the same ASR and NLU technology that powers Amazon Alexa, so callers can simply say what they want instead of having to listen to long lists of menu options and guess which one is most closely related to what they want to do.

- **Real-time and historical analytics:** Understanding your contact centre at the most granular level is key to improving performance and lowering costs. Amazon Connect offers powerful analytics tools, including a visual dashboard with customizable real-time and historical metrics. With Amazon Connect you can also stream your most detailed contact metrics to the data lake of your choice where you can join and analyse them with other data like conversion rates or customer satisfaction. This enables your contact centre manager to make data-driven decisions to increase agent productivity and reduce customer wait times. Historical metrics also provide longer-term insights to identify common trends with customer issues and overall operational performance.
- **Call recording:** Amazon Connect also comes with integrated call recording for agent performance assessment to help monitor and improve customer experiences.

18.1.2. Benefits

- **Deliver omnichannel customer service:** Build high-quality omnichannel voice and interactive chat experiences to support your customers from anywhere. Use a single intuitive user interface (UI) for contact routing, queuing, and analytics.
- **Use built-in AI and ML to personalize interactions:** With embedded artificial intelligence (AI) and machine learning (ML), Amazon Connect makes it easy to automate interactions, understand customer sentiment, authenticate callers, and enable capabilities like interactive voice response (IVR) and chatbots.
- **Improve agent productivity:** Empower your agents to be more proactive and productive. Surface unified customer profiles and recommended answers in real time, and track follow-up tasks to quickly resolve customer issues.
- **Easy to use:** In just a few clicks, you can set up and make changes to your contact centre so that agents can begin helping customers right away.
- **Cost effective:** Save up to 80 percent compared to traditional contact centre solutions with no minimum fees, long-term commitments, or upfront license charges.
- **Flexible and scalable:** Easily scale up or down to meet demand, with the flexibility to onboard tens of thousands of agents working from anywhere.

18.2. Backup/Restore and Disaster Recovery

This requirement is not applicable for this service. For additional information beyond what is described herein, please refer to http://aws.amazon.com/documentation/.

18.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

18.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/connect/
- **Service quotas:** https://docs.aws.amazon.com/connect/latest/adminguide/amazon-connect-service-limits.html
- **Service FAQs:** https://aws.amazon.com/connect/faqs/

18.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/connect/ and the following links for comprehensive technical documentation regarding this service:

- Administrator Guide: Helps you get started using Amazon Connect. Learn how to provision, configure, monitor, and scale a virtual contact centre.

19. Amazon Detective

19.1. Service Overview

Amazon Detective makes it easy to analyse, investigate, and quickly identify the root cause of potential security issues or suspicious activities. Amazon Detective automatically collects log data from your AWS resources and uses machine learning, statistical analysis, and graph theory to build a linked set of data that enables you to easily conduct faster and more efficient security investigations.

AWS security services like Amazon GuardDuty, Amazon Macie, and AWS Security Hub as well as partner security products can be used to identify potential security issues, or findings. These services are really helpful in alerting you when something is wrong and pointing out where to go to fix it. But sometimes there might be a security finding where you need to dig a lot deeper and analyse more information to isolate the root cause and take action. Determining the root cause of security findings can be a complex process that often involves collecting and combining logs from many separate data sources, using extract, transform, and load (ETL) tools or custom scripting to organize the data, and then security analysts having to analyse the data and conduct lengthy investigations.

Amazon Detective simplifies this process by enabling your security teams to easily investigate and quickly get to the root cause of a finding. Amazon Detective can analyse trillions of events from multiple data sources such as Virtual Private Cloud (VPC) Flow Logs, AWS CloudTrail, and Amazon GuardDuty, and automatically creates a unified, interactive view of your resources, users, and the interactions between them over time. With this unified view, you can visualize all the details and context in one place to identify the underlying reasons for the findings, drill down into relevant historical activities, and quickly determine the root cause.

You can get started with Amazon Detective in just a few clicks in the AWS Console. There is no software to deploy, or data sources to enable and maintain.

19.1.1. Features

- **Automatic data collection across all your AWS accounts:** Amazon Detective automatically ingests and processes relevant data from all enabled accounts. You don't have to configure or enable any data sources. Amazon Detective collects and analyses events from data sources, such as AWS CloudTrail, VPC Flow Logs, and Amazon GuardDuty findings, and maintains up to a year of aggregated data for analysis.
- **Consolidates disparate events into a graph model:** Amazon Detective can analyse trillions of events from many separate data sources about the IP traffic, AWS management operations, and malicious or unauthorized activity to construct a graph model that distils log data using machine learning, statistical analysis, and graph theory to build a linked set of data for security investigations.
- **Interactive visualizations for efficient investigation:** Amazon Detective provides interactive visualizations that makes it easy to investigate issues faster and more thoroughly with less effort. With an unified view that enables you to visualize all the context and details in one place, it is easier to identify patterns that may validate or refute a security issue, and to understand all of the resources impacted by a security finding.
- **Newly observed geolocations:** The Amazon Detective geolocation map shows you activity coming from newly observed locations that weren't previously observed. This helps you to identify unusual activity and investigate if it is legitimate or suspicious activity.
- **Overall API call volume:** The Overall API call volume shows you successful and failed calls in a specific time period and compares it to the established baseline. This helps you to identify patterns of abnormal activity and validate a security finding.
- **Seamless integration for investigating a security finding:** Amazon Detective is integrated with AWS security services such as Amazon GuardDuty and AWS Security Hub as well as AWS partner security products to help quickly investigate security findings identified in these services.
- **Simple deployment with no upfront data source integration or complex configurations to maintain:** With few clicks in the AWS Management Console, you can enable Amazon Detective. There is no software to deploy, agents to install, or complex configurations to maintain. There are also no data sources to enable, which means you do not have to incur the costs of data source enablement, data transfer, and data storage.

19.1.2. Benefits

- **Faster and more effective investigations:** Amazon Detective presents a unified view of user and resource interactions over time, with all the context and details in one place to help you quickly analyse and get to the root cause of a security finding. For example, an Amazon GuardDuty finding, like an unusual Console Login API call, can be quickly investigated in Amazon Detective with details about the API call trends over time, and user login attempts on a geolocation map. These details enable you to quickly identify if you think it is legitimate or an indication of a compromised AWS resource.

- **Save time and effort with continuous data updates:** Amazon Detective automatically processes terabytes of event data records about IP traffic, AWS management operations, and malicious or unauthorized activity. It organizes the data into a graph model that summarizes all the security-related relationships in your AWS environment. Amazon Detective then queries this model to create visualizations used in investigations. The graph model is continuously updated as new data becomes available from AWS resources, so you spend less time managing constantly changing data.
- **Easy to use visualizations:** Amazon Detective produces visualizations with the information you need to investigate and respond to security findings. It helps you answer questions like 'is this normal for this role to have so many failed API calls?' or 'is this spike in traffic from this instance expected?' without having to organize any data or develop, configure, or tune your own queries and algorithms. Amazon Detective maintains up to a year of aggregated data that shows changes in the type and volume of activity over a selected time window, and links those changes to security findings.

19.2. Backup/Restore and Disaster Recovery

This requirement is not applicable for this service. For additional information beyond what is described herein, please refer to http://aws.amazon.com/documentation/.

19.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

19.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/detective/
- **Service quotas:** https://docs.aws.amazon.com/general/latest/gr/detective.html
- **Service FAQs:** https://aws.amazon.com/detective/faqs/

19.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/detective/ and the following links for comprehensive technical documentation regarding this service.

- Administration Guide: Provides information on how to enable and disable the Detective service and how to manage the accounts in a Detective behaviour graph. Describes other administrative functions such as logging, tracking usage, and maintaining security.
- User Guide: Provides information on how to investigate suspicious activity using the visualizations generated from a Detective behaviour graph. Identifies the possible investigation starting points and shows how to use the Detective tools to conduct an investigation.
- API Reference: Describes the API operations for Amazon Detective. Provides example requests and responses for each operation.

20. Amazon DevOps Guru

20.1. Service Overview

Amazon DevOps Guru is a service powered by machine learning (ML) that is designed to make it easy to improve an application's operational performance and availability. DevOps Guru helps detect behaviours that deviate from normal operating patterns so you can identify operational issues long before they impact your customers.

DevOps Guru uses ML models informed by years of Amazon.com and AWS operational excellence to identify anomalous application behaviour (for example, increased latency, error rates, resource constraints, and others) and helps surface critical issues that could cause potential outages or service disruptions. When DevOps Guru identifies a critical issue, it automatically sends an alert and provides a summary of related anomalies, the likely root cause, and context for when and where the issue occurred. When possible, DevOps Guru also helps provide recommendations on how to remediate the issue.

With one-click deployment, DevOps Guru automatically ingests operational data from your AWS applications and provides a single dashboard to visualize issues in your operational data. You can get started by enabling DevOps Guru for all resources in your AWS account, resources in your AWS CloudFormation Stacks, or resources grouped together by AWS Tags, with no manual setup or ML expertise required.

20.1.1. Features

- **Consolidate operational data from multiple sources:** Amazon DevOps Guru continuously analyzes and consolidates streams of operational data from multiple sources such as Amazon CloudWatch metrics, AWS Config, AWS CloudFormation, and AWS X-Ray and provides you with a single-console dashboard to search for and visualize anomalies in your operational data, thereby reducing the need to use multiple tools. This delegated administrator can then view, sort, and filter insights from all accounts within your organization to develop an org-wide view of the health of all monitored applications—without the need for any additional customization.
- **Save time with ML-powered insights:** Amazon DevOps Guru improves application availability and remediates operational issues faster with less manual effort by using ML-powered recommendations. It continuously ingests and analyzes metrics, logs, events, and traces to establish normal bounds for application behavior. DevOps Guru then looks for deviations from normal behavior and aggregates anomalies to create operational insights based on component relationships in your application. Operational insights include information on which components are impacted, identification of related anomalies, and recommendations on how to remediate using contextual data such as AWS CloudTrail events.
- **Automatically configure alarms:** Developers and operators can enable Amazon DevOps Guru to configure and set up alarms for their applications. As applications evolve and you adopt new services, DevOps Guru automatically recognizes the new resources and ingests related metrics. It then alerts you when a deviation occurs from normal operating patterns without requiring any manual updates to rules and alarms.
- **Detect the most critical issues with minimal noise:** Amazon DevOps Guru leverages years of experience operating universally available applications such as Amazon.com and uses ML models trained on internal AWS operational data to provide accurate operational insights for critical issues that impact applications.

- **One-click deployment, with no additional software to deploy and manage:** With one click in the AWS Management Console or a single API call, you can enable Amazon DevOps Guru on a single account. Amazon DevOps Guru also supports multi-account insight visibility through AWS Organizations integration. Once enabled, Amazon DevOps Guru uses ML to automatically collect and analyze data such as application metrics, logs, events, and behaviors that deviate from normal operating patterns. There are no additional services to deploy or manage.
- **Integrate with AWS services and third-party tools:** Amazon DevOps Guru natively integrates with Amazon CloudWatch, AWS Config, AWS CloudFormation, and AWS X-Ray to discover and track connections and dependencies between application components. DevOps Guru also integrates with AWS Systems Manager and Amazon EventBridge. The integration with AWS Systems Manager enables you to automatically receive an OpsItem in OpsCenter for each insight that DevOps Guru generates. This allows you to leverage OpsCenter functionality to further view, investigate, and resolve operational issues faster. The integration with Amazon EventBridge enables you to set up routing rules to determine where to send notifications, use pre-defined DevOps Guru patterns to only send notifications or trigger actions that match that pattern (e.g., only send for “New Insights Open”), or create custom patterns to send notifications. DevOps Guru is also integrated with third-party incident management tools from PagerDuty and Atlassian who are able to ingest SNS notifications from DevOps Guru, so you can automatically manage incidents within their platform.

20.1.2. Benefits

Copy and paste benefits from service landing page

- **Automatically detect operational issues:** Using ML, Amazon DevOps Guru automatically collects and analyzes data such as application metrics, logs, events, and behaviors that deviate from normal operating patterns. The service is designed to automatically detect and alert on operational issues and risks, such as impending resource exhaustion, code and configuration changes that may cause outages, memory leaks, under-provisioned compute capacity, and database input/output (I/O) overutilization.
- **Resolve issues quickly with ML-powered insights:** Amazon DevOps Guru helps reduce time to identify and resolve the root cause of issues by correlating anomalous behavior and operational events. When an issue occurs, DevOps Guru is designed to generate insights with a summary of related anomalies and contextual information about the issue. When possible, it helps provide actionable recommendations for remediation.
- **Easily scale and maintain availability:** Amazon DevOps Guru saves you the time and effort involved in manually updating static rules and alarms so you can effectively monitor complex and evolving applications. When you migrate or adopt new AWS services, DevOps Guru automatically analyzes their metrics, logs, and events. Then it produces insights, helping you easily adapt to changing behavior and evolving system architecture.
- **Reduce noise and alarm fatigue:** Amazon DevOps Guru helps developers and IT operators reduce alarm noise and overcome alarm fatigue by using pre-trained ML models to correlate and group related anomalies and surface the most critical alerts. With DevOps Guru, you can reduce the need to manage multiple monitoring tools and alarms, which means you can focus on the root cause of the issue and remediation.

20.2. Backup/Restore and Disaster Recovery

This requirement is not applicable for this service. For additional information beyond what is described herein, please refer to http://aws.amazon.com/documentation/.

20.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

20.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/devops-guru/index.html
- **Service quotas:** https://docs.aws.amazon.com/devops-guru/latest/userguide/quotas.html
- **Service FAQs:** https://aws.amazon.com/devops-guru/faqs/?nc=sn&loc=5

20.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/devops-guru/index.html and the following links for comprehensive technical documentation regarding this service.

- **User Guide:** Describes how to use Amazon DevOps Guru to generate insights using machine learning to help you improve the performance of your operational applications.
- **API Reference:** Describes the API operations for Amazon DevOps Guru. Also provides details of related request and response syntax and errors.
- **AWS CLI Reference for Amazon DevOps Guru:** Describes the AWS CLI commands that you can use to automate building your source code.

21. Amazon DocumentDB (with MongoDB compatibility)

21.1. Service Overview

Amazon DocumentDB (with MongoDB compatibility) is a database service that is purpose-built for JSON data management at scale, fully managed and integrated with AWS, and enterprise-ready with high durability. This scalable service offers you the durability you need when operating mission-critical MongoDB workloads. In Amazon DocumentDB, storage scales automatically up to 64 TB without any impact to your application. DocumentDB supports millions of requests per second with up to 15 low latency read replicas in minutes, without any application downtime, regardless of the size of your data. Amazon DocumentDB is designed for 99.99% availability and replicates six copies of your data across three AWS Availability Zones (AZs). You can use AWS Database Migration Service (DMS) for free (for six months) to easily migrate your self-managed MongoDB databases to Amazon DocumentDB with virtually no downtime.

21.1.1. Features

- **MongoDB-compatible:** Amazon DocumentDB is compatible with MongoDB 3.6 and 4.0 drivers and tools. A vast majority of the applications, drivers, and tools that customers already use today with their open-source MongoDB non-relational database can be used with Amazon DocumentDB. Amazon DocumentDB emulates the responses that a client expects from a MongoDB server by implementing the Apache 2.0 open source

MongoDB 3.6 and 4.0 APIs on a purpose-built, distributed, fault-tolerant, and self-healing storage system that gives customers the performance, scalability, and availability they need when operating mission-critical MongoDB workloads at scale.

- **Fully Managed:** Getting started with Amazon DocumentDB is easy. Just launch a new Amazon DocumentDB cluster using the AWS Management Console. Amazon DocumentDB instances are pre-configured with parameters and settings appropriate for the instance class you have selected. You can launch a cluster and connect your application within minutes without additional configuration. Amazon DocumentDB will keep your database up-to-date with the latest patches. You can control if and when your cluster is patched via Database Engine Version Management.
- **Performance at scale:** Amazon DocumentDB has a flexible JSON document model, data types, and efficient indexing. The service uses a scale-up, in-memory optimized architecture to allow for fast query evaluation over large documents sets. With a few clicks in the AWS Management Console, you can scale the compute and memory resources, up or down, by creating new replica instances of the desired size or by removing instances. Compute scaling operations typically complete in a few minutes. Amazon DocumentDB will automatically grow the size of your storage volume as your cluster storage needs grow. Your storage volume will grow in increments of 10 GB up to a maximum of 64 TB. You don't need to provision excess storage for your NoSQL database to handle future growth.
- **Highly Secure and Compliant:** Amazon DocumentDB runs in Amazon Virtual Private Cloud (VPC), which allows you to isolate your cluster in your own virtual network and connect to your on-premises IT infrastructure using industry- standard encrypted IPsec virtual private networks (VPNs). In addition, using Amazon DocumentDB's VPC configuration, you can configure firewall settings and control network access to your cluster.
- **Encryption:** Amazon DocumentDB allows you to encrypt your databases using keys you create and control through AWS Key Management Service (KMS). On a cluster running with Amazon DocumentDB encryption, data stored at rest in the underlying storage is encrypted, as are the automated backups, snapshots, and replicas in the same cluster. By default, connections between a client and Amazon DocumentDB are encrypted-in-transit with TLS.
- **Highly Available:** Amazon DocumentDB Global Clusters provides disaster recovery from region-wide outages and enables low-latency global reads. Amazon DocumentDB Global Clusters replicates your data to clusters in up to 5 AWS regions with little to no impact on performance, with a typical lag of less than one second. Learn more about setting up Global Clusters in the Amazon DocumentDB user guide.
- **Instance Monitoring and Repair:** The health of your Amazon DocumentDB cluster and its instances are continuously monitored. If the instance powering your database fails, the instance and associated processes are automatically restarted. Amazon DocumentDB recovery does not require the potentially lengthy replay of database redo logs, so your instance restart times are typically 30 seconds or less. It also isolates the database cache from database processes, allowing the cache to survive a database restart.
- **Fault-tolerant and Self-healing Storage:** Each 10 GB portion of your storage volume is replicated six ways, across three Availability Zones (AZs). Amazon DocumentDB uses

fault-tolerant storage that transparently handles the loss of up to two copies of data without affecting database write availability and up to three copies without affecting read availability. Amazon DocumentDB's storage is also self-healing; data blocks and disks are continuously scanned for errors and replaced automatically.

21.1.2. Benefits

- **Scalable:** Support millions of document read requests per second by scaling compute and storage independently.
- **Automated:** Automate hardware provisioning, patching, setup, and other database management tasks.
- **Durable:** Achieve 99.999999999% durability with automatic replication, continuous backup, and strict network isolation.
- **Compatible:** Use existing MongoDB drivers and tools with the Apache 2.0 open-source MongoDB 3.6 and 4.0 APIs.

21.2. Backup/Restore and Disaster Recovery

This service allows for backup and recovery. Automated backups are always enabled on Amazon DocumentDB clusters. Amazon DocumentDB automatically maintains six copies of your data across three Availability Zones and will automatically attempt to recover your instance in a healthy AZ with no data loss. In the unlikely event your data is unavailable within Amazon DocumentDB storage, you can restore from a cluster snapshot or perform a point-in-time restore operation to a new cluster. Amazon DocumentDB automatically divides your storage volume into 10GB segments spread across many disks. Each 10GB chunk of your storage volume is replicated six ways, across three Availability Zones.

21.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

21.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/documentdb/
- **Service quotas:** https://docs.aws.amazon.com/documentdb/latest/developerguide/limits.html
- **Service FAQs:** https://aws.amazon.com/documentdb/faqs/

21.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/documentdb/ and the following links for comprehensive technical documentation regarding this service.

- Developer Guide: Provides a conceptual overview of Amazon DocumentDB and provides instructions on using the various features with both the console and the command line interface.
- Amazon DocumentDB section of the AWS CLI Reference: Describes the command line interface for Amazon DocumentDB (with MongoDB compatibility) in detail. Provides basic syntax, options, and usage examples for each command.

22. Amazon DynamoDB

22.1. Service Overview

Amazon DynamoDB is a NoSQL database that supports key-value and document data models. Developers can use DynamoDB to build modern, serverless applications that can start small and scale globally to support petabytes of data and tens of millions of read and write requests per second. DynamoDB is designed to run high-performance, internet-scale applications that would overburden traditional relational databases.

22.1.1. Features

- **Performance at scale:** DynamoDB is a key-value and document database that can support tables of virtually any size with horizontal scaling. This enables DynamoDB to scale to more than 10 trillion requests per day with peaks greater than 20 million requests per second, over petabytes of storage.
- **Key-value and document data models:** DynamoDB supports both key-value and document data models. This enables DynamoDB to have a flexible schema, so each row can have any number of columns at any point in time. This allows you to easily adapt the tables as your business requirements change, without having to redefine the table schema as you would in relational databases.
- **Microsecond latency with DynamoDB Accelerator:** DynamoDB Accelerator (DAX) is an in-memory cache that delivers fast read performance for your tables at scale by enabling you to use a fully managed in-memory cache. Using DAX, you can improve the read performance of your DynamoDB tables by up to 10 times—taking the time required for reads from milliseconds to microseconds, even at millions of requests per second.
- **Automated global replication with global tables:** DynamoDB global tables replicate your data automatically across your choice of AWS Regions and automatically scale capacity to accommodate your workloads. With global tables, your globally distributed applications can access data locally in the selected regions to get single-digit millisecond read and write performance.
- **Advanced streaming applications with Kinesis Data Streams for DynamoDB:** Amazon Kinesis Data Streams for DynamoDB captures item-level changes in your DynamoDB tables as a Kinesis data stream. This feature enables you to build advanced streaming applications such as real-time log aggregation, real-time business analytics, and Internet of Things data capture. Through Kinesis Data Streams, you also can use Amazon Kinesis Data Firehose to deliver DynamoDB data automatically to other AWS services.
- **Serverless:** With DynamoDB, there are no servers to provision, patch, or manage, and no software to install, maintain, or operate. DynamoDB automatically scales tables to adjust for capacity and maintains performance with zero administration. Availability and fault tolerance are built in, eliminating the need to architect your applications for these capabilities.
- **Read/write capacity modes:** DynamoDB provides capacity modes for each table: on-demand and provisioned. For workloads that are less predictable for which you are unsure that you will have high utilization, on-demand capacity mode takes care of managing capacity for you, and you only pay for what you consume. Tables using provisioned capacity mode require you to set read and write capacity. Provisioned

capacity mode is more cost effective when you're confident you'll have decent utilization of the provisioned capacity you specify.

- **On-demand mode:** For tables using on-demand capacity mode, DynamoDB instantly accommodates your workloads as they ramp up or down to any previously reached traffic level. If a workload's traffic level hits a new peak, DynamoDB adapts rapidly to accommodate the workload. You can use on-demand capacity mode for both new and existing tables, and you can continue using the existing DynamoDB APIs without changing code.
- **Auto scaling:** For tables using provisioned capacity, DynamoDB delivers automatic scaling of throughput and storage based on your previously set capacity by monitoring the performance usage of your application. If your application traffic grows, DynamoDB increases throughput to accommodate the load. If your application traffic shrinks, DynamoDB scales down so that you pay less for unused capacity.
- **Change tracking with triggers:** DynamoDB integrates with AWS Lambda to provide triggers. Using triggers, you can automatically execute a custom function when item-level changes in a DynamoDB table are detected. With triggers, you can build applications that react to data modifications in DynamoDB tables. The Lambda function can perform any actions you specify, such as sending a notification or initiating a workflow.

22.1.2. Benefits

- **Performance:** Deliver apps with consistent single-digit millisecond performance, nearly unlimited throughput and storage, and automatic multi-region replication.
- **Security:** Secure your data with encryption at rest, automatic backup and restore, and guaranteed reliability with an SLA of up to 99.999% availability.
- **Fully Managed:** Focus on innovation and optimize costs with a fully managed serverless database that automatically scales up and down to fit your needs.
- **Integration:** Integrate with AWS services to do more with your data. Use built-in tools to perform analytics, extract insights, and monitor traffic trends.

22.2. Backup/Restore and Disaster Recovery

This service allows for backup and recovery. The service can back up tables and data. Users control this via a manual or scheduled API call. Users schedule and recover backups through a web interface.

22.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

22.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/dynamodb/
- **Service quotas:** https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/ServiceQuotas.html

- **Service FAQs:** https://aws.amazon.com/dynamodb/faqs/

22.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/dynamodb/ and the following links for comprehensive technical documentation regarding this service.

- Developer Guide: Provides a conceptual overview of Amazon DynamoDB and includes detailed development instructions for using the various features.
- API Reference: Describes all the API operations for Amazon DynamoDB in detail. Also provides sample requests, responses, and errors for the supported web services protocols.

23. Amazon ECS Anywhere

23.1. Service Overview

Amazon Elastic Container Service (ECS) Anywhere enables you to easily run and manage container workloads on customer-managed infrastructure.

ECS Anywhere builds upon the ease and simplicity of Amazon ECS to provide a consistent tooling and API experience across your container-based applications. Whether on-premises or in the cloud, you'll have similar cluster management, workload scheduling, and monitoring you've come to know from Amazon ECS. Reduce costs and mitigate complex local container orchestration by taking advantage of the completely managed solution that ECS Anywhere provides. ECS Anywhere helps you meet compliance requirements and scale your business without sacrificing your on-premises investments.

23.1.1. Features

- **Containerize existing on-premises workloads:** Rather than installing and operating a local control plane, you can use the same hyperscale, trusted, and fully managed Amazon ECS control plane for your on-premises container workloads.
- **Data processing workloads at the edge:** Run containerized data processing workloads at edge locations on your own hardware with Amazon ECS Anywhere so that you can stay close to your end customers and maintain reduced latency.
- **Burst to cloud as needed:** With Amazon ECS Anywhere, you can use your on-premises infrastructure as base capacity while bursting into the AWS cloud for additional capacity to meet peaks in demand and as your business grows. ECS Anywhere offers a cost-effective way for you to meet your compute needs as your business scales.
- **Make use of existing capital investments:** Amazon ECS Anywhere enables you to leverage your existing capital investments while simultaneously taking advantage of running workloads in the cloud. Consistent Amazon ECS tooling on-premises and in the cloud makes it easier for you to migrate your containers workloads to the cloud in the future, if you choose.
- **On-premises ML and video processing workloads:** Utilize your on-premises GPU compute capacity to run GPU-based container workloads while keeping the simplicity of a fully managed orchestration service provided by ECS. Run machine learning, image processing, 3D visualization, big data, among other applications without the need to transfer your data to the cloud or manage a third party orchestration software.

23.1.2. Benefits

- **Fully managed cloud control plane:** With Amazon ECS Anywhere, you don't need to run and operate separate container management software for your on-premises container workloads. With the familiar managed in-region ECS control plane to orchestrate your containers and run tasks on your infrastructure, you'll spend less time on operational overhead and more time focusing on driving innovation for your business.
- **Consistent tooling and governance:** Amazon ECS Anywhere makes it easy for you to use the same ECS APIs, cluster management, workload scheduling, and monitoring for all of your container-based applications. This ensures a simple and consistent operator experience no matter where your applications are running.
- **Manage your hybrid footprint:** Amazon ECS Anywhere enables you to run applications in both on-premises environments and the cloud with a standardized container orchestrator, removing the need for your team to learn multiple domains and skillsets and manage complex software on their own.
- **Helps fulfil your compliance and business requirements:** Satisfy compliance, data gravity, and other business requirements by running your workloads on infrastructure you own while enjoying simple and familiar ECS tooling.

23.2. Backup/Restore and Disaster Recovery

This service allows for backup and recovery. The service can back up machine images. Users control this via a manual or scheduled API call. Users schedule and recover backups through a web interface.

23.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

23.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/AmazonECS/latest/developerguide/ecs-anywhere.html
- **Service quotas:** https://docs.aws.amazon.com/AmazonECS/latest/developerguide/service-quotas.html
- **Service FAQs:** https://aws.amazon.com/ecs/anywhere/faqs/

23.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/AmazonECS/latest/developerguide/ecs-anywhere.html for comprehensive technical documentation regarding this service.

24. Amazon EKS Anywhere

24.1. Service Overview

Amazon EKS Anywhere helps simplify the creation and operation of on-premises Kubernetes clusters with default component configurations while providing tools for automating cluster management. It builds on the strengths of Amazon EKS Distro: the same Kubernetes

distribution that powers Amazon EKS on AWS. AWS supports all Amazon EKS Anywhere components including the integrated 3rd-party software, so that customers can reduce their support costs and avoid maintenance of redundant open-source and third-party tools. In addition, Amazon EKS Anywhere gives customers on-premises Kubernetes operational tooling that's consistent with Amazon EKS.

24.1.1. Features

- **Hybrid cloud consistency:** You may have lots of Kubernetes workloads on Amazon EKS but also need to operate Kubernetes clusters on-premises. Amazon EKS Anywhere offers strong operational consistency with Amazon EKS so you can standardize your Kubernetes operations based on a unified toolset.
- **Disconnected environment:** You may need to secure your applications in disconnected environment or run applications in areas without internet connectivity. Amazon EKS Anywhere allows you to deploy and operate highly-available clusters with the same Kubernetes distribution that powers Amazon EKS on AWS.
- **Application modernization:** Amazon EKS Anywhere empowers you to modernize your on-premises applications, removing the heavy lifting of keeping up with upstream Kubernetes and security patches, so you can focus on your core business value.
- **Data sovereignty:** You may want to keep your large data sets on-premises due to legal requirements concerning the location of the data. Amazon EKS Anywhere brings the trusted Amazon EKS Kubernetes distribution and tools to where your data needs to be.

24.1.2. Benefits

- **Simplify on-premises Kubernetes management:** Amazon EKS Anywhere helps simplify the creation and operation of on-premises Kubernetes clusters with default component configurations while providing tools for automating cluster management.
- **One stop support:** AWS supports all Amazon EKS Anywhere components including the integrated 3rd-party software, so that customers can reduce their support costs and avoid maintenance of redundant open-source and third-party tools.
- **Consistent and reliable:** Amazon EKS Anywhere gives you on-premises Kubernetes operational tooling that's consistent with Amazon EKS. It builds on the strengths of Amazon EKS Distro and provides open-source software that's up-to-date and patched, so you can have a Kubernetes environment on-premises that is more reliable than self-managed Kubernetes offerings.

24.2. Backup/Restore and Disaster Recovery

This service allows for backup and recovery. The service can back up data on the cluster by publishing snapshots to S3. Users control this via User intervention and schedules. Users schedule and recover backups through a web interface.

24.3. Pricing Overview

"Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace."

24.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/eks/index.html

- **Service quotas:** https://docs.aws.amazon.com/eks/latest/userguide/service-quotas.html
- **Service FAQs:** https://aws.amazon.com/eks/eks-anywhere/faqs/

24.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/eks/index.html and the following links for comprehensive technical documentation regarding this service.

- **User Guide:** Describes key concepts of Amazon EKS and provides instructions for using the features of Amazon EKS.
- **API Reference:** Documents the Amazon EKS API.
- **Amazon EKS section of AWS CLI Reference:** Documents the Amazon EKS commands available in the AWS Command Line Interface (AWS CLI).
- **Amazon EKS module of AWS Cloud Development Kit:** Describes how to define Amazon EKS clusters using the AWS Cloud Development Kit.
- **Amazon EKS Workshop:** Interactive workshop that shows host to use Amazon EKS features with other AWS services.
- **Amazon EKS best practices guides:** Describes best practices for using Amazon EKS.

25. Amazon Elastic Block Store (EBS)

25.1. Service Overview

Amazon EBS allows you to create storage volumes and attach them to Amazon EC2 instances. Once attached, you can create a file system on top of these volumes, run a database, or use them in any other way you would use block storage. Amazon EBS volumes are placed in a specific Availability Zone where they are automatically replicated to protect you from the failure of a single component. All EBS volume types offer durable snapshot capabilities and are designed for 99.999% availability.

25.1.1. Features

- **Amazon EBS volume types:** Amazon EBS provides a range of options that allow you to optimize storage performance and cost for your workload. These options are divided into two major categories: SSD-backed storage for transactional workloads, such as databases and boot volumes (performance depends primarily on IOPS), and HDD-backed storage for throughput intensive workloads, such as MapReduce and log processing (performance depends primarily on MB/s).
- **Amazon data lifecycle manager for EBS snapshots:** Data Lifecycle Manager for EBS snapshots provides a simple, automated way to back up data stored on EBS volumes by ensuring that EBS snapshots are created and deleted on a custom schedule. You no longer need to use scripts or other tools to comply with data backup and retention policies specific to your organization or industry.
- **Amazon EBS Elastic Volumes:** Elastic Volumes is a feature that allows you to easily adapt your volumes as the needs of your applications change. Elastic Volumes allows you to dynamically increase capacity, tune performance, and change the type of any new or existing current generation volume with no downtime or performance impact. Easily right-size your deployment and adapt to performance changes.
- **Amazon EBS Snapshots:** Amazon EBS provides the ability to save point-in-time snapshots of your volumes to Amazon S3. Amazon EBS Snapshots are stored incrementally: only the blocks that have changed after your last snapshot are saved, and

you are billed only for the changed blocks. If you have a device with 100 GB of data but only 5 GB has changed after your last snapshot, a subsequent snapshot consumes only 5 additional GB and you are billed only for the additional 5 GB of snapshot storage, even though both the earlier and later snapshots appear complete.

- **Amazon EBS-Optimized instances:** For an additional low, hourly fee, customers can launch certain Amazon EC2 instance types as EBS-optimized instances. EBS-optimized instances enable EC2 instances to fully use the IOPS provisioned on an EBS volume.
- **Amazon EBS availability and durability:** Amazon EBS volumes are designed to be highly available, reliable, and durable. At no additional charge to you, Amazon EBS volume data is replicated across multiple servers in an Availability Zone to prevent the loss of data from the failure of any single component. Amazon EBS offers a higher durability volume (io2), that is designed to provide 99.999% durability with an annual failure rate (AFR) of 0.001%, where failure refers to a complete or partial loss of the volume. For example, if you have 100,000 EBS io2 volumes running for 1 year, you should expect only one io2 volume to experience a failure. This makes io2 ideal for business-critical applications such as SAP HANA, Oracle, Microsoft SQL Server and IBM DB2 that will benefit from higher uptime. io2 volumes are 2000 times more reliable than typical commodity disk drives, which fail with an AFR of around 2%. All other Amazon EBS volumes are designed to provided 99.8%-99.9% durability with an AFR of between 0.1% - 0.2%,
- **Amazon EBS encryption and AWS Identity and Access Management:** Amazon EBS encryption offers seamless encryption of EBS data volumes, boot volumes and snapshots, eliminating the need to build and manage a secure key management infrastructure. EBS encryption enables data at rest security by encrypting your data volumes, boot volumes and snapshots using Amazon-managed keys or keys you create and manage using the AWS Key Management Service (KMS). In addition, the encryption occurs on the servers that host EC2 instances, providing encryption of data as it moves between EC2 instances and EBS data and boot volumes.

25.1.2. Benefits

- **Scalable:** Scale fast for your most demanding, high-performance workloads, including mission-critical applications such as SAP, Oracle, and Microsoft products.
- **Available and durable:** Protect against failures with 99.999% availability, including replication within Availability Zone (AZs), and 99.999% durability with io2 Block Express volumes.
- **Range of storage types:** Select the storage that best fits your workload. Volumes range from cost-effective dollar-per-GB to high performance with the fastest IOPS and throughput.

25.2. Backup/Restore and Disaster Recovery

This service allows for backup and recovery. The service can back up volumes. Users control this via a manual or scheduled API call. Users schedule and recover backups through a web interface.

25.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

25.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/ebs/
- **Service quotas:** https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-resource-quotas.html
- **Service FAQs:** https://aws.amazon.com/ebs/faqs/

25.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/ebs/ and the following links for comprehensive technical documentation regarding this service.

- User Guide: Describes key concepts of Amazon EBS and provides instructions for using the Amazon EBS volumes.
- API Reference: Describes the API operations for Amazon EBS.

26. Amazon Elastic Compute Cloud (EC2)

26.1. Service Overview

Amazon EC2 provides the broadest and deepest instance choice to match your workload's needs. General purpose, compute optimized, memory optimized, storage optimized, and accelerated computing instance types are available that provide the optimal compute, memory, storage, and networking balance for your workloads. Processors from Intel, AMD, NVIDIA and AWS power these instance types and provide additional performance and cost optimizations. Local storage and enhanced networking options available with instance types further help optimize performance for workloads that are disk or network I/O bound. Many instance types also offer bare metal instances that provide your applications with direct access to the processor and memory of the underlying server for running in non-virtualized environments or for applications where you want to use your own hypervisor.

26.1.1. Features

- **Multiple Locations:** Amazon EC2 provides the ability to place instances in multiple locations. Amazon EC2 locations are composed of Regions and Availability Zones. Availability Zones are distinct locations that are engineered to be insulated from failures in other Availability Zones and provide inexpensive, low latency network connectivity to other Availability Zones in the same Region. By launching instances in separate Availability Zones, you can protect your applications from failure of a single location. Regions consist of one or more Availability Zones and are geographically dispersed. The Amazon EC2 Service Level Agreement commitment is 99.99% availability for each Amazon EC2 Region. Please refer to Regional Products and Services for more details of our product and service availability by region.
- **High Precision Time with Amazon Time Sync Service:** The Amazon Time Sync Service provides a highly accurate, reliable and available time source to AWS services including EC2 instances. For instructions on how to access the service, see Setting the Time sections of the Linux and Windows User Guides.
- **Choice of operating systems and software:** Amazon Machine Images (AMIs) are preconfigured with an ever-growing list of operating systems, including Microsoft

Windows and Linux distributions such as Amazon Linux 2, Ubuntu, Red Hat Enterprise Linux, CentOS, SUSE and Debian. We work with our partners and community to provide you with the most choice possible. The AWS Marketplace features a wide selection of commercial and free software from well-known vendors, designed to run on your EC2 instances.

- **Pay for What You Use:** With per-second billing, you only pay for what you use. It takes the cost of unused minutes and seconds in an hour off of the bill, so you can focus on improving your applications instead of maximizing usage to the hour. Learn more about EC2 pricing.
- **Scale Seamlessly with Amazon EC2 Auto Scaling:** Amazon EC2 Auto Scaling allows you to automatically scale your Amazon EC2 capacity up or down according to conditions you define. You can use the dynamic and predictive scaling policies within EC2 Auto Scaling to add or remove EC2 instances. Predictive scaling uses machine learning to proactively allocate instances based on anticipated demand, and dynamic scaling allows you to scale compute based on defined metrics. With EC2 Auto Scaling, you can ensure that the number of Amazon EC2 instances you're using scales up seamlessly during demand spikes to maintain performance, and scales down automatically during demand lulls to minimize costs. See Amazon EC2 Auto Scaling for more details.
- **Optimize Compute Performance and Cost with Amazon EC2 Fleet:** With a single API call, Amazon EC2 Fleet lets you provision compute capacity across EC2 instance types, Availability Zones, and purchase models to help optimize scale, performance and cost. Read FAQs and this AWS blog to learn more. You can also access EC2 Fleet capabilities via Amazon EC2 Auto Scaling to provision and automatically scale compute capacity across EC2 instance types, Availability Zones, and purchase options in a single Auto Scaling Group.
- **Optimized CPU Configurations:** The Optimize CPUs feature gives you greater control of your Amazon EC2 instances on two fronts. First, you can specify a custom number of vCPUs when launching new instances to save on vCPU-based licensing costs. Second, you can disable Intel Hyper-Threading Technology (Intel HT Technology) for workloads that perform well with single-threaded CPUs, such as certain high-performance computing (HPC) applications. To learn more about how Optimize CPUs can help you, visit the Optimize CPUs documentation here.
- **Pause and Resume Your Instances:** You can hibernate your Amazon EC2 instances backed by Amazon EBS, and resume them from this state at a later time. Applications that take a while to bootstrap and persist state into memory (RAM) can benefit from this feature.
- **Optimal storage for every workload:** Different Amazon EC2 workloads can have vastly different storage requirements. Beyond the built-in instance storage, we also offer Amazon Elastic Block Store (Amazon EBS) and Amazon Elastic File System (Amazon EFS) to suit other cloud storage workload requirements. Amazon EBS provides persistent, highly available, consistent, low-latency block storage volumes for use with Amazon EC2 instances, while Amazon EFS provides simple, scalable, persistent, fully managed cloud file storage for shared access.
- **High Packet-Per-Second Performance and Low Latency with Enhanced Networking:** Enhanced Networking enables you to get significantly higher packet per

second (PPS) performance, lower network jitter and lower latencies. This feature uses a network virtualization stack that provides higher I/O performance and lower CPU utilization compared to traditional implementations. For instructions on how to enable Enhanced Networking on EC2 instances, see the Enhanced Networking on Linux and Enhanced Networking on Windows tutorials. For availability of this feature by instance, or to learn more, visit the Enhanced Networking FAQ section.

- **Run High Levels of Inter-Node Communications with Elastic Fabric Adapter:** Elastic Fabric Adapter (EFA) is a network interface for Amazon EC2 instances that enables customers to run applications requiring high levels of inter-instance communications, like machine learning, computational fluid dynamics, weather modelling, and reservoir simulation, at scale on AWS. EFA is available as an optional EC2 networking feature that you can enable on any supported EC2 instance at no additional cost.
- **Manage Dynamic Cloud Computing Services with Elastic IP Addresses:** Elastic IP addresses are static IP addresses designed for dynamic cloud computing. An Elastic IP address is associated with your account, not with a particular instance, and you control that address until you choose to explicitly release it. Unlike traditional static IP addresses, however, Elastic IP addresses allow you to mask instance or Availability Zone failures by programmatically remapping your public IP addresses to any instance in your account. You can also optionally configure the reverse DNS record of any of your Elastic IP addresses by filling out this form.
- **High Throughput and Low Latency with High Performance Computing (HPC) Clusters:** Customers with complex computational workloads such as tightly coupled parallel processes, or with applications sensitive to network performance, can achieve the same high compute and network performance provided by custom-built infrastructure while benefiting from the elasticity, flexibility and cost advantages of Amazon EC2.
- **Access Services Hosted on AWS Easily and Securely with AWS PrivateLink:** AWS PrivateLink is a purpose-built technology designed for customers to access Amazon services in a highly performant and highly available manner, while keeping all the network traffic within the AWS network.

26.1.2. Benefits

- **Broadest and deepest provision:** Amazon Elastic Compute Cloud (Amazon EC2) offers the broadest and deepest compute platform, with over 475 instances and choice of the latest processor, storage, networking, operating system, and purchase model to help you best match the needs of your workload.
- **Wide range of instances:** We are the first major cloud provider that supports Intel, AMD, and Arm processors, the only cloud with on-demand EC2 Mac instances, and the only cloud with 400 Gbps Ethernet networking.
- **Price performance:** We offer the best price performance for machine learning training, as well as the lowest cost per inference instances in the cloud.
- **Widely used:** More SAP, high performance computing (HPC), ML, and Windows workloads run on AWS than any other cloud.
- **Reliable and scalable:** Access reliable, scalable infrastructure on demand. Scale capacity within minutes with SLA commitment of 99.99% availability.

- **Security:** Provide secure compute for your applications. Security is built into the foundation of Amazon EC2 with the AWS Nitro System.
- **Flexibility:** Optimize performance and cost with flexible options like AWS Graviton-based instances, Amazon EC2 Spot instances, and AWS Savings Plans.
- **Ease of migration:** Migrate and build apps with ease using AWS Migration Tools, AWS Managed Services, or Amazon Lightsail. Learn how AWS can help.

We offer four different ways to buy instances, each with their own cost benefits:

- **On-Demand Instances:** On-Demand Instances let you pay for compute capacity by the hour with no long-term commitments. This frees you from the costs and complexities of planning, purchasing, and maintaining hardware and transforms what are commonly large, fixed costs into much smaller variable costs. On-Demand Instances also remove the need to buy “safety net” capacity to handle periodic traffic spikes.
- **Reserved Instances:** A Reserved Instance provides you with a significant discount (up to 75%) compared to On-Demand Instance pricing. There are three Reserved Instance payment options—No Upfront, Partial Upfront, and All Upfront—that enable you to balance the amount you pay up front with your effective hourly price. The Reserved Instance Marketplace is also available, which provides you with the opportunity to sell Reserved Instances if your needs change (e.g., want to move instances to a new AWS Region, change to a new instance type, or sell capacity for projects that end before your Reserved Instance term expires).
- **Spot Instances:** Spot Instances allow customers to bid on unused Amazon EC2 capacity and run those instances for as long as their bid exceeds the current Spot Price. The Spot Price changes periodically based on supply and demand, and customers whose bids meet or exceed it gain access to the available Spot Instances. If you have flexibility in when your applications can run, Spot Instances can significantly lower your Amazon EC2 costs.
- **Savings Plans:** Savings Plans is a flexible pricing model that provides savings of up to 72% on your AWS compute usage. This pricing model offers lower prices on Amazon EC2 instances usage, regardless of instance family, size, OS, tenancy or AWS Region, and also applies to AWS Fargate and AWS Lambda usage. Savings Plans offer significant savings over On Demand, just like EC2 Reserved Instances, in exchange for a commitment to use a specific amount of compute power (measured in $/hour) for a one or three year period. You can sign up for Savings Plans for a 1- or 3-year term and easily manage your plans by taking advantage of recommendations, performance reporting and budget alerts in the AWS Cost Explorer.

26.1.3. Instance types

Amazon EC2 provides a wide selection of instance types optimised to fit different use cases. Instance types comprise varying combinations of CPU, memory, storage, and networking capacity and give you the flexibility to choose the appropriate mix of resources for your applications. Each instance type includes one or more instance sizes, allowing you to scale your resources to the requirements of your target workload. We have provided details of some of the instance types available below.

26.1.3.1. General Purpose

A1 - Amazon EC2 A1 instances deliver significant cost savings and are ideally suited for scale-out and Arm-based workloads that are supported by the extensive Arm ecosystem. A1 instances are the first EC2 instances powered by AWS Graviton Processors that feature 64-bit Arm Neoverse cores and custom silicon designed by AWS.

- Custom built AWS Graviton Processor with 64-bit Arm Neoverse cores
- Support for Enhanced Networking with Up to 10 Gbps of Network bandwidth
- EBS-optimized by default
- Powered by the AWS Nitro System, a combination of dedicated hardware and lightweight hypervisor

T3 - T3 instances are the next generation burstable general-purpose instance type that provide a baseline level of CPU performance with the ability to burst CPU usage at any time for as long as required. T3 instances offer a balance of compute, memory, and network resources and are designed for applications with moderate CPU usage that experience temporary spikes in use.

T3 instances accumulate CPU credits when a workload is operating below baseline threshold. Each earned CPU credit provides the T3 instance the opportunity to burst with the performance of a full CPU core for one minute when needed. T3 instances can burst at any time for as long as required in Unlimited mode.

- Burstable CPU, governed by CPU Credits, and consistent baseline performance
- Unlimited mode by default to ensure performance during peak periods and Standard mode option for a predictable monthly cost
- Powered by the AWS Nitro System, a combination of dedicated hardware and lightweight hypervisor
- AWS Nitro System and high frequency Intel Xeon Scalable processors result in up to a 30% price performance improvement over T2 instances

T3a - T3a instances are the next generation burstable general-purpose instance type that provide a baseline level of CPU performance with the ability to burst CPU usage at any time for as long as required. T3a instances offer a balance of compute, memory, and network resources and are designed for applications with moderate CPU usage that experience temporary spikes in use. T3a instances deliver up to 10% cost savings over comparable instance types.

T3a instances accumulate CPU credits when a workload is operating below baseline threshold. Each earned CPU credit provides the T3a instance the opportunity to burst with the performance of a full CPU core for one minute when needed. T3a instances can burst at any time for as long as required in Unlimited mode.

- AMD EPYC 7000 series processors with an all core turbo clock speed of 2.5 GHz
- Burstable CPU, governed by CPU Credits, and consistent baseline performance
- Unlimited mode by default to ensure performance during peak periods and Standard mode option for a predictable monthly cost
- Powered by the AWS Nitro System, a combination of dedicated hardware and lightweight hypervisor

T2 – T2 instances are Burstable Performance Instances that provide a baseline level of CPU performance with the ability to burst above the baseline. T2 Unlimited instances can sustain high CPU performance for as long as a workload needs it. For most general-purpose workloads, T2 Unlimited instances will provide ample performance without any additional charges.

The baseline performance and ability to burst are governed by CPU Credits. T2 instances receive CPU Credits continuously at a set rate depending on the instance size, accumulating CPU Credits when they are idle, and consuming CPU credits when they are active. T2 instances are a good choice for a variety of general-purpose workloads including micro-services, low-latency interactive applications, small and medium databases, virtual desktops, development, build and stage environments, code repositories, and product prototypes. For more information see Burstable Performance Instances. Features include:

- High frequency Intel Xeon processors
- Burstable CPU, governed by CPU Credits, and consistent baseline performance
- Lowest-cost general purpose instance type, and Free Tier eligible*
- Balance of compute, memory, and network resources

M6g - Amazon EC2 M6g instances are powered by Arm-based AWS Graviton2 processors. They deliver up to 40% better price/performance over current generation M5 instances and offer a balance of compute, memory, and networking resources for a broad set of workloads.

- Custom built AWS Graviton2 Processor with 64-bit Arm Neoverse cores
- Support for Enhanced Networking with Up to 25 Gbps of Network bandwidth
- EBS-optimized by default
- Powered by the AWS Nitro System, a combination of dedicated hardware and lightweight hypervisor

M5 - M5 instances are the latest generation of General Purpose Instances. This family provides a balance of compute, memory, and network resources, and it is a good choice for many applications. Features include:

- 2.5 GHz Intel Xeon® Platinum 8175 processors with new Intel Advanced Vector Extension (AXV-512) instruction set
- New larger instance size, m5.24xlarge, offering 96 vCPUs and 384 GiB of memory
- EBS-optimized by default and higher EBS performance on smaller instance sizes
- Up to 25 Gbps network bandwidth using Enhanced Networking
- Requires HVM AMIs that include drivers for ENA and NVMe
- Powered by the new light-weight Nitro system, a combination of dedicated hardware and lightweight hypervisor

M5a - M5a instances are the latest generation of General Purpose Instances powered by AMD EPYC 7000 series processors. M5a instances deliver up to 10% cost savings over comparable instance types.

- AMD EPYC 7000 series processors with an all core turbo clock speed of 2.5 GHz
- Up to 20 Gbps network bandwidth using Enhanced Networking

- Requires HVM AMIs that include drivers for ENA and NVMe
- Powered by the AWS Nitro System, a combination of dedicated hardware and lightweight hypervisor
- Instance storage offered via EBS or NVMe SSDs that are physically attached to the host server
- With M5ad instances, local NVMe-based SSDs are physically connected to the host server and provide block-level storage that is coupled to the lifetime of the M5a instance

M5n - M5 instances are ideal for workloads that require a balance of compute, memory, and networking resources including web and application servers, small and mid-sized databases, cluster computing, gaming servers, and caching fleet. The higher bandwidth, M5n and M5dn, instance variants are ideal for applications that can take advantage of improved network throughput and packet rate performance.

- 2nd generation Intel Xeon Scalable Processors (Cascade Lake) with a sustained all-core Turbo CPU frequency of 3.1 GHz and maximum single core turbo frequency of 3.5 GHz
- Support for the new Intel Vector Neural Network Instructions (AVX-512 VNNI) which will help speed up typical machine learning operations like convolution, and automatically improve inference performance over a wide range of deep learning workloads
- 25 Gbps of peak bandwidth on smaller instance sizes
- 100 Gbps of network bandwidth on the largest instance size
- Requires HVM AMIs that include drivers for ENA and NVMe
- Powered by the AWS Nitro System, a combination of dedicated hardware and lightweight hypervisor
- Instance storage offered via EBS or NVMe SSDs that are physically attached to the host server
- With M5dn instances, local NVMe-based SSDs are physically connected to the host server and provide block-level storage that is coupled to the lifetime of the M5 instance

M4 – M4 instances provide a balance of compute, memory, and network resources, and it is a good choice for many applications. Features include:

- 2.3 GHz Intel Xeon® E5-2686 v4 (Broadwell) processors or 2.4 GHz Intel Xeon® E5-2676 v3 (Haswell) processors
- EBS-optimized by default at no additional cost
- Support for Enhanced Networking
- Balance of compute, memory, and network resources

26.1.3.2. Compute-Optimised

C5 - C5 instances are optimized for compute-intensive workloads and deliver very cost-effective high performance at a low price per compute ratio. Features include:

- C5 instances offer a choice of processors based on the size of the instance.

- New C5 and C5d 12xlarge, 24xlarge, and metal instance sizes feature custom 2nd generation Intel Xeon Scalable Processors (Cascade Lake) with a sustained all core Turbo frequency of 3.6GHz and single core turbo frequency of up to 3.9GHz.
- Other C5 instance sizes will launch on the 2nd generation Intel Xeon Scalable Processors (Cascade Lake) or 1st generation Intel Xeon Platinum 8000 series (Skylake-SP) processor with a sustained all core Turbo frequency of up to 3.4GHz, and single core turbo frequency of up to 3.5 GHz.
- New larger 24xlarge instance size offering 96 vCPUs, 192 GiB of memory, and optional 3.6TB local NVMe-based SSDs
- Requires HVM AMIs that include drivers for ENA and NVMe
- With C5d instances, local NVMe-based SSDs are physically connected to the host server and provide block-level storage that is coupled to the lifetime of the C5 instance
- Elastic Network Adapter (ENA) provides C5 instances with up to 25 Gbps of network bandwidth and up to 14 Gbps of dedicated bandwidth to Amazon EBS.
- Powered by the AWS Nitro System, a combination of dedicated hardware and lightweight hypervisor

C5n - C5n instances are ideal for high compute applications (including High Performance Computing (HPC) workloads, data lakes, and network appliances such as firewalls and routers) that can take advantage of improved network throughput and packet rate performance. C5n instances offers up to 100 Gbps network bandwidth and increased memory over comparable C5 instances. C5n.18xlarge instances support Elastic Fabric Adapter (EFA), a network interface for Amazon EC2 instances that enables customers to run applications requiring high levels of inter-node communications, like High Performance Computing (HPC) applications using the Message Passing Interface (MPI), at scale on AWS.

- 3.0 GHz Intel Xeon Platinum processors with Intel Advanced Vector Extension 512 (AVX-512) instruction set
- Run each core at up to 3.5 GHz using Intel Turbo Boost Technology
- Larger instance size, c5n.18xlarge, offering 72 vCPUs and 192 GiB of memory
- Requires HVM AMIs that include drivers for ENA and NVMe
- Network bandwidth increases to up to 100 Gbps, delivering increased performance for network intensive applications.
- EFA support on c5n.18xlarge instances
- 33% higher memory footprint compared to C5 instances
- Powered by the AWS Nitro System, a combination of dedicated hardware and lightweight hypervisor

C4 – C4 instances are optimized for compute-intensive workloads and deliver very cost-effective high performance at a low price per compute ratio. Features include:

- High frequency Intel Xeon E5-2666 v3 (Haswell) processors optimized specifically for EC2
- Default EBS-optimized for increased storage performance at no additional cost

- Higher networking performance with Enhanced Networking supporting Intel 82599 VF
- Requires Amazon VPC, Amazon EBS and 64-bit HVM AMIs

26.1.3.3. Memory Optimised

R5 - R5 instances deliver 5% additional memory per vCPU than R4 and the largest size provides 768 GiB of memory. In addition, R5 instances deliver a 10% price per GiB improvement and a ~20% increased CPU performance over R4.

- Up to 3.1 GHz Intel Xeon® Platinum 8175 processors with new Intel Advanced Vector Extension (AVX-512) instruction set
- Up to 768 GiB of memory per instance
- Powered by the AWS Nitro System, a combination of dedicated hardware and lightweight hypervisor
- With R5d instances, local NVMe-based SSDs are physically connected to the host server and provide block-level storage that is coupled to the lifetime of the R5 instance
- New 8xlarge and 16xlarge sizes now available.

R5a - R5a instances are the latest generation of Memory Optimized instances ideal for memory-bound workloads and are powered by AMD EPYC 7000 series processors. R5a instances deliver up to 10% lower cost per GiB memory over comparable instances.

- AMD EPYC 7000 series processors with an all core turbo clock speed of 2.5 GHz
- Up to 20 Gbps network bandwidth using Enhanced Networking
- Up to 768 GiB of memory per instance
- Powered by the AWS Nitro System, a combination of dedicated hardware and lightweight hypervisor
- Instance storage offered via EBS or NVMe SSDs that are physically attached to the host server
- With R5ad instances, local NVMe-based SSDs are physically connected to the host server and provide block-level storage that is coupled to the lifetime of the R5a instance

R5n - R5 instances are ideal for memory-bound workloads including high performance databases, distributed web scale in-memory caches, mid-sized in-memory database, real time big data analytics, and other enterprise applications. The higher bandwidth, R5n and R5dn, instance variants are ideal for applications that can take advantage of improved network throughput and packet rate performance.

- 2nd generation Intel Xeon Scalable Processors (Cascade Lake) with a sustained all-core Turbo CPU frequency of 3.1 GHz and maximum single core turbo frequency of 3.5 GHz
- Support for the new Intel Vector Neural Network Instructions (AVX-512 VNNI) which will help speed up typical machine learning operations like convolution, and automatically improve inference performance over a wide range of deep learning workloads
- 25 Gbps of peak bandwidth on smaller instance sizes
- 100 Gbps of network bandwidth on the largest instance size
- Requires HVM AMIs that include drivers for ENA and NVMe

- Powered by the AWS Nitro System, a combination of dedicated hardware and lightweight hypervisor
- Instance storage offered via EBS or NVMe SSDs that are physically attached to the host server
- With R5dn instances, local NVMe-based SSDs are physically connected to the host server and provide block-level storage that is coupled to the lifetime of the R5 instance

R4 - R4 instances are optimised for memory-intensive applications and offer better price per GiB of RAM than R3. Features include:

- High Frequency Intel Xeon E5-2686 v4 (Broadwell) processors
- DDR4 Memory
- Support for Enhanced Networking

X1e - X1e instances are optimized for high-performance databases, in-memory databases and other memory intensive enterprise applications. X1e instances offer one of the lowest price per GiB of RAM among Amazon EC2 instance types. Features include:

- High frequency Intel Xeon E7-8880 v3 (Haswell) processors
- One of the lowest price per GiB of RAM
- Up to 3,904 GiB of DRAM-based instance memory
- SSD storage and EBS-optimized by default and at no additional cost
- Ability to control processor C-state and P-state configurations on x1e.32xlarge, x1e.16xlarge and x1e.8xlarge instances

X1 - X1 instances are optimised for large-scale, enterprise-class and in-memory applications, and offer one of the lowest price per GiB of RAM among Amazon EC2 instance types. Features include:

- High frequency Intel Xeon E7-8880 v3 (Haswell) processors
- One of the lowest price per GiB of RAM
- Up to 1,952 GiB of DRAM-based instance memory
- SSD storage and EBS-optimized by default and at no additional cost
- Ability to control processor C-state and P-state configuration

High Memory - High memory instances are purpose built to run large in-memory databases, including production deployments of SAP HANA, in the cloud.

- 6, 9, 12, 18, and 24 TiB of instance memory, the largest of any EC2 instance
- Powered by the AWS Nitro System, a combination of dedicated hardware and lightweight hypervisor
- Bare metal performance with direct access to host hardware
- EBS-optimized by default at no additional cost
- Available in Amazon Virtual Private Clouds (VPCs)

z1d - Amazon EC2 z1d instances offer both high compute capacity and a high memory footprint. High frequency z1d instances deliver a sustained all core frequency of up to 4.0 GHz, the fastest of any cloud instance.

- A custom Intel® Xeon® Scalable processor with a sustained all core frequency of up to 4.0 GHz
- Up to 1.8TB of instance storage
- High memory with up to 384 GiB of RAM
- Powered by the AWS Nitro System, a combination of dedicated hardware and lightweight hypervisor
- With z1d instances, local NVMe-based SSDs are physically connected to the host server and provide block-level storage that is coupled to the lifetime of the z1d instance

26.1.3.4. Accelerated Computing Instances

P3 - P3 instances are the latest generation of general purpose GPU instances. Features include:

- Up to 8 NVIDIA Tesla V100 GPUs, each pairing 5,120 CUDA Cores and 640 Tensor Cores
- High frequency Intel Xeon E5-2686 v4 (Broadwell) processors for p3.2xlarge, p3.8xlarge, and p3.16xlarge.
- High frequency 2.5 GHz (base) Intel Xeon P-8175M processors for p3dn.24xlarge.
- Supports NVLink for peer-to-peer GPU communication
- Provides up to 100 Gbps of aggregate network bandwidth.
- EFA support on p3dn.24xlarge instances

P2 - P2 instances are intended for general-purpose GPU compute applications. Features include:

- High frequency Intel Xeon E5-2686 v4 (Broadwell) processors
- High-performance NVIDIA K80 GPUs, each with 2,496 parallel processing cores and 12GiB of GPU memory
- Supports GPUDirect™ for peer-to-peer GPU communications
- Provides Enhanced Networking using the Amazon EC2 Elastic Network
- Adaptor with up to 20Gbps of aggregate network bandwidth within a Placement Group
- Amazon EBS-optimised by default at no additional cost

Inf1 - Amazon EC2 Inf1 instances are built from the ground up to support machine learning inference applications.

- Up to 16 AWS Inferentia Chips
- AWS Neuron SDK
- High frequency 2nd Gen Intel® Xeon® Scalable processors
- Up to 100 Gbps networking

G4 - G4 instances are designed to help accelerate machine learning inference and graphics-intensive workloads.

- 2nd Generation Intel Xeon Scalable (Cascade Lake) processors
- NVIDIA T4 Tensor Core GPUs
- Up to 100 Gbps of networking throughput
- Up to 1.8 TB of local NVMe storage

G3 - G3 instances are optimized for graphics-intensive applications. Features include:

- High frequency Intel Xeon E5-2686 v4 (Broadwell) processors
- NVIDIA Tesla M60 GPUs, each with 2048 parallel processing cores and 8 GiB of video memory
- Enables NVIDIA GRID Virtual Workstation features, including support for 4 monitors with resolutions up to 4096x2160. Each GPU included in your instance is licensed for one “Concurrent Connected User"
- Enables NVIDIA GRID Virtual Application capabilities for application virtualization software like Citrix XenApp Essentials and VMware Horizon, supporting up to 25 concurrent users per GPU
- Each GPU features an on-board hardware video encoder designed to support up to 10 H.265 (HEVC) 1080p30 streams and up to 18 H.264 1080p30 streams, enabling low-latency frame capture and encoding, and high-quality interactive streaming experiences
- Enhanced Networking using the Elastic Network Adapter (ENA) with 25 Gbps of aggregate network bandwidth within a Placement Group

F1 - F1 instances offer customizable hardware acceleration with field programmable gate arrays (FPGAs). Features include:

Instances Features:

- High frequency Intel Xeon E5-2686 v4 (Broadwell) processors
- NVMe SSD Storage
- Support for Enhanced Networking

FPGA Features:

- Xilinx Virtex UltraScale+ VU9P FPGAs
- 64 GiB of ECC-protected memory on 4x DDR4 o Dedicated PCI-Express x16 interface
- Approximately 2.5 million logic elements
- Approximately 6,800 Digital Signal Processing (DSP) engines
- FPGA Developer AMI

26.1.3.5. Storage-Optimised

I3 - I3 is instance family provides Non-Volatile Memory Express (NVMe) SSD-backed Instance storage optimized for low latency, very high random I/O performance, high sequential read throughput and provide high IOPS at a low cost. Features include:

- High Frequency Intel Xeon E5-2686 v4 (Broadwell) Processors with base frequency of 2.3 GHz
- Up to 25 Gbps of network bandwidth using Elastic Network Adapter (ENA)based Enhanced Networking
- High Random I/O performance and High Sequential Read throughput

I3en - This instance family provides dense Non-Volatile Memory Express (NVMe) SSD instance storage optimized for low latency, high random I/O performance, high sequential disk throughput, and offers the lowest price per GB of SSD instance storage on Amazon EC2. I3en also offers Bare Metal instances (i3en.metal), powered by the Nitro System, for non-virtualized workloads, workloads that benefit from access to physical resources, or workloads that may have license restrictions.

- Up to 60 TB of NVMe SSD instance storage
- Up to 100 Gbps of network bandwidth using Elastic Network Adapter (ENA)-based Enhanced Networking
- High random I/O performance and high sequential disk throughput
- Up to 3.1 GHz Intel® Xeon® Scalable (Skylake) processors with new Intel Advanced Vector Extension (AVX-512) instruction set
- Powered by the AWS Nitro System, a combination of dedicated hardware and lightweight hypervisor
- Support bare metal instance size for workloads that benefit from direct access to physical processor and memory
- Support for Elastic Fabric Adapter on i3en.24xlarge

D2 – D2 instances feature up to 48 TB of HDD-based local storage, deliver high disk throughput, and offer the lowest price per disk throughput performance on Amazon EC2. Features include:

- High-frequency Intel Xeon E5-2676 v3 (Haswell) processors
- HDD storage
- Consistent high performance at launch time
- High disk throughput
- Support for Enhanced Networking

H1 - H1 instances feature up to 16 TB of HDD-based local storage, deliver high disk throughput, and a balance of compute and memory. Features include:

- Powered by 2.3 GHz Intel® Xeon® E5 2686 v4 processors (codenamed Broadwell)
- Up to 16TB of HDD storage

- High disk throughput
- ENA enabled Enhanced Networking up to 25 Gbps

26.1.3.6. Previous Generation Instances

AWS offers Previous Generation Instances for users who have optimized their applications around these instances and have yet to upgrade. Previous Generation Instances are still fully supported and retain the same features and functionality.

Previous Generation Instances are available through the AWS Management Console, AWS CLI, and EC2 API tools. For more information, see Previous Generation Instances.

26.1.3.7. Instance Features

Amazon EC2 instances provide a number of additional features to help you deploy, manage, and scale your applications.

- **Burstable Performance Instances –** Amazon EC2 allows you to choose between Fixed Performance Instances (e.g. M3, C3, and R3) and Burstable Performance Instances (e.g. T3). Burstable Performance Instances provide a baseline level of CPU performance with the ability to burst above the baseline.

  T Unlimited instances can sustain high CPU performance for as long as a workload needs it. For most general-purpose workloads, T Unlimited instances will provide ample performance without any additional charges. The hourly T instance price automatically covers all interim spikes in usage when the average CPU utilization of a T instance is at or less than the baseline over a 24-hour window. If the instance needs to run at higher CPU utilization for a prolonged period, it can do so at a flat additional charge of 5 cents per vCPU-hour.

  T instances' baseline performance and ability to burst are governed by CPU Credits. Each T instance receives CPU Credits continuously, the rate of which depends on the instance size. T instances accrue CPU Credits when they are idle, and use CPU credits when they are active. A CPU Credit provides the performance of a full CPU core for one minute.

  For example, a t2.small instance receives credits continuously at a rate of 12 CPU Credits per hour. This capability provides baseline performance equivalent to 20% of a CPU core (20% x 60 mins = 12 mins). If the instance does not use the credits it receives, they are stored in its CPU Credit balance up to a maximum of 288 CPU Credits. When the t2.small instance needs to burst to more than 20% of a core, it draws from its CPU Credit balance to handle this surge automatically.

  With T2 Unlimited enabled, the t2.small instance can burst above the baseline even after its CPU Credit balance is drawn down to zero. For a vast majority of general purpose workloads where the average CPU utilization is at or below the baseline performance, the basic hourly price for t2.small covers all CPU bursts. If the instance happens to run at an average 25% CPU utilization (5% above baseline) over a period of 24 hours after its CPU Credit balance is drawn to zero, it will be charged an additional 6 cents (5 cents/vCPU-hour x 1 vCPU x 5% x 24 hours).

  Many applications such as web servers, developer environments and small databases don't need consistently high levels of CPU, but benefit significantly from having full access to very fast CPUs when they need them. T2 instances are engineered specifically for these use cases. If you need consistently high CPU performance for

applications such as video encoding, high volume websites or HPC applications, we recommend you use Fixed Performance Instances. T2 instances are designed to perform as if they have dedicated high speed Intel cores available when your application really needs CPU performance, while protecting you from the variable performance or other common side-effects you might typically see from oversubscription in other environments.

- **Bare Metal Instances –** Amazon EC2 bare metal instances provide your applications with direct access to the Intel® Xeon® Scalable processor and memory resources of the underlying server. Bare metal instances come up with up to 448 vCPU (u-xxtb1.metal), and up to 24TB RAM (u-24tb1.metal).

  These instances are ideal for workloads that require access to the hardware feature set (such as Intel® VT-x), for applications that need to run in non-virtualized environments for licensing or support requirements, or for customers who wish to use their own hypervisor. Bare metal instances allow EC2 customers to run applications that benefit from deep performance analysis tools, specialized workloads that require direct access to bare metal infrastructure, legacy workloads not supported in virtual environments, and licensing-restricted Tier 1 business critical applications. Bare metal instances also make it possible for customers to run virtualization secured containers such as Clear Linux Containers. Workloads on bare metal instances continue to take advantage of all the comprehensive services and features of the AWS Cloud, such as Amazon Elastic Block Store (EBS), Elastic Load Balancer (ELB) and Amazon Virtual Private Cloud (VPC).

- **Multiple Storage Options –** Amazon EC2 allows you to choose between multiple storage options based on your requirements. Amazon EBS is a durable, block-level storage volume that you can attach to a single, running Amazon EC2 instance. You can use Amazon EBS as a primary storage device for data that requires frequent and granular updates. For example, Amazon EBS is the recommended storage option when you run a database on Amazon EC2. Amazon EBS volumes persist independently from the running life of an Amazon EC2 instance. Once a volume is attached to an instance you can use it like any other physical hard drive. Amazon EBS provides three volume types to best meet the needs of your workloads: General Purpose (SSD), Provisioned IOPS (SSD), and Magnetic. General Purpose (SSD) is the new, SSD-backed, general purpose EBS volume type that we recommend as the default choice for customers. General Purpose (SSD) volumes are suitable for a broad range of workloads, including small to medium sized databases, development and test environments, and boot volumes. Provisioned IOPS (SSD) volumes offer storage with consistent and low latency performance, and are designed for I/O intensive applications such as large relational or NoSQL databases. Magnetic volumes provide the lowest cost per gigabyte of all EBS volume types. Magnetic volumes are ideal for workloads where data is accessed infrequently, and applications where the lowest storage cost is important.

  Many Amazon EC2 instances can also include storage from disks that are physically attached to the host computer. This disk storage is referred to as instance store. Instance store provides temporary block-level storage for Amazon EC2 instances. The data on an instance store volume persists only during the life of the associated Amazon EC2 instance.

  In addition to block level storage via Amazon EBS or instance store, you can also use Amazon S3 for highly durable, highly available object storage. Learn more about Amazon EC2 storage options from the Amazon EC2 documentation.

- **EBS-Optimised Instances –** For an additional, low, hourly fee, customers can launch selected Amazon EC2 instances types as EBS-optimized instances. For C5, C4, M5, M4, P3, P2, G3, and D2 instances, this feature is enabled by default at no additional cost. EBS-optimized instances enable EC2 instances to fully use the IOPS provisioned on an EBS volume. EBS-optimized instances deliver dedicated throughput between Amazon EC2 and Amazon EBS, with options between 500 and 4,000 Megabits per second (Mbps) depending on the instance type used. The dedicated throughput minimizes contention between Amazon EBS I/O and other traffic from your EC2 instance, providing the best performance for your EBS volumes. EBS-optimized instances are designed for use with both Standard and Provisioned IOPS Amazon EBS volumes. When attached to EBS-optimized instances, Provisioned IOPS volumes can achieve single digit millisecond latencies and are designed to deliver within 10% of the provisioned IOPS performance 99.9% of the time. We recommend using Provisioned IOPS volumes with EBS-optimized instances or instances that support cluster networking for applications with high storage I/O requirements.
- **Cluster Networking –** Select EC2 instances support cluster networking when launched into a common cluster placement group. A cluster placement group provides low latency networking between all instances in the cluster. The bandwidth an EC2 instance can utilize depends on the instance type and its networking performance specification. Inter instance traffic within the same region can utilize up to 5 Gbps for single-flow and up to 25 Gbps for multi-flow traffic in each direction (full duplex). Traffic to and from S3 buckets in the same region can also utilize all available instance aggregate bandwidth. When launched in a placement group, instances can utilize up to 10 Gbps for single flow traffic and up to 25 Gbps for multi-flow traffic. Network traffic to the Internet is limited to 5 Gbps (full duplex). Cluster networking is ideal for high performance analytics systems and many science and engineering applications, especially those using the MPI library standard for parallel programming.
- **Dedicated Instances –** Dedicated Instances are Amazon EC2 instances that run in a VPC on hardware that's dedicated to a single customer. They are ideal for workloads where corporate policies or industry regulations require that your Amazon EC2 instances be physically isolated at host hardware level from instances that belong to other AWS accounts. Dedicated Instances let you take full advantage of the benefits of the AWS Cloud: on-demand elastic provisioning, pay only for what you use, all while ensuring that your Amazon EC2 compute instances are isolated at the hardware level.

  You can also use Dedicated Hosts to launch Amazon EC2 instances on physical servers that are dedicated for your use. Dedicated Hosts give you additional visibility and control over how instances are placed on a physical server, and you can reliably use the same physical server over time. As a result, Dedicated Hosts enable you to use your existing server-bound software licenses like Windows Server and address corporate compliance and regulatory requirements. Visit this page to compare Dedicated Instances and Dedicated Hosts.

26.2. Backup/Restore and Disaster Recovery

This service allows for backup and recovery. The service can back up machine images. Users control this via a manual or scheduled API call. Users schedule and recover backups through a web interface.

26.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

26.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/ec2/index.html
- **Service quotas:** https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ec2-resource-limits.html
- **Service FAQs:** https://aws.amazon.com/ec2/faqs/

26.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/ec2/index.html and the following links for comprehensive technical documentation regarding this service.

- User Guide for Linux Instances: Describes key concepts of Amazon EC2 and provides instructions for using the features of Amazon EC2.
- User Guide for Windows Instances: Describes key concepts for Amazon EC2 and provides instructions for launching and using your Windows instance.
- User Guide for AWS Nitro Enclaves: Describes key concepts for AWS Nitro Enclaves and provides instructions for using enclaves.
- API Reference: Documents the Amazon EC2 Query API.

27. Amazon Elastic Container Registry (ECR)

27.1. Service Overview

Amazon Elastic Container Registry (Amazon ECR) is a fully managed Docker container registry that makes it easy for developers to store, manage, and deploy Docker container images.

27.1.1. Features

- **Amazon container orchestrator integration:** Amazon Elastic Container Registry (Amazon ECR) is integrated with Amazon Elastic Container Service (Amazon ECS) and Amazon Elastic Kubernetes Service (Amazon EKS), which means you can easily store and run container images for applications with either orchestrator. All you need to do is specify the Amazon ECR repository in your task or pod definition for Amazon ECS or Amazon EKS to retrieve the appropriate images for your applications.
- **OCI and Docker support:** Amazon ECR supports Open Container Initiative (OCI) standards and the Docker Registry HTTP API V2. This allows you to use Docker CLI commands (e.g., push, pull, list, tag) or your preferred Docker tools to interact with Amazon ECR, maintaining your existing development workflow. You can easily access

Amazon ECR from any Docker environment, whether in the cloud, on-premises, or on your local machine. Amazon ECR lets you store Docker container images and related OCI artefacts in your repositories.

- **Public container image and artefact gallery:** You can discover and use container software that vendors, open source projects, and community developers share publicly in the Amazon ECR public gallery. Popular base images such as operating systems, AWS-published images, Kubernetes add-ons, and files, such as Helm charts, can be found in the gallery. You don't need to use an AWS account to search or pull a public image, however using your account makes it easier and faster to use public container software.
- **High availability and durability:** Amazon ECR stores your container images and artefacts in Amazon S3. Amazon S3 is designed for 99.999999999% (11 9's) of data durability because it automatically creates and stores copies of all S3 objects across multiple systems. This means that your data is available when needed and protected against failures, errors, and threats. Amazon ECR can also automatically replicate your data to multiple AWS Regions for your high availability applications.
- **Team and public collaboration:** Amazon ECR supports the ability to define and organize repositories in your registry using namespaces. This allows you to organize your repositories based on your team's existing workflows. You can set which API actions another user may perform on your repository (e.g., create, list, describe, delete, and get) through resource-level policies, allowing you to share your repositories easily with different users and AWS accounts. You can easily share your container artefacts with anyone in the world by storing them in a public repository.
- **Access control:** Amazon ECR uses AWS Identity and Access Management (IAM) to control and monitor who and what (e.g., EC2 instances) can access your container images. Through IAM, you can define policies to allow users within the same AWS account or other accounts to access your container images in private repositories. You can also further refine these policies by specifying different permissions for different users and roles (e.g., push, pull, or full administrator access). Anyone in the world can access your container images stored in public repositories for worldwide collaboration.
- **Encryption:** You can transfer your container images to and from Amazon ECR via HTTPS. Your images are also automatically encrypted at rest using Amazon S3 server-side encryption. Amazon ECR also lets you choose your own key managed by AWS Key Management Service (AWS KMS) to encrypt images at rest.
- **Third-party integrations:** Amazon ECR is integrated with third-party developer tools. You can integrate Amazon ECR into your continuous integration and delivery process allowing you to maintain your existing development workflow. Learn more about our third-party integration on our Partners page.
- **Pull through cache repositories:** With Amazon ECR's pull through cache repositories, you can retrieve, store, and sync container artefacts stored in publicly accessible container registries. It offers the high download rates that you need and the availability, security, and scale that you've come to depend on. With frequent registry syncs and no additional tools to manage, pull through cache repositories help you keep container images sourced from public registries up to date.

27.1.2. Benefits

- **Ready to use:** Push container images to Amazon ECR without installing or scaling infrastructure, and pull images using any management tool.
- **Secure:** Share and download images securely over Hypertext Transfer Protocol Secure (HTTPS) with automatic encryption and access controls.
- **Efficient:** Access and distribute your images faster, reduce download times, and improve availability using a scalable, durable architecture.

27.2. Backup/Restore and Disaster Recovery

This requirement is not applicable for this service. For additional information beyond what is described herein, please refer to http://aws.amazon.com/documentation/.

27.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

27.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/ecr/
- **Service quotas:** https://docs.aws.amazon.com/AmazonECR/latest/userguide/service-quotas.html
- **Service FAQs:** https://aws.amazon.com/ecr/faqs/

27.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/ecr/ and the following links for comprehensive technical documentation regarding this service.

- User Guide: Describes key concepts of Amazon ECR and provides instructions for using the features of Amazon ECR.
- API Reference: Describes all the API operations for managing your private registry and private repositories on Amazon ECR.

28. Amazon Elastic Container Service (ECS)

28.1. Service Overview

Amazon Elastic Container Service (Amazon ECS) allows you to easily deploy containerized workloads on AWS. The powerful simplicity of Amazon ECS enables you to grow from a single Docker container to managing your entire enterprise application portfolio. Run and scale your container workloads across availability zones, in the cloud, and on-premises, without the complexity of managing a control plane or nodes.

28.1.1. Features

- **Serverless by default with AWS Fargate:** AWS Fargate is built into Amazon ECS, which means you no longer have to worry about managing servers, handling capacity planning, or figuring out how to isolate container workloads for security. Just define your application's requirements, select Fargate as your launch type in the console or

Command Line Interface (CLI), and Fargate takes care of all the scaling and infrastructure management required to run your containers.

- **Amazon ECS Anywhere:** With ECS Anywhere, you can use the same familiar Amazon ECS console and operator tools to manage your on-premises container workloads for a consistent experience across your container-based applications. The AWS Systems Manager (SSM) integration automatically and securely establishes trust between your on-premises hardware and the AWS control plane.
- **Security and isolation by design:** Amazon ECS natively integrates with the Security, Identity, and Management and Governance tools you already trust, which helps you get to production quickly and successfully. You can assign granular permissions for each of your containers, giving you a high level of isolation when building your applications. Launch your containers with the security and compliance levels you have come to expect from AWS.
- **Autonomous control plane operations:** Amazon ECS is a fully-managed container orchestration service, with AWS configuration and operational best practices built-in, and no control plane, nodes, or add-ons for you to manage. It natively integrates with both AWS and third-party tools to make it easier for teams to focus on building the applications, not the environment.
- **Docker Support:** Amazon ECS supports Docker and enables you to run and manage Docker containers. It even integrates into the Docker Compose CLI, so you can define and run multi-container applications. Applications you package locally as a container will deploy and run on Amazon ECS without the need for any configuration changes.
- **Windows Containers Compatibility:** Amazon ECS supports management of Windows containers. An Amazon ECS-optimized Windows Amazon Machine Image (AMI) provides enhanced instance and container launch time performance and visibility into CPU, memory utilization, and reservation metrics.
- **AWS Copilot:** The AWS Copilot CLI is a tool for developers to build, release, and operate production ready containerized applications on Amazon ECS and AWS Fargate. Copilot takes best practices, from infrastructure to continuous delivery, and makes them available to customers from the comfort of their command line. You can also monitor the health of your service by viewing your service's status or logs, scale up or down production services, and spin up a new environment for automated testing. Download AWS Copilot.
- **Repository Support:** Amazon ECS can be used with any third-party hosted Docker image repository or accessible private Docker registry, such as Docker Hub and Amazon Elastic Container Registry (ECR). All you need to do is specify the repository in your task definition and Amazon ECS retrieves the appropriate images for your applications.
- **Task Definitions:** Amazon ECS allows you to define tasks through a JavaScript Object Notation (JSON) template called a Task Definition. Within a Task Definition, you can specify one or more containers that are required for your task, including the Docker repository and image, memory and CPU requirements, shared data volumes, and how the containers are linked to each other. You can launch as many tasks as you want from a single Task Definition file that you can register with the service. Task Definition files also give you version control over your application specification.

- **Programmatic Control:** Amazon ECS provides you with a set of simple API actions to allow you to integrate and extend the service. The API actions allow you to create and delete clusters, register and deregister tasks, launch, and terminate Docker containers, and provide detailed information about the state of your cluster and its instances. You can also use AWS CloudFormation to provision Amazon ECS clusters, register task definitions, and schedule containers.
- **Container Deployments:** Amazon ECS allows you to easily update your containers to new versions. You can upload a new version of your application task definition, and the Amazon ECS scheduler automatically starts new containers using the updated image and stop containers running the previous version. Amazon ECS automatically registers and deregisters your containers from the associated Application Load Balancer.
- **Blue/Green Deployments:** Blue/green deployments with AWS CodeDeploy help you minimize downtime during application updates. You can launch a new version of your Amazon ECS service alongside the old version and test the new version before you reroute traffic. You can also monitor the deployment process and rapidly rollback if there is an issue.
- **Container Auto-Recovery:** The Amazon ECS will automatically recover unhealthy containers to ensure that you have the desired number of containers supporting your application.
- **Capacity Providers:** Capacity Providers allow you to define flexible rules for how containerized workloads run on different types of compute capacity, and manage the scaling of the capacity. Capacity Providers work with both Amazon Elastic Compute Cloud (Amazon EC2) and AWS Fargate. When running tasks and services, you can split them across multiple Capacity Providers, enabling new capabilities such as running a service in a predefined split percentage across Fargate and Fargate Spot.
- **Storage:** Amazon Elastic File System (Amazon EFS) is a simple, scalable, fully managed elastic file system, enabling you to build modern applications, and persist and share data and state, from your Amazon ECS and AWS Fargate deployments. All aspects of using Amazon EFS with containers, including connectivity, is cared for, zero management required. You can simply focus on your applications, not infrastructure.
- **Task Scheduling:** Amazon ECS task scheduling allows you to run processes that perform work and then stop, such as batch processing jobs. Task scheduling starts tasks automatically from a queue of jobs, or based on a time interval that you define.
- **Service Scheduling:** Amazon ECS service scheduling allows you to run stateless services and applications. This scheduling strategy ensures that a specified number of tasks are constantly running and restarts tasks if failure occurs. Customers can ensure that tasks are registered against an Elastic Load Balancing load balancer and can perform health checks that users define for running tasks.
- **Daemon Scheduling:** Amazon ECS daemon scheduling automatically runs the same task on each selected instance in your ECS cluster. This makes it easy to run tasks that provide common management functionality for a service like logging, monitoring, or backups.
- **Task Placement:** Amazon ECS allows users to customize how tasks are placed onto a cluster of Amazon EC2 instances based on built-in attributes such as instance type, Availability Zone, or user-defined custom attributes. Use attributes such as environment

= production to label resources, list API actions to find those resources, and the RunTask and CreateService API actions to schedule tasks on those resources. With Amazon ECS, use placement strategies such as bin pack and spread to further define where tasks are placed. Policies can be chained together to achieve sophisticated placement capabilities without writing any code.

- **Service Discovery:** Amazon ECS is integrated with AWS Cloud Map to make it easy for your containerized services to discover and connect with each other. AWS Cloud Map is a cloud resource discovery service that lets you define custom names for your application resources. It increases your application availability because your web service will always discover the most up-to-date locations of these dynamically changing resources.
- **Service Mesh:** Service mesh makes it easy to build and run complex microservices applications by standardizing how every microservice in the application communicates. AWS App Mesh is a service that makes it easy to configure part of your application for end-to-end visibility and high-availability. To use App Mesh, add the Envoy proxy image to the ECS task definition. App Mesh manages Envoy configuration to provide service mesh capabilities. App Mesh exports metrics, logs, and traces to the endpoints specified in the Envoy bootstrap configuration provided. App Mesh provides an API to configure traffic routes, circuit breaking, retries, and other controls between microservices that are mesh-enabled.
- **Task Networking:** Amazon ECS supports Docker networking and integrates with Amazon VPC to provide isolation for containers. This gives you control over how containers connect with other services and external traffic. With Amazon ECS, you can choose between four networking modes for your containers that cater towards different use cases:
- **Load Balancing:** Amazon ECS is integrated with Elastic Load Balancing, allowing you to distribute traffic across your containers using Application Load Balancers or Network Load Balancers. You specify the task definition and the load balancer to use, and Amazon ECS automatically adds and removes containers from the load balancer. Specify a dynamic port in the task definition, which gives your container an unused port when it is scheduled on an EC2 instance. In addition, use path-based routing to share a load balancer with multiple services.
- **Monitoring:** Amazon ECS provides monitoring capabilities for your containers and clusters through Amazon CloudWatch. You can monitor average and aggregate CPU and memory utilization of running tasks as grouped by task definition, service, or cluster. Set CloudWatch alarms to alert you when your containers or clusters need to scale up or down.
- **Logging:** Amazon ECS allows you to record all your Amazon ECS API calls and have the log files delivered to you through AWS CloudTrail. The recorded information includes the identity of the API caller, the time of the API call, the source IP address of the API caller, the request parameters, and the response elements returned by Amazon ECS. CloudTrail provides you a history of API calls made from the AWS Management Console, AWS SDKs, and AWS CLI. It enables security analysis, resource change tracking, and compliance auditing.
- **AWS Config:** AWS Config integrates with Amazon ECS to provide you visibility into your configuration of AWS resources in your AWS account. AWS Config allows users to

monitor and track how resources were configured, how they relate to one another, and how the configurations and relationships change over time. AWS Config enables you to simplify compliance and security, operational troubleshooting, and resource administration.

28.1.2. Benefits

- **CI/CD and automation:** Launch thousands of containers across the cloud using your preferred continuous integration and delivery (CI/CD) and automation tools.
- **Serverless:** Optimize your time with AWS Fargate serverless compute for containers, which eliminates the need to configure and manage control plane, nodes, and instances.
- **Value:** Save up to 50 percent on compute costs with autonomous provisioning, auto-scaling, and pay-as-you-go pricing.
- **Integration and compliance:** Integrate seamlessly with AWS management and governance solutions, standardized for compliance with virtually every regulatory agency around the globe.

28.2. Backup/Restore and Disaster Recovery

This service allows for backup and recovery. The service can back up machine images. Users control this via a manual or scheduled API call. Users schedule and recover backups through a web interface.

28.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

28.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/ecs/
- **Service quotas:** https://docs.aws.amazon.com/AmazonECS/latest/developerguide/service-quotas.html
- **Service FAQs:** https://aws.amazon.com/ecs/faqs/

28.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/ecs/ and the following links for comprehensive technical documentation regarding this service.

- Developer Guide: Describes key concepts of Amazon ECS and provides instructions for using the features of Amazon ECS.
- User Guide for AWS Fargate: Describes key concepts of Amazon ECS on AWS Fargate and provides instructions for launching containers on the serverless infrastructure provided by Fargate.
- Best Practices Guide: Provides recommendations and best practices for building and managing your Amazon ECS based applications.

29. Amazon Elastic Kubernetes Service (EKS)

29.1. Service Overview

Amazon Elastic Kubernetes Service (Amazon EKS) is a managed Kubernetes service that makes it easy for you to run Kubernetes on AWS and on-premises. Kubernetes is an open-source system for automating deployment, scaling, and management of containerized applications. Amazon EKS is certified Kubernetes-conformant, so existing applications that run on upstream Kubernetes are compatible with Amazon EKS.

Amazon EKS automatically manages the availability and scalability of the Kubernetes control plane nodes responsible for scheduling containers, managing application availability, storing cluster data, and other key tasks.

Amazon EKS lets you run your Kubernetes applications on both Amazon Elastic Compute Cloud (Amazon EC2) and AWS Fargate. With Amazon EKS, you can take advantage of all the performance, scale, reliability, and availability of AWS infrastructure, as well as integrations with AWS networking and security services, such as application load balancers (ALBs) for load distribution, AWS Identity and Access Management (IAM) integration with role-based access control (RBAC), and AWS Virtual Private Cloud (VPC) support for pod networking.

29.1.1. Features

- **Managed Control Plane:** Amazon EKS provides a scalable and highly-available Kubernetes control plane running across multiple AWS Availability Zones (AZs). Amazon EKS automatically manages availability and scalability of Kubernetes API servers and etcd persistence layer. Amazon EKS runs the Kubernetes control plane across three AZs to ensure high availability, and automatically detects and replaces unhealthy control plane nodes.
- **Service Integrations:** AWS Controllers for Kubernetes (ACK) gives you direct management control over AWS services from within your Kubernetes environment. ACK makes it simple to build scalable and highly available Kubernetes applications utilizing AWS services.
- **Hosted Kubernetes Console:** EKS provides an integrated console for Kubernetes clusters. Cluster operators and application developers can use EKS as a single place to organize, visualize, and troubleshoot your Kubernetes applications running on Amazon EKS. The EKS console is hosted by AWS and is available automatically for all EKS clusters.
- **EKS Add-Ons:** EKS add-ons are common operational software for extending the Kubernetes operational functionality. You can use EKS to install and keep the add-on software up-to-date. When you start an Amazon EKS cluster, select the add-ons you would like to run in the cluster, including Kubernetes tools for observability, networking, auto-scaling, and AWS service integrations.
- **Managed Node Groups:** Amazon EKS lets you create, update, scale, and terminate nodes for your cluster with a single command. These nodes can also leverage Amazon EC2 Spot Instances to reduce costs. Managed node groups run Amazon EC2 instances using the latest EKS-optimized or custom Amazon Machine Images (AMIs) in your AWS account, while updates and terminations gracefully drain nodes to ensure your applications remain available.

- **Use eksctl for launching nodes and single line management:** Use the eksctl command-line tool to get up and running with Amazon EKS in minutes. Simply run an "eksctl create cluster" command to create your EKS cluster. You can use eksctl to simplify cluster management and operations including managing nodes and add ons.
- **Windows Support:** Amazon EKS supports Windows worker nodes and Windows container scheduling. EKS supports running Windows worker nodes alongside Linux worker nodes, allowing you to use the same cluster for managing applications on either operating system.
- **ARM Support:** AWS Graviton2 processors power Arm-based EC2 instances, delivering a major leap in performance and capabilities as well as significant cost savings. Improving application cost efficiency is a primary goal of running containers. Combine both, and you get great price performance. For example, workload testing shows instance types based on Graviton2 processors deliver up to 40% better price performance than their equivalent x86-based M5, C5, and R5 families. Amazon EKS on AWS Graviton2 is generally available where both services are available Regionally.
- **Networking and Security:** Amazon EKS makes it easy to provide security for your Kubernetes clusters, with advanced features and integrations to AWS services and technology partner solutions. For example, IAM provides fine-grained access control and Amazon VPC isolates your Kubernetes clusters from other customers.
- **Support for IPv6:** Amazon Elastic Kubernetes Service (EKS) supports IPv6, enabling customers to scale containerized applications on Kubernetes far beyond limits of private IPv4 address space. With EKS support for IPv6, pods are assigned only a globally routable IPv6 address, allowing you to scale applications in your cluster without consuming limited private IPv4 address space.
- **Service Discovery:** AWS Cloud Map is a cloud resource discovery service. With Cloud Map, you can define custom names and maintain updated locations of dynamically changing application resources. This increases your application availability, because your web service always discovers the most up-to-date resource locations.
- **Service Mesh:** Service mesh standardizes how every microservice within your application communicates, making it easy to build and run complex microservices applications. AWS App Mesh configures your application for end-to-end visibility and high-availability. You can use the AWS App Mesh controller for Kubernetes to create new services connected to the mesh, define traffic routing, and configure security features like encryption. Additionally, App Mesh allows you to automatically register your Kubernetes pods in AWS Cloud Map for service discovery.
- **VPC Native Networking:** Your EKS clusters run in an Amazon VPC, allowing you to use your own VPC security groups and network access control lists (ACLs). No compute resources are shared with other customers, which provides you a high level of isolation to build secure and reliable applications.
- **AWS IAM Authenticator:** Amazon EKS integrates Kubernetes RBAC (the native role based access control system for Kubernetes) with AWS IAM. You can assign RBAC roles directly to each IAM entity, allowing granular access permission control over your Kubernetes control plane nodes.
- **IAM for Service Accounts:** Amazon EKS allows you to assign IAM permissions to your Kubernetes service accounts. The IAM role can control access to other containerized

services, AWS resources external to the cluster such as databases and secrets, or third-party services and applications running outside of AWS. This gives you fine-grained, pod-level access control when running clusters with multiple co-located services while simplifying cluster availability and cost optimization.

- **Serverless Compute:** EKS supports AWS Fargate to run your Kubernetes applications using serverless compute. Fargate removes the need to provision and manage servers, lets you specify and pay for resources per application, and improves security through application isolation by design.
- **Certified Conformant:** Amazon EKS runs upstream Kubernetes and is certified Kubernetes-conformant, so you can use all the existing plug-ins and tooling from the Kubernetes community. Applications running on Amazon EKS are fully compatible with applications running on any standard Kubernetes environment, whether running in on-premises data centres or public clouds. This means that you can easily migrate any standard Kubernetes application to Amazon EKS without refactoring your code.
- **Managed Cluster Updates:** Amazon EKS makes it easy to update running clusters to the latest Kubernetes version without managing the update process. Kubernetes version updates are done in place, removing the need to create new clusters or migrate applications to a new cluster. As new Kubernetes versions are released and validated for use with Amazon EKS, we will support three stable Kubernetes versions at any given time as part of the update process. You can initiate new version installation and review in-flight update status via the SDK, CLI or AWS Console.
- **Advanced Workload Support:** Amazon EKS provides an optimized Amazon Machine Image (AMI) that includes configured NVIDIA drivers for GPU-enabled P2 and P3 Amazon EC2 instances. This makes it easy to use Amazon EKS to run computationally advanced workloads, including machine learning (ML), Kubeflow, deep learning (DL) containers, high performance computing (HPC), financial analytics, and video transcoding.
- **Open-Source Compatibility:** Amazon EKS is fully compatible with Kubernetes community tools and supports popular Kubernetes add-ons. These include CoreDNS, which creates a DNS service for your cluster, and both the Kubernetes Dashboard web-based UI and the kubectl command line tool, which help access and manage your cluster on Amazon EKS.
- **EKS Connector:** Amazon EKS allows you to connect any conformant Kubernetes cluster to AWS and visualize it in the Amazon EKS console. You can connect any conformant Kubernetes cluster, including Amazon EKS Anywhere clusters running on-premises, self-managed clusters on Amazon Elastic Compute Cloud (Amazon EC2), and other Kubernetes clusters running outside of AWS. Regardless where your cluster is running, you can use the Amazon EKS console to view all connected clusters and the Kubernetes resources running on them.

29.1.2. Benefits

- **Flexibility:** Explore multiple ways to configure VPC, ALB, EC2 Kubernetes worker nodes, and Amazon EKS.
- **Value:** Reduce costs with efficient compute resource provisioning and automatic Kubernetes application scaling.

- **Security:** Ensure a more secure Kubernetes environment with security patches automatically applied to your cluster's control plane.

29.2. Backup/Restore and Disaster Recovery

This service allows for backup and recovery. The service can back up machine images. Users control this via a manual or scheduled API call. Users schedule and recover backups through a web interface.

29.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

29.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/eks/index.html
- **Service quotas:** https://docs.aws.amazon.com/general/latest/gr/eks.html
- **Service FAQs:** https://aws.amazon.com/eks/faqs/

29.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/eks/index.html and the following links for comprehensive technical documentation regarding this service.

- User Guide: Describes key concepts of Amazon EKS and provides instructions for using the features of Amazon EKS.
- API Reference: Documents the Amazon EKS API.
- Amazon EKS best practices guides: Describes best practices for using Amazon EKS.

30. Amazon Elastic File System (EFS)

30.1. Service Overview

Amazon Elastic File System (Amazon EFS) provides a simple, serverless, set-and-forget elastic file system that lets you share file data without provisioning or managing storage. It can be used with AWS services and on-premises resources and is built to scale on demand to petabytes without disrupting applications.

Amazon EFS is well suited to support a broad spectrum of use cases from home directories to business-critical applications. Use cases include storage for containerized and serverless applications, big data analytics, web serving and content management, application development and testing, media and entertainment workflows, and database backups.

30.1.1. Features

- **Fully managed:** Amazon EFS is a fully managed service providing NFS shared file system storage for Linux workloads. EFS makes it simple to create and configure file systems. You don't have to worry about managing file servers or storage, updating hardware, configuring software, or performing backups. In seconds, you can create a

fully managed file system by using the AWS Management Console, the AWS Command Line Interface (CLI), or an AWS SDK.

- **Highly available and durable:** Amazon EFS is designed to be highly available and is designed for 99.999999999% (11 nines) durability. By default, every EFS file system object (such as directory, file, and link) is redundantly stored across multiple Availability Zones (AZs) for file systems using Standard storage classes.
- **Storage classes and lifecycle management:** Amazon EFS offers Standard and One Zone storage classes for both frequently accessed and infrequently accessed files. The Standard and One Zone storage classes are performance-optimized to deliver consistent low latencies. The Amazon EFS Standard-Infrequent Access (EFS Standard-IA) and Amazon EFS One Zone-Infrequent Access (EFS One Zone-IA) storage classes are cost-optimized for files accessed less frequently. You can start saving on your storage costs by simply enabling EFS Lifecycle Management for your file system and choosing an age-off policy (7, 14, 30, 60, or 90 days).
- **Scalable performance:** Amazon EFS is designed to provide the throughput, IOPS, and low latency needed for a broad range of workloads. Throughput and IOPS scale as a file system grows and can burst to higher throughput levels for short periods of time to support the unpredictable performance needs of file workloads. For the most demanding workloads, Amazon EFS can support performance over 10 GB/second and over 500,000 IOPS.
- **Shared file system with NFS v4.0 and v4.1 support:** Amazon EFS provides secure access for thousands of connections for Amazon Elastic Compute Cloud (EC2) instances, AWS container and serverless compute services, and on-premises servers simultaneously using a traditional file permissions model, file locking, and hierarchical directory structure through the NFS v4 protocol. Amazon EC2 instances can access your file system across AZs and Regions while on-premises servers can access it through AWS Direct Connect or AWS VPN.
- **Performance modes:** Amazon EFS is designed to provide the throughput, IOPS, and low latency needed for a broad range of workloads and offers two performance modes: General Purpose and Max I/O.
- **Throughput modes:** Amazon EFS offers two throughput modes: Bursting and Provisioned. The throughput mode helps determine the overall throughput a file system can achieve. With Bursting Throughput, the throughput scales with the size of the file system, dynamically bursting as needed to support the spiky nature of many file-based workloads. Provisioned Throughput is designed to support applications that require higher dedicated throughput than the default Bursting mode and can be configured independently of the amount of data stored on the file system.
- **Elastic and scalable:** With Amazon EFS, storage capacity is elastic, growing and shrinking automatically as you add and remove files, dynamically providing the storage capacity to applications as they need it. Since capacity is elastic, no provisioning is necessary, and you'll be billed only for what you use.
- **Encryption:** Amazon EFS offers encryption for data at rest and data in transit, providing a comprehensive encryption solution to secure both types.
- **Containers and serverless file storage:** Amazon EFS is integrated with containers and serverless compute services from AWS that require shared storage for latency-sensitive

and IOPS-heavy workloads at any scale. In a single step, EFS provides applications running on Amazon Elastic Container Service (ECS), Amazon Elastic Kubernetes Service (EKS), AWS Fargate, and AWS Lambda with access to shared file systems for stateful workloads.

30.1.2. Benefits

- **Quick and simple:** Create and configure shared file systems simply and quickly for AWS compute services—no provisioning, deploying, patching, or maintenance required.
- **Scalable:** Scale your file system automatically as files are added, removed, and burst to higher throughput levels when necessary.
- **Cost efficient:** Pay only for the storage you use and reduce costs up to 92 percent by automatically moving infrequently accessed files.
- **Secure and reliable:** Securely and reliably access your files with a fully managed file system designed for high availability and 99.999999999 percent (11 9s) durability.

30.2. Backup/Restore and Disaster Recovery

This service allows for backup and recovery. The service can back up files on EFS. Users control this via a manual or scheduled API call. Users schedule and recover backups through a web interface.

30.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

30.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/efs/
- **Service quotas:** https://docs.aws.amazon.com/efs/latest/ug/limits.html
- **Service FAQs:** https://aws.amazon.com/efs/faq/

30.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/efs/ and the following links for comprehensive technical documentation regarding this service.

- User Guide: Walks through how to set up Amazon EFS, move data into it, and integrate it with other services. Includes the API reference.

31. Amazon Elastic Inference

31.1. Service Overview

Amazon Elastic Inference allows you to attach low-cost GPU-powered acceleration to Amazon EC2 and Sagemaker instances or Amazon ECS tasks, to reduce the cost of running deep learning inference by up to 75%. Amazon Elastic Inference supports TensorFlow, Apache MXNet, PyTorch and ONNX models.

Inference is the process of making predictions using a trained model. In deep learning applications, inference accounts for up to 90% of total operational costs for two reasons. Firstly,

standalone GPU instances are typically designed for model training - not for inference. While training jobs batch process hundreds of data samples in parallel, inference jobs usually process a single input in real time, and thus consume a small amount of GPU compute. This makes standalone GPU inference cost-inefficient. On the other hand, standalone CPU instances are not specialized for matrix operations, and thus are often too slow for deep learning inference. Secondly, different models have different CPU, GPU, and memory requirements. Optimizing for one resource can lead to underutilization of other resources and higher costs.

Amazon Elastic Inference solves these problems by allowing you to attach just the right amount of GPU-powered inference acceleration to any EC2 or SageMaker instance type or ECS task, with no code changes. With Amazon Elastic Inference, you can choose any CPU instance in AWS that is best suited to the overall compute and memory needs of your application, and then separately configure the right amount of GPU-powered inference acceleration, allowing you to efficiently utilize resources and reduce costs.

31.1.1. Features

- **Integrated with Amazon SageMaker, Amazon EC2, and Amazon ECS:** There are multiple ways to run inference workloads on AWS: deploy your model on Amazon SageMaker for a fully managed experience, or run it on Amazon EC2 instances or Amazon ECS tasks and manage it yourself. Amazon Elastic Inference is integrated to work seamlessly with Amazon SageMaker, Amazon EC2, and Amazon ECS, allowing you to add inference acceleration in all scenarios. You can specify the desired amount of inference acceleration when you create your model's HTTPS endpoint in Amazon SageMaker, when you launch your Amazon EC2 instance, and when you define your Amazon ECS task.
- **TensorFlow, Apache MXNet and PyTorch support:** Amazon Elastic Inference is designed to be used with AWS's enhanced versions of TensorFlow Serving, Apache MXNet and PyTorch. These enhancements enable the frameworks to automatically detect the presence of inference accelerators, optimally distribute the model operations between the accelerator's GPU and the instance's CPU, and securely control access to your accelerators using AWS Identity and Access Management (IAM) policies. The enhanced TensorFlow Serving, MXNet and PyTorch libraries are provided automatically in Amazon SageMaker, AWS Deep Learning AMIs, and AWS Deep Learning Containers, so you don't have to make any code change to deploy your models in production. You can also download them separately by following the instructions here.
- **Open Neural Network Exchange (ONNX) format support:** ONNX is an open format that makes it possible to train a model in one deep learning framework and then transfer it to another for inference. This allows you to take advantage of the relative strengths of different frameworks. ONNX is integrated into PyTorch, MXNet, Chainer, Caffe2, and Microsoft Cognitive Toolkit, and there are connectors for many other frameworks including TensorFlow. To use ONNX models with Amazon Elastic Inference, your trained models need to be transferred to the AWS-optimized version of Apache MXNet for production deployment.
- **Choice of single or mixed precision operations:** Amazon Elastic Inference accelerators support both single-precision (32-bit floating point) operations and mixed precision (16-bit floating point) operations. Single precision provides an extremely large numerical range to represent the parameters used by your model. However, most models don't actually need this much precision and calculating numbers that large results in unnecessary loss of performance. To avoid that problem, mixed-precision

operations allow you to reduce the numerical range by half to gain up to 8x greater inference performance.

- **Available in multiple amounts of acceleration:** Amazon Elastic Inference is available in multiple throughput sizes ranging from 1 to 32 trillion floating point operations per second (TFLOPS) per accelerator, making it efficient for accelerating a wide range of inference models including computer vision, natural language processing, and speech recognition. Compared to standalone Amazon EC2 P3 instances that start at 125 TFLOPS (the smallest P3 instance available), Amazon Elastic Inference starts at a single TFLOPS per accelerator. This allows you to scale up inference acceleration in more appropriate increments. You can also select from larger accelerator sizes, up to 32 TFLOPS per accelerator, for more complex models.
- **Auto-scaling:** Amazon Elastic Inference can be part of the same Amazon EC2 Auto Scaling group you use to scale your Amazon SageMaker, Amazon EC2, and Amazon ECS instances. When EC2 Auto Scaling adds more EC2 instances to meet the demands of your application, it also scales up the accelerator attached to each instance. Similarly, when Auto Scaling reduces your EC2 instances as demand goes down, it also scales down the attached accelerator for each instance. This makes it easy to scale your inference acceleration alongside your application's compute capacity to meet the demands of your application.

31.1.2. Benefits

- **Reduce inference costs by up to 75%:** Amazon Elastic Inference allows you to choose the instance type that is best suited to the overall compute and memory needs of your application. You can then separately specify the amount of inference acceleration that you need. This reduces inference costs by up to 75% because you no longer need to over-provision GPU compute for inference.
- **Get exactly what you need:** Amazon Elastic Inference can provide as little as a single-precision TFLOPS (trillion floating point operations per second) of inference acceleration or as much as 32 mixed-precision TFLOPS. This is a much more appropriate range of inference compute than the range of up to 1,000 TFLOPS provided by a standalone Amazon EC2 P3 instance. For example, a simple language processing model might require only one TFLOPS to run inference well, while a sophisticated computer vision model might need up to 32 TFLOPS.
- **Respond to changes in demand:** You can easily scale the amount of inference acceleration up and down using Amazon EC2 Auto Scaling groups to meet the demands of your application without over-provisioning capacity. When EC2 Auto Scaling increases your EC2 instances to meet increasing demand, it also automatically scales up the attached accelerator for each instance. Similarly, when it reduces your EC2 instances as demand goes down, it also automatically scales down the attached accelerator for each instance. This helps you pay only for what you need when you need it.

31.2. Backup/Restore and Disaster Recovery

This requirement is not applicable for this service. For additional information beyond what is described herein, please refer to http://aws.amazon.com/documentation/.

31.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

31.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/elastic-inference.html
- **Service quotas:** https://docs.aws.amazon.com/general/latest/gr/elastic-inference.html
- **Service FAQs:** https://aws.amazon.com/machine-learning/elastic-inference/faqs/

31.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/elastic-inference.html comprehensive technical documentation regarding this service.

32. Amazon Elastic MapReduce (EMR)

32.1. Service Overview

Amazon EMR is a cloud big data platform for running large-scale distributed data processing jobs, interactive SQL queries, and machine learning (ML) applications using open-source analytics frameworks such as Apache Spark, Apache Hive, and Presto.

32.1.1. Features

- **Easy to use:** Amazon EMR simplifies building and operating big data environments and applications. Related EMR features include easy provisioning, managed scaling, and reconfiguring of clusters, and EMR Studio for collaborative development.
- **Elastic:** Amazon EMR enables you to quickly and easily provision as much capacity as you need, and automatically or manually add and remove capacity. This is very useful if you have variable or unpredictable processing requirements.
- **Low cost:** Amazon EMR is designed to reduce the cost of processing large amounts of data. Some of the features that make it low cost include low per-second pricing, Amazon EC2 Spot integration, Amazon EC2 Reserved Instance integration, elasticity, and Amazon S3 integration.
- **Flexible data stores:** With Amazon EMR, you can leverage multiple data stores, including Amazon S3, the Hadoop Distributed File System (HDFS), and Amazon DynamoDB.
- **Use your favourite open source applications:** With versioned releases on Amazon EMR, you can easily select and use the latest open source projects on your EMR cluster, including applications in the Apache Spark and Hadoop ecosystems.
- **Big Data Tools:** Amazon EMR supports powerful and proven Hadoop tools such as Apache Spark, Apache Hive, Presto, and Apache HBase. Data scientists use EMR to run deep learning and machine learning tools such as TensorFlow, Apache MXNet, and, using bootstrap actions, add use case-specific tools and libraries. Data analysts use EMR Studio, Hue and EMR Notebooks for interactive development, authoring Apache

Spark jobs, and submitting SQL queries to Apache Hive and Presto. Data Engineers use EMR for data pipeline development and data processing, and use Apache Hudi to simplify incremental data management and data privacy use cases requiring record-level insert, updates, and delete operations.

- **Data access control:** By default, Amazon EMR application processes use EC2 instance profile when they call other AWS services. For multi-tenant clusters, Amazon EMR offers three options to manage user access to Amazon S3 data.
- **Select the right instance for your cluster:** You choose what types of EC2 instances to provision in your cluster (standard, high memory, high CPU, high I/O, etc.) based on your application's requirements. You have root access to every instance and you can fully customize your cluster to suit your requirements. Learn more about supported Amazon EC2 Instance Types. Amazon EMR now provides up to 30% lower cost and up to 15% improved performance for Spark workloads on Graviton2-based instances.
- **Debug your applications:** When you enable debugging on a cluster, Amazon EMR archives the log files to Amazon S3 and then indexes those files. You can then use a graphical interface in the console to browse the logs and view job history in an intuitive way.
- **Deep learning:** Use popular deep learning frameworks like Apache MXNet to define, train, and deploy deep neural networks. You can use these frameworks on Amazon EMR clusters with GPU instances.
- **Install additional software:** You can use bootstrap actions or a custom Amazon Machine Image (AMI) running Amazon Linux to install additional software on your cluster. Bootstrap actions are scripts that are run on the cluster nodes when Amazon EMR launches the cluster. They run before Hadoop starts and before the node begins processing data. You can also preload and use software on a custom Amazon Linux AMI. Learn more about Amazon EMR Bootstrap Actions and custom Amazon Linux AMIs.
- **Efficiently copy data:** You can quickly move large amounts of data from Amazon S3 to HDFS, from HDFS to Amazon S3, and between Amazon S3 buckets using Amazon EMR's S3DistCp, an extension of the open source tool Distcp, which uses MapReduce to efficiently move large amounts of data.
- **Custom JAR:** Write a Java program, compile against the version of Hadoop you want to use, and upload to Amazon S3. You can then submit Hadoop jobs to the cluster using the Hadoop JobClient interface.

32.1.2. Benefits

- **Efficient:** Run big data applications and petabyte-scale data analytics faster, and at less than half the cost of on-premises solutions.
- **Broad support:** Build applications using the latest open-source frameworks, with options to run on customized Amazon EC2 clusters, Amazon EKS, AWS Outposts, or Amazon EMR Serverless.
- **Fast:** Get up to 2X faster time-to-insights with performance-optimized and open-source API-compatible versions of Spark, Hive, and Presto.
- **Tools:** Easily develop, visualize, and debug your applications using EMR Notebooks and familiar open-source tools in EMR Studio.

32.2. Backup/Restore and Disaster Recovery

This service allows for backup and recovery. The service can back up block volumes. Users control this via a manual or scheduled API call. Users schedule and recover backups through a web interface.

32.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

32.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/emr/
- **Service quotas:** https://docs.aws.amazon.com/general/latest/gr/emr.html
- **Service FAQs:** https://aws.amazon.com/emr/faqs/

32.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/emr/ and the following links for comprehensive technical documentation regarding this service.

- Management Guide: Describes key concepts of Amazon EMR and provides instructions for using the service.
- Release Guide: Provides information about Amazon EMR releases, including installed cluster software such as Hadoop and Spark.

33. Amazon ElastiCache

33.1. Service Overview

Amazon ElastiCache is a fully managed, in-memory caching service supporting flexible, real-time use cases. You can use ElastiCache for caching, which accelerates application and database performance, or as a primary data store for use cases that don't require durability like session stores, gaming leaderboards, streaming, and analytics. ElastiCache is compatible with Redis and Memcached.

33.1.1. Features

- **Pay for what you use:** Amazon ElastiCache has no upfront costs. With on-demand nodes you pay only for the resources you consume by the hour without any long-term commitments. With Reserved Nodes, you can make a low, one-time, up-front payment for each node you wish to reserve for a 1 or 3 year term. In return, you receive a significant discount off the ongoing hourly usage rate for the Node(s) you reserve.
- **Improve latency and throughput:** Amazon ElastiCache can be used to significantly improve latency and throughput for many read-heavy application workloads (such as social networking, gaming, media sharing and Q&A portals) or compute-intensive workloads (such as a recommendation engine) by allowing you to store the objects that are often read in cache. Moreover, with Redis support for advanced data structures, you can augment the database tier to provide features (such as leaderboard, counting,

session and tracking) that are not easily achievable via databases in a cost-effective way.

- **Simplify management:** Amazon ElastiCache simplifies and offloads the management, monitoring, and operation of in-memory cache environments, enabling you to focus on the differentiating parts of your applications.
- **Support for two engines:** Memcached and Redis
- **Ease of management:** With a few clicks via the AWS Management Console you can configure and launch cache nodes for the engine you wish to use.
- **Compatibility with the specific engine protocol:** This means most of the client libraries will work with the respective engines they were built for - no additional changes or tweaking required.
- **Monitoring:** Detailed monitoring statistics for the engine nodes at no extra cost via Amazon CloudWatch

33.1.2. Benefits

- **Performance:** Boost application performance, reducing latency to microseconds.
- **Scalability:** Scale with just a few clicks to meet the needs of your most demanding, internet-scale applications.
- **Value:** Reduce costs and eliminate the operational overhead of self-managed caching.
- **Choice:** Build with your choice of Redis or Memcached, two popular open-source caching technologies.

33.2. Backup/Restore and Disaster Recovery

This service allows for backup and recovery. The service can back up memory snapshots to disk. Users control this via a manual or scheduled API call. Users schedule and recover backups through a web interface.

33.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

33.4. Service Constraints

Information on service features and constraints is available online here:

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

33.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/elasticache/ and the following links for comprehensive technical documentation regarding this service.

- ElastiCache for Redis User Guide: For ElastiCache for Redis, helps you understand the components and features provided and how to use them. Learn how to access

ElastiCache for Redis through a web-based GUI, command line tools, and the ElastiCache API.

- ElastiCache for Memcached User Guide: For ElastiCache for Memcached, helps you understand the components and features provided and how to use them. Learn how to access ElastiCache for Memcached through a web-based GUI, command line tools, and the ElastiCache API.

34. Amazon EventBridge

34.1. Service Overview

Amazon EventBridge is a serverless event bus that makes it easier to build event-driven applications at scale using events generated from your applications, integrated Software-as-a-Service (SaaS) applications, and AWS services. EventBridge delivers a stream of real-time data from event sources such as Zendesk or Shopify to targets like AWS Lambda and other SaaS applications. You can set up routing rules to determine where to send your data to build application architectures that react in real-time to your data sources with event publisher and consumer completely decoupled.

34.1.1. Features

- **API Destinations:** API Destinations is a new feature for EventBridge that enables developers to send events back to many on-premises or software as a service (SaaS) applications with the ability to control throughput and authentication. Customers can send events to any web-based application with a web address without worrying about writing custom code, or using additional infrastructure. Customers can configure rules with input transformations that will map the event's format to the receiving service format, and you can use EventBridge to take care of security and delivery.
- **Archive and Replay Events:** Event Replay is a new feature for Amazon EventBridge that allows customers to reprocess past events back to an event bus or a specific EventBridge rule. This feature enables developers to debug their applications quickly, extend them by hydrating targets with historic events, and recover from errors.
- **Schema Registry:** The EventBridge schema registry stores event schema in a registry that other developers can easily search and access in your organization, so you don't have to find events and their structure manually. The registry also allows you to generate code bindings for programming languages such as Java, Python, or TypeScript directly in your IDE so that the event can be used as an object in your code. By turning on schema discovery for an event bus, the schemas of events are automatically discovered and added to the registry, removing the need to create a schema for an event manually. Schemas for all AWS services are automatically visible in your schema registry, and the schemas for integrated SaaS applications are visible when you turn on schema discovery for the SaaS partner event bus. Learn more in this blog post.
- **Fully managed and scalable event bus:** Amazon EventBridge is a serverless, fully managed, and scalable event bus that allows applications to communicate using events. There is no infrastructure to manage and no capacity to provision.
- **SaaS Integration:** Your AWS applications can take action based on events that SaaS applications generate. Amazon EventBridge is natively integrated with SaaS applications from many providers including Datadog, OneLogin, PagerDuty, Savyint, Segment, SignalFX, SugarCRM, Symantec, Whispir, and Zendesk, with additional integrations

planned. You don't need to manage any integration setup such as authentication events from your SaaS provider simply appear on your event bus.

- **Over 100 built-in event sources and targets:** Amazon EventBridge is directly integrated with over 130 event sources and over 35 targets, including AWS Lambda, Amazon SQS, Amazon SNS, AWS Step Functions, Amazon Kinesis Data Streams, Amazon Kinesis Data Firehose, with additional sources and targets planned. All mutating API calls (i.e., all calls except Describe*, List*, and Get*) across all AWS services generate events through AWS CloudTrail.
- **Decoupled event publishers and subscribers:** Amazon EventBridge makes it easy for you to build event-driven application architectures. Applications or microservices can publish events to the event bus without awareness of subscribers. Applications or microservices can subscribe to events without awareness of the publisher. You can also send events from your own applications to an event bus via the service's PutEvents API. Other applications can then receive events through any of the many supported AWS target services. This decoupling allows teams to work independently, leading to faster development and improved agility.
- **Event filtering:** You can filter events with rules. A rule matches incoming events for a given event bus and routes them to targets for processing. A single rule can route to multiple targets, all of which are processed in parallel. Rules allow different application components to look for and process the events that are of interest to them. A rule can customize an event before it is sent to the target by passing along only certain parts or by overwriting it with a constant. You can also have multiple rules that match on the same event, so different microservices or applications can choose to match events based on specific filters.
- **Reliable event delivery:** Amazon EventBridge provides at-least-once event delivery to targets, including retry with exponential backoff for up to 24 hours. Events are stored durably across multiple Availability Zones (AZs), providing additional assurance your events will be delivered to their destination. Amazon EventBridge also provides a 99.99% availability service level agreement (SLA), ensuring your applications are able to access the service reliably.
- **Automatic response to operational changes in AWS services:** Amazon EventBridge extends its predecessor, Amazon CloudWatch Events, and provides a near- real time stream of system events that describe changes to your AWS resources. It allows you to respond quickly to operational changes and take corrective action. You simply write rules to indicate which events are of interest to your application and what automated actions to take when a rule matches an event. You can, for example, set a rule to invoke an AWS Lambda function to remediate an issue, or notify an Amazon Simple Notification Service (SNS) topic to alert an operator.
- **Scheduled events:** You can set up scheduled events using the popular Unix cron syntax. Scheduled events are generated on a periodic basis and invoke any of the supported target AWS services.
- **Monitoring and auditing:** You can monitor your event bus using Amazon CloudWatch metrics, such as the number of times an event matches a rule, or the number of times a target is invoked. You can use Amazon CloudWatch Logs to store, monitor, and analyse events that are triggered in your environment. AWS CloudTrail enables you to monitor the calls made to the Amazon EventBridge API.

- **Security and compliance:** Amazon EventBridge integrates with AWS Identity and Access Management (IAM) so that you can control which users and resources have permission to access your data and how they can access it. EventBridge supports VPC endpoints and encryption in transit using TLS 1.2. Amazon EventBridge is GDPR, SOC, ISO, DoD CC SRG, and FedRamp compliant and is also HIPAA-eligible.
- **Pay per event:** Events generated by AWS services are free. You only pay for events generated by your own applications or SaaS applications.

34.1.2. Benefits

- **Build event-driven architectures:** EventBridge simplifies the process of building event-driven architectures. With EventBridge, your event targets don't need to be aware of event sources because you can filter and publish directly to EventBridge. There is no setup required. Improve developer agility as well as application resiliency with loosely coupled event-driven architectures.
- **Connect SaaS apps:** EventBridge ingests data from supported SaaS applications and routes it to AWS services and SaaS targets (through API destinations - an HTTP invocation endpoint target for events) without writing custom integration code. You can use EventBridge to connect your SaaS apps, or use events from your SaaS apps to trigger workflows for customer support, business operations, and more. Learn more about integrated SaaS partners.
- **Write less custom code:** EventBridge makes it easier to connect applications. You can ingest, filter, transform and deliver events without writing custom code. The EventBridge schema registry stores a collection of easy-to-find event schemas and enables you to download code bindings for those schemas in your IDE so you can represent events as a strongly-typed objects in your code. Automatically add schemas discovered from your event bus to the registry through the schema discovery feature.
- **Reduce operational overhead:** With EventBridge, there are no servers to provision, patch, and manage. There is no additional software to install, maintain, or operate. EventBridge automatically scales based on the number of events ingested, and you pay only for events published by your AWS or SaaS applications. EventBridge has built-in distributed availability and fault-tolerance. EventBridge also has a native event archive and replay capability that makes it easier to recover from failures or build a new application state from old events.

34.2. Backup/Restore and Disaster Recovery

This requirement is not applicable for this service. For additional information beyond what is described herein, please refer to http://aws.amazon.com/documentation/.

34.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

34.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/eventbridge/

- **Service quotas:** https://docs.aws.amazon.com/eventbridge/latest/userguide/eb-quota.html
- **Service FAQs:** https://aws.amazon.com/eventbridge/faqs/

34.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/eventbridge/ and the following links for comprehensive technical documentation regarding this service.

- User Guide: Describes key concepts of Amazon EventBridge and provides instructions for using the features of Amazon EventBridge.
- API Reference: Describes the core API operations for Amazon EventBridge in detail.

35. Amazon FinSpace

35.1. Service Overview

Amazon FinSpace is a data management and analytics service purpose-built for the financial services industry (FSI). FinSpace reduces the time you spend finding and preparing petabytes of financial data to be ready for analysis from months to minutes.

Financial services organizations analyse data from internal data stores like portfolio, actuarial, and risk management systems as well as petabytes of data from third-party data feeds, such as historical securities prices from stock exchanges. It can take months to find the right data, get permissions to access the data in a compliant way, and prepare it for analysis.

FinSpace removes the heavy lifting of building and maintaining a data management system for financial analytics. With FinSpace, you collect data and catalogue it by relevant business concepts such as asset class, risk classification, or geographic region. FinSpace makes it easy to discover and share data across your organization in accordance with your compliance requirements. You define your data access policies in one place and FinSpace enforces them while keeping audit logs to allow for compliance and activity reporting. FinSpace also includes a library of 100+ functions, like time bars and Bollinger bands, for you to prepare data for analysis.

35.1.1. Features

- **Flexible data import and secure storage:** FinSpace’s APIs allow you to ingest data from your internal and third-party data feeds, such as risk management systems or historical securities prices from stock exchanges, and existing S3 buckets. You can also drag and drop files using the web application. Your data is then encrypted in FinSpace using an AWS Key Management Service (KMS) key that you create and manage.
- **Business Data Catalog:** Data loaded into FinSpace is automatically tracked in the business data catalogue. FinSpace's business data catalogue makes it easy for you to find, learn about, and access data directly instead of relying on a technical team to provision data for analysis.
- **Metadata and classifications:** In FinSpace, you can use metadata to provide business context and meaning to your data so that it is easier to organize and understand. In order to describe your data in more structured ways, like by provider, frequency of delivery, licensing terms, and more, you can add classifications and associate attributes by writing in text or selecting classifications from premade drop downs. With FinSpace, you can also define classifications based on your business terms and usage and set rules to ensure that all metadata is captured in the same way.

- **Automatic application of corrections to time series data:** FinSpace is designed to simplify the collecting and processing of data common in the financial services industry, such as time series data and reference data. As data is periodically collected, FinSpace can apply corrections or revisions to your time series data so that no additional post-processing is required in order for you to perform analysis. FinSpace also automatically converts all collected raw data into compute optimized formats like Apache Parquet that are more conducive to analysis.
- **Financial time series data processing library:** FinSpace's financial time series data processing library allows you to transform and enrich time series data. The library of 100+ functions can be used to run financial analytics in your Jupyter notebooks, such as the computation of statistical and technical indicators to support investment and risk management decisions. It also provides functions to filter and normalize time series data such as a stock's open, high, low, and close prices into time bars.
- **Historical data views:** FinSpace makes it easy to validate modelling assumptions using historical data by enabling you to explore and use datasets at any point in time, such as when the data was originally collected.
- **Integrated Jupyter notebooks:** When you want to perform data preparation and run analysis in FinSpace, you can access all data directly from Jupyter notebooks built-in to integrate with your data. FinSpace allows you to securely share data created in your notebook with others on your team via the FinSpace catalogue.
- **Data access controls and audit reports:** FinSpace lets you define access policies in one place and enforces them across data search, visualization, and analysis. FinSpace also records access and operations that you and your systems perform on data. You can use FinSpace audit reports to demonstrate compliance with your data governance policies.
- **Fully managed service:** FinSpace eliminates the need to integrate a financial data management system with analysis tools and build all the components required to do so, such as bi-temporal storage, a data catalogue, Spark clusters, and more. Your teams only need to integrate FinSpace with your data sources to get started.
- **Managed Spark compute:** With FinSpace, you can launch managed Spark clusters from the Jupyter notebook to easily process data at scale, perform data transformations, and run analytics. You can select from one of the pre-configured cluster templates to match the size and complexity of the operation you need to perform.

35.1.2. Benefits

- **Find data with just a few clicks:** Amazon FinSpace makes it easy to store, catalogue, and manage your data according to concepts common in the financial services industry like asset class and instrument type. You can search the FinSpace catalogue using familiar phrasing like "options trades for the S&P 500" or "loan delinquency rates" to easily find the data you are looking for.
- **Get insights in minutes:** Amazon FinSpace reduces the time it takes to prepare your financial services data by providing over 100 pre-built data preparation functions, including machine learning functions, that are commonly used in the financial services industry. With FinSpace, you can maintain previous versions of your data and track how this data changes over time. You can also load your existing libraries of proprietary and

open source functions into FinSpace so you can continue to use them for your data preparation and analysis.

- **Ensure regulatory compliance:** Amazon FinSpace helps you meet your regulatory compliance requirements by enforcing data access controls and tracking data usage to generate compliance and activity reports. You define data access controls in FinSpace and they are automatically enforced.
- **Eliminate operational overhead:** Amazon FinSpace lets you focus on financial data analysis without having to worry about the operational overhead of building and maintaining the infrastructure needed for financial data management and analytics.

35.2. Backup/Restore and Disaster Recovery

This requirement is not applicable for this service. For additional information beyond what is described herein, please refer to http://aws.amazon.com/documentation/.

35.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

35.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/finspace/
- **Service quotas:** https://docs.aws.amazon.com/finspace/latest/userguide/finspace-quotas.html
- **Service FAQs:** https://aws.amazon.com/finspace/faqs/

35.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/finspace/ and the following links for comprehensive technical documentation regarding this service.

- User Guide: Provides a conceptual overview of Amazon FinSpace, and detailed information about how to configure FinSpace catalogue, load data, organize and prepare data, and perform analysis.

36. Amazon Forecast

36.1. Service Overview

Amazon Forecast uses machine learning (ML) to generate more accurate demand forecasts with just a few clicks, without requiring any prior ML experience. Amazon Forecast includes algorithms that are based on over twenty years of forecasting experience and developed expertise used by Amazon.com bringing the same technology used at Amazon to developers as a fully managed service, removing the need to manage resources. Amazon Forecast uses ML to learn not only the best algorithm for each item, but the best ensemble of algorithms for each item, automatically creating the best model for your data.

36.1.1. Features

- **Forecast Explainability:** Explore what factors, such as price, holidays or weather, are driving your forecasts with Amazon Forecast, which provides forecast Explainability

report in the form of impact scores for all your forecasts, specific time series of interest or specific time durations. Explainability provides you more insight into better managing your business operations.

- **Automatically include local weather information:** With Weather Index, Amazon Forecast can increase your forecasting accuracy by automatically ingesting local weather information in your demand forecasts with one click and at no extra cost. Weather conditions influence consumer demand patterns, product merchandizing decisions, staffing requirements, and energy consumption needs. When you use the Weather Index, Forecast trains a model with historical weather information for the locations of your operations and uses the latest 14-day weather forecasts on items that are influenced by day-to-day variations to create more accurate demand forecasts.
- **Generate probabilistic forecasts:** Unlike most other forecasting solutions that generate point forecasts, Amazon Forecast generates probabilistic forecasts at three different quantiles by default: 10%, 50% and 90%. In addition, you can choose any quantile between 1% and 99%, including the 'mean' forecast. This allows you to choose a forecast that suits your business needs depending on whether the cost of capital (over forecasting) or missing customer demand (under forecasting) is of importance.
- **Works with any historical time series data to create accurate forecasts:** Amazon Forecast can use virtually any historical time series data (e.g., price, promotions, economic performance metrics) to create accurate forecasts for your business. For example, in a retail scenario, Amazon Forecast uses machine learning to process your time series data (such as price, promotions, and store traffic) and combines that with associated data (such as product features, floor placement, and store locations) to determine the complex relationships between them. By combining time series data with additional variables, Amazon Forecast can be 50% more accurate than non-machine learning forecasting tools.
- **Easily evaluate the accuracy of your forecasting models:** Amazon Forecast provides six different comprehensive accuracy metrics to help you understand the performance of your forecasting model and compare it to previous forecasting models you've created that may have looked at a different set of variables or used a different period of time for the historical data. Amazon Forecast automatically splits your data into a training and testing set allowing you to download the forecasts it generates for the testing set for you to use a custom metric to evaluate the accuracy or allows you to create multiple backtest windows and visualize the metrics, helping you evaluate model accuracy over different start dates.
- **Integrate with your existing tools:** Amazon Forecast can be easily imported into common business and supply chain applications, such as SAP and Oracle Supply Chain. This makes it easy to integrate more accurate forecasting into your existing business processes with little to no change.

36.1.2. Benefits

- **Scale:** Scale operations by forecasting millions of items, using the same technology as Amazon.com.
- **Optimise:** Optimize inventory and reduce waste with accurate forecasts at a granular level.

- **Efficiency:** Improve capital utilization and make long-term decisions with more confidence.
- **Satisfaction:** Increase customer satisfaction with optimal staffing to meet varying demand levels.

36.2. Backup/Restore and Disaster Recovery

This requirement is not applicable for this service. For additional information beyond what is described herein, please refer to http://aws.amazon.com/documentation/.

36.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

36.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/forecast/
- **Service quotas:** https://docs.aws.amazon.com/forecast/latest/dg/limits.html
- **Service FAQs:** https://aws.amazon.com/forecast/faqs/

36.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/forecast/ and the following links for comprehensive technical documentation regarding this service.

- Developer Guide: Provides a conceptual overview of Amazon Forecast. Includes detailed instructions for using the features and provides a complete API reference for developers.

37. Amazon Fraud Detector

37.1. Service Overview

Amazon Fraud Detector is a fully managed service that makes it easy to identify potentially fraudulent online activities such as online payment fraud and fake account creation. Amazon Fraud Detector uses machine learning (ML) and 20 years of fraud detection expertise from Amazon Web Services (AWS) and Amazon.com to automatically identify potentially fraudulent activity and catch more fraud faster. With Amazon Fraud Detector, you can create a fraud detection model with just a few clicks and no prior ML experience. Amazon Fraud Detector handles all of the ML heavy lifting for you.

37.1.1. Features

- **Automated model creation:** Amazon Fraud Detector fully automates the creation of machine learning models that identify potential fraud for common online activities such as new account creations, online payments, and guest checkouts. The automated model-building process takes care of all the heavy lifting such as data validation and enrichment, feature engineering, algorithm selection, hyperparameter tuning, and model deployment. You simply upload your dataset, select the model type, and Amazon Fraud Detector automatically finds the best-fitting fraud detection ML model. No coding or previous machine learning experience is required.

- **Models that continuously learn:** Your model maintains its performance longer between retrainings because Amazon Fraud Detector automatically calculates information like account age, time since last activity, and counts of activities. This means that your model can learn the difference between trusted customers who frequently make transactions and fraudsters' continued attempts.
- **Insights into your model performance:** For each model you train, you can see all of the inputs you provided ranked by their impact on model performance. Using the importance values and relative ranking, you can gain insight into what inputs are driving your model performance.
- **Trigger rule-based actions:** Once you create an Amazon Fraud Detector fraud detection model, you can use the Amazon Fraud Detector console or application programming interface (API) to create rules based on model predictions. Customers can create rules to take actions such as accept, review, or collect more information for specific model scores. For example, you can easily create a rule to flag suspicious customer accounts for review if the model score is greater than your predetermined threshold and the account's phone number country and IP address country do not match.
- **Real-time fraud prediction API:** You can use the Amazon Fraud Detector API to perform real-time fraud predictions and evaluate online activities in your application as they occur. For example, you can call the fraud predictions API to check every new account sign-up for potential fraud risk, using your model and rules to trigger an action.
- **A single interface to review and audit your predictions and detection logic:** Using the Amazon Fraud Detector console, you can easily search and review your past fraud evaluations to audit detection logic. View event data, detection logic applied during the evaluation, and the conditions that resulted in a fraud prediction outcome.
- **Amazon SageMaker integration:** If you have already created a fraud detection model in Amazon SageMaker, you can integrate it with Amazon Fraud Detector to stop even more fraud. You can use both your Amazon SageMaker and Amazon Fraud Detector models in your application to detect different types of fraud. For example, your application can use the Amazon Fraud Detector model to assess the fraud risk of customer accounts, and simultaneously use your Amazon SageMaker model to check for account compromise risk.

37.1.2. Benefits

- **No experience needed:** Build, deploy, and manage fraud detection models without previous machine learning (ML) experience.
- **Learn more:** Gain insights from your historical data, plus 20+ years of Amazon experience, to construct an accurate, customized fraud detection model.
- **Start immediately:** Start detecting fraud immediately, easily enhance models with customized business rules, and deploy results to generate critical predictions.

37.2. Backup/Restore and Disaster Recovery

This requirement is not applicable for this service. For additional information beyond what is described herein, please refer to http://aws.amazon.com/documentation/.

37.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

37.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/frauddetector/
- **Service quotas:** https://docs.aws.amazon.com/frauddetector/latest/ug/limits.html
- **Service FAQs:** https://aws.amazon.com/fraud-detector/faqs/

37.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/frauddetector/ and the following links for comprehensive technical documentation regarding this service.

- User Guide: Provides a conceptual overview of how to use Amazon Fraud Detector.
- API Reference: Describes all the API operations for Amazon Fraud Detector, with sample requests, responses, and errors for the supported web service protocols.

38. FreeRTOS

38.1. Service Overview

FreeRTOS is an open source, real-time operating system for microcontrollers that makes small, low-power edge devices easy to program, deploy, secure, connect, and manage. Distributed freely under the MIT open source license, FreeRTOS includes a kernel and a growing set of software libraries suitable for use across industry sectors and applications. This includes securely connecting your small, low-power devices to AWS Cloud services like AWS IoT Core or to more powerful edge devices running AWS IoT Greengrass. FreeRTOS is built with an emphasis on reliability and ease of use, and offers the predictability of long term support releases.

A microcontroller contains a simple, resource-constrained processor that can be found in many devices, including appliances, sensors, fitness trackers, industrial automation, and automobiles. Many of these small devices can benefit from connecting to the cloud or locally to other devices, but have limited compute power and memory capacity and typically perform simple, functional tasks. Microcontrollers frequently run operating systems that may not have built-in functionality to connect to local networks or the cloud, making IoT applications a challenge. FreeRTOS helps solve this problem by providing the kernel to run low-power devices as well as software libraries that make it easy to securely connect to the cloud or other edge devices, so you can collect data from them for IoT applications and take action.

38.1.1. Features

- **Local connectivity:** Local connectivity to an edge device running AWS IoT Greengrass allows FreeRTOS devices to continue communicating, collecting data, and taking actions without a cloud connection. FreeRTOS devices can connect to the local network via Wi-Fi and Ethernet using local connectivity libraries such as Wi-Fi management. The Wi-Fi management library implements an abstraction layer for Wi-Fi features such as setup, configuration, provisioning, security, and power management.

- **Cloud connectivity:** Cloud connectivity allows you to easily collect data and take actions on microcontroller-based devices for use in IoT applications and with other AWS cloud services. You can connect FreeRTOS devices to AWS IoT Core using MQTT-based messaging or HTTP. MQTT is a lightweight protocol with a small footprint, enabling efficient communication for constrained, microcontroller-based devices. FreeRTOS facilitates easy onboarding with standard, vendor-independent library interfaces. MQTT is a lightweight protocol with a small footprint, enabling efficient communication for constrained, microcontroller-based devices. Cloud connectivity allows devices like smart electricity meters to send back information on consumption and analyse that data with other AWS services like AWS IoT Analytics.
- **Support for AWS IoT Core Device Shadows:** FreeRTOS also supports the AWS IoT Core Device Shadow API with a Device Shadow library. Device Shadows create a persistent, virtual version, or "shadow," of each device that includes the device's latest state so that applications or other devices can read messages and interact with the device. Microcontroller-based devices, like a temperature controlled fan, can benefit from a device shadow by saving the latest state in the cloud, such as "rotating," and then update the state to "stop," so when the device is back online, it implements the action to stop.
- **Support for AWS IoT Device Defender:** FreeRTOS provides an AWS IoT Device Defender library. The integration with AWS IoT Device Defender makes it easy to report on device-side metrics to detect anomalies when these metrics deviate from expected behaviour. AWS IoT Device Defender also continuously audits the IoT configurations associated with your FreeRTOS devices to make sure that they comply with security best practices.
- **Secure device, connection, and updates:** FreeRTOS comes with libraries for security, including secure cloud connection, certificate authentication, key management, and a code signing feature.
- **Over-the-air updates:** You can use AWS IoT Device Management with FreeRTOS devices for an integrated OTA update solution. FreeRTOS makes deploying OTA updates for microcontroller-based devices less memory intensive by communicating those updates over a single TLS connection, shared with other AWS IoT Core communications. You provide a firmware image, select the devices to update, select a code-signing method, and schedule the update, all within the AWS IoT Device Management console. You can use OTA updates to deploy security updates, bug fixes, and new firmware updates to devices in the field.
- **FreeRTOS Long Term Support:** With FreeRTOS Long Term Support (LTS) releases, you can rely on a FreeRTOS version that provides feature stability, and security updates and critical bug fixes for two years. This makes it easier to identify and include only recommended changes to the FreeRTOS kernel and libraries, without adding risk of introducing updates that could break an existing application.
- **FreeRTOS Extended Maintenance Plan:** FreeRTOS Extended Maintenance Plan (EMP) allows you to receive security patches and critical bug fixes on your chosen FreeRTOS Long Term Support (LTS) version for up to 10 years* beyond the expiry of the initial LTS period. FreeRTOS EMP can help you keep your microcontroller-based devices secure for years, save operating system upgrade costs, and reduce risks associated with patching your devices in the field.

38.1.2. Benefits

- **Reduce product liability risks:** Run firmware that receives security patches on a feature-stable codebase throughout the lifecycle of your product. A feature-stable codebase ensures that you receive security patches on the same LTS version, so you don't need to upgrade to the latest FreeRTOS version with potentially breaking changes.
- **Save operating system upgrade costs:** Continue to use FreeRTOS libraries that provide feature and API stability for the term of your subscription so you don't need to incur development, testing, and QA costs to migrate to the latest FreeRTOS release.
- **Improve device security for the long term:** Receive security patches and critical bug fixes on your chosen FreeRTOS LTS libraries to improve security of your IoT devices throughout their lifecycle.
- **Reduce the risk of delayed updates:** Updating devices with critical fixes involves project planning, release readiness testing, and over-the-air (OTA) update scheduling. Reduce delayed deployment risks by receiving timely notification of upcoming patches and support with integration issues.

38.2. Backup/Restore and Disaster Recovery

This requirement is not applicable for this service. For additional information beyond what is described herein, please refer to http://aws.amazon.com/documentation/.

38.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

38.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/freertos/
- **Service quotas:** https://docs.aws.amazon.com/general/latest/gr/amazon-freertos.html
- **Service FAQs:** https://aws.amazon.com/freertos/faqs/

38.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/freertos/ and the following links for comprehensive technical documentation regarding this service.

- User Guide: Provides detailed information about the microcontroller operating system that makes small, low-powered edge devices easy to program, deploy, secure, and maintain.

39. Amazon FSx for Lustre

39.1. Service Overview

Amazon FSx for Lustre is a fully managed service that provides cost-effective, high-performance, scalable storage for compute workloads. Powered by Lustre, the world's most popular high-performance file system, FSx for Lustre offers shared storage with sub-ms latencies, up to terabytes per second of throughput, and millions of IOPS. FSx for Lustre file systems can also be linked to Amazon Simple Storage Service (S3) buckets, allowing you to

access and process data concurrently from both a high-performance file system and from the S3 API.

39.1.1. Features

- **World's most popular high-performance file system:** The Lustre open source file system was built to solve the problem of quickly and cheaply processing the world's ever-growing data sets, and it's the most widely used file system for the 500 fastest computers in the world. It is battle-tested across a broad set of industries, from energy to life sciences, to media production to financial services, for workloads ranging from genome sequencing to video transcoding to machine learning to fraud detection.
- **Low-latency data access:** The average first-byte latency when you access file data is sub-millisecond on SSD-based file systems and single-digit millisecond on HDD-based file systems. Every Amazon FSx for Lustre file system, regardless of the deployment type, storage type, or throughput performance level, is supported by a metadata server backed by low-latency SSD storage. The SSD-based metadata server ensures that all metadata operations, which represent the majority of file system operations, are delivered with sub-millisecond latencies.
- **Simple to deploy with compute management services:** Amazon FSx for Lustre integrates with AWS Batch though EC2 Launch Templates. AWS Batch is a cloud-native batch scheduler for HPC, ML, and other asynchronous workloads. AWS Batch will automatically and dynamically size instances to job resource requirements, and use existing FSx for Lustre file systems when launching instances and running jobs.
- **Multiple deployment types:** FSx for Lustre offers a choice between scratch and persistent file systems for short-term and longer-term data processing. Scratch file systems are ideal for temporary storage and shorter-term processing of data. Data is not replicated and does not persist if a file server fails. Persistent file systems are ideal for longer-term storage and workloads. With persistent file systems, data is replicated, and file servers are replaced if they fail.
- **Multiple storage options:** FSx for Lustre offers Solid-State Disk (SSD) and Hard Disk Drive (HDD) storage options to optimize cost and performance for your workload. For low-latency, IOPS-intensive workloads that typically feature small, random file operations, you change choose one of the SSD storage options. For throughput-intensive workloads that typically feature large, sequential file operations, you can choose one of the HDD storage options.
- **Manage consumption with storage quotas:** You can use storage quotas to monitor and control user-and group-level storage consumption on your file systems, and to ensure that no user or group is able to consume excessive amounts of capacity. Storage quotas are intended for storage administrators who manage multi-user file systems such as user shares for data scientists, computational engineers, and genomics researchers.
- **Meet security and compliance requirements:** All Amazon FSx for Lustre file systems are encrypted at-rest, and in-transit encryption is available in select regions.
- **Network isolation:** You access your Amazon FSx file system from endpoints in your Amazon VPC, which enables you to isolate your file system in your own virtual network. You can configure security group rules and control network access to your Amazon FSx file systems.

- **Resource-level permissions:** Amazon FSx is integrated with AWS Identity and Access Management (IAM). This integration means that you can control the actions your AWS IAM users and groups can take to manage your file systems (such as creating and deleting file systems). You can also tag your Amazon FSx resources and control the actions that your IAM users and groups can take based on those tags.
- **Centralized backup and compliance management with AWS Backup:** Amazon FSx is integrated with AWS Backup enabling fully managed, policy-based backup and restore capabilities for your Amazon FSx file systems. The integration with AWS Backup allows you to protect customer data and ensure compliance across AWS services for business continuity purposes.

39.1.2. Benefits

- **Improve workload performance:** Amazon FSx for Lustre file systems scale to terabytes per second of throughput and millions of IOPS. FSx for Lustre also supports concurrent access to the same file or directory from thousands of compute instances. FSx for Lustre provides consistent, low latencies for file operations.
- **Use for any compute workload:** FSx for Lustre is compatible with the most popular Linux-based AMIs, including Amazon Linux, Red Hat Enterprise Linux (RHEL), CentOS, Ubuntu, and SUSE Linux.
- **Easily import and export Amazon S3 data:** Amazon FSx for Lustre integrates natively with Amazon S3, making it easy to access your S3 data to run data-processing workloads.
- **Simple to use with compute services:** Amazon FSx for Lustre is accessible from workloads running on Amazon EC2 instances or on on-premises computers/servers. Once mounted, you can work with the files and directories in your file system just like you would with a local file system. FSx for Lustre file systems also are accessible from containers running on Amazon Elastic Kubernetes Service (EKS).
- **Accelerate Amazon SageMaker training jobs:** Amazon FSx for Lustre integrates with Amazon Sagemaker as an input data source. When using Amazon SageMaker with Amazon FSx for Lustre, your machine learning training jobs are accelerated by eliminating the initial download step from S3, and your TCO is reduced by avoiding the repeated download of common objects (saving S3 request costs) for iterative jobs on the same data set.
- **Reduce the effective cost of storage using data compression:** You can use data compression to reduce storage consumption of both your file system storage and your file system backups. The data compression feature uses the LZ4 compression algorithm which is optimized to deliver high levels of compression without adversely impacting file system performance. Once data compression is enabled, newly written files are automatically compressed by FSx for Lustre before they are written to disk and automatically uncompressed when they are read.
- **Eliminate administrative burden and scale capacity on demand:** With a few clicks in the Amazon FSx console, CLI, or API you can create and scale a high-performance Lustre file system. With Amazon FSx file systems, you don't have to worry about managing file servers and storage volumes, updating hardware, configuring software, running out of capacity, or tuning performance -- Amazon FSx automates these time-consuming administration tasks.

39.2. Backup/Restore and Disaster Recovery

Amazon FSx for Lustre is designed for high-performance workloads where your long-term data is stored in a data repository such as Amazon S3 or an on-premises data store. As a result you inherit all of the availability and durability benefits that Amazon S3 provides. Amazon S3 Standard, S3 Standard–IA, S3 One Zone-IA, and S3 Glacier are all designed to provide 99.999999999% durability of objects over a given year.

39.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

39.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/fsx/
- **Service quotas:** https://docs.aws.amazon.com/fsx/latest/LustreGuide/limits.html
- **Service FAQs:** https://aws.amazon.com/fsx/lustre/faqs/

39.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/fsx/ and the following links for comprehensive technical documentation regarding this service.

- Amazon FSx for Lustre User Guide: Describes key concepts of Amazon FSx for Lustre and provides instructions for launching an FSx for Lustre file system and using the features of FSx for Lustre.

40. Amazon FSx for OpenZFS

40.1. Service Overview

Amazon FSx for OpenZFS is a storage service that lets you launch, run, and scale fully managed OpenZFS file systems on AWS. It provides the familiar features, performance, and capabilities of OpenZFS file systems with the agility, scalability, and simplicity of a fully managed AWS service.

FSx for OpenZFS file systems are accessible from Linux, Windows, and macOS compute instances and containers via the industry-standard NFS protocol (v3, v4, v4.1, and v4.2). Powered by the AWS Graviton family of processors along with the latest AWS disk and networking technologies, FSx for OpenZFS delivers up to 1 million IOPS, with latencies of a few hundred microseconds for your high-performance workloads.

40.1.1. Features

- **High-speed, low-latency file storage in the cloud:** Amazon FSx for OpenZFS delivers one of the lowest file storage latencies available in the cloud. It is built on the latest AWS compute, disk, and networking technologies, providing fast, consistent latencies of a few hundred microseconds for your high-performance workloads with even lower latencies for frequently accessed data.
- **Throughput and IOPS performance:** Amazon FSx for OpenZFS is built on the AWS Graviton family of processors, supporting up to 12.5 gigabytes per second (GB/s) of throughput and up to 1 million IOPS for frequently accessed cached data. For data

accessed from persistent disk storage, FSx for OpenZFS file systems deliver up to 4 GB/s and up to 160,000 IOPS. You can also enable data compression on your file system to help increase your effective throughput.

- **Scalable performance for up to thousands of clients:** FSx for OpenZFS supports simultaneous access from up to thousands of clients so that you can deliver shared, high-performance file storage for users or applications at scale. With support for multiple parallel connections per client, you can deliver your file system's maximum levels of throughput and IOPS even for just a single client.
- **Flexible storage and performance capacity:** Amazon FSx for OpenZFS lets you independently set your storage and performance capacity to optimize your file system for your specific workload needs. You can scale your provisioned throughput at any time to adapt your file system as your needs evolve.
- **ZFS-powered storage efficiency capabilities:** Amazon FSx for OpenZFS is powered by the popular OpenZFS file system, which was designed to provide high levels of storage efficiency. Amazon FSx for OpenZFS supports the latest Z-Standard and LZ4 compression technologies, which can help you reduce your storage costs for high-performance FSx for OpenZFS file systems and backup storage. It also supports multiple data containers (volumes) per file system, thin provisioning, and user or group storage quotas to enable you to efficiently support multiple teams, applications, and use cases within a single file system.
- **Pay-as-you-go billing model:** With Amazon FSx, you pay only for the resources you use. You are billed for file systems based on the storage capacity (per GB-month), SSD IOPS (per IOPS-month), and throughput capacity (per MBps-month) that you provision. You are billed by the second, ensuring that you only pay for resources for the period of time when you're using them.
- **Support for the latest versions of NFS:** Amazon FSx for OpenZFS provides full support for NFS v3, v4.0, v4.1, and v4.2 to help you migrate your NFS-based or other ZFS-based on-premises applications and storage to AWS, and build new cloud-native applications easily using a simple, highly available, high-performance NFS endpoint.
- **Accessible in AWS and on premises:** With Amazon FSx for OpenZFS, you can access file systems from another VPC (including a VPC in another region) using AWS Transit Gateway or VPC Peering, and you can access file systems from on premises using AWS Direct Connect or VPN.
- **Rich ZFS capabilities for working with data:** Amazon FSx for OpenZFS provides rich ZFS capabilities for working with data, like point-in-time snapshots and in-place data cloning, natively via the FSx API. With snapshots, you can easily track historical versions of your data and applications, and restore these versions with a click of a button. With data cloning, you can clone data in place to quickly test features and changes without interrupting your existing users or applications, and without needing to duplicate your data. Clones are created almost instantly, they consume no additional capacity upon creation, and any data modifications are isolated from your original dataset.

40.1.2. Benefits

- **Simple, flexible administration:** With Amazon FSx for OpenZFS, you have full flexibility and control over how you administer your file systems. You can manage your file systems using the AWS Management Console, AWS Command Line Interface (AWS CLI), and AWS SDK.

- **Automatic file system backups for disaster recovery:** Amazon FSx for OpenZFS automatically takes highly durable, daily file-system backups to Amazon S3. You can modify your backup schedule, take additional backups of your file system, or restore your existing backups to a new file system at any time. Backups are point-in-time consistent, incremental, easy to manage, and quick to create and restore. To provide additional layers of data protection and meet business continuity, disaster recovery, and compliance requirements, you can also copy these file system backups across AWS Regions.
- **Easy file-level restore with snapshots:** Amazon FSx for OpenZFS supports near instant point-in-time volume snapshots that are stored directly within your file system. End users can easily restore volumes to past snapshots, or even undo changes and compare versions of individual files or directories.
- **Encryption:** All Amazon FSx for OpenZFS file system data is automatically encrypted at rest using keys managed with AWS Key Management Service (AWS KMS). Data is automatically encrypted in transit with 256-bit encryption when accessed from supported Amazon EC2 instance types.
- **Network isolation:** Amazon FSx for OpenZFS is integrated with IAM, which lets you control who can administer your file systems, volumes, and backups (create, update, delete, etc.). You can also tag your Amazon FSx for OpenZFS resources and control the actions that your IAM users and groups can take based on those tags.
- **Resource-level permissions and tagging:** Amazon FSx for OpenZFS is integrated with IAM, which lets you control who can administer your file systems, volumes, and backups (create, update, delete, etc.). You can also tag your Amazon FSx for OpenZFS resources and control the actions that your IAM users and groups can take based on those tags.
- **API activity monitoring:** You can monitor and secure API calls using AWS CloudTrail and IAM and detect and flag suspicious API usage patterns using Amazon GuardDuty.
- **File- and directory-level access control:** Amazon FSx for OpenZFS supports POSIX permissions and POSIX ACLs so that you can define fine-grained file- and directory-level access permissions for individual users and groups.
- **Compliance:** AWS has the longest-running compliance program in the cloud and is committed to helping you navigate your requirements. Amazon FSx for OpenZFS has been assessed to meet global and industry security standards, complying with PCI DSS, ISO 9001, 27001, 27017, and 27018; SOC 1, 2, and 3; in addition to being HIPAA eligible. For more information and resources, visit our compliance page. You can also refer to the Services in Scope by Compliance Program page to see a full list of services and certifications.

40.2. Backup/Restore and Disaster Recovery

This service allows for backup and recovery. The service can take automatic daily backups of the file system. Users can automate this when setting up the fsx file system for the first time or it can be done incrementally with the user manually interacting with it.

40.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

40.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/fsx/index.html
- **Service quotas:** https://docs.aws.amazon.com/fsx/latest/OpenZFSGuide/limits.html
- **Service FAQs:** https://aws.amazon.com/fsx/openzfs/faqs/

40.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/fsx/index.html and the following links for comprehensive technical documentation regarding this service.

- Amazon FSx for Lustre User Guide: Describes key concepts of Amazon FSx for Lustre and provides instructions for launching an FSx for Lustre file system and using the features of FSx for Lustre.
- Amazon FSx for NetApp ONTAP User Guide: Describes key concepts of Amazon FSx for NetApp ONTAP and provides instructions for launching an FSx for ONTAP file system and using the features of FSx for ONTAP.
- Amazon FSx for OpenZFS User Guide: Describes key concepts of Amazon FSx for OpenZFS and provides instructions for launching an FSx for OpenZFS file system and using the features of FSx for OpenZFS.
- Amazon FSx for Windows File Server User Guide: Describes key concepts of Amazon FSx for Windows File Server and provides instructions for launching an FSx for Windows File Server file system and using the features of FSx for Windows File Server.
- Amazon FSx API Reference: Describes the Amazon FSx API operations and data types for all Amazon FSx file system types. Also provides sample requests, responses, and errors for the supported web services protocols.
- Amazon FSx section of the AWS CLI Reference: Documents the AWS CLI commands for Amazon FSx.

41. Amazon FSx for Windows File Server

41.1. Service Overview

Amazon FSx for Windows File Server makes it easy for you to launch and scale reliable, performant, and secure shared file storage for your applications and end users. With Amazon FSx, you can launch highly durable and available file systems that can span multiple availability zones (AZs) and can be accessed from up to thousands of compute instances using the industry-standard Server Message Block (SMB) protocol. It provides a rich set of administrative and security features, and integrates with Microsoft Active Directory (AD). To serve a wide spectrum of workloads, Amazon FSx provides high levels of file system throughput and IOPS and consistent sub-millisecond latencies.

41.1.1. Features

- **Performance:** Amazon FSx is designed to deliver fast, predictable, and consistent performance. Amazon FSx provides multiple GB/s of throughput per file system, hundreds of thousands of IOPS per file system, and consistent sub-millisecond latencies for file operations. To get the right performance for your workload, you can choose a

throughput level for your file system and scale this throughput level up or down at any time.

- **Scale:** Amazon FSx provides storage of up to 64 TB per file system. You can use DFS Namespaces to create shared common namespaces spanning multiple Amazon FSx file systems to scale out storage and throughput to virtually unlimited levels.
- **Encryption:** All Amazon FSx file system data is automatically encrypted at rest and in transit. Encryption of data at-rest uses keys managed with AWS Key Management Service (AWS KMS). Data is automatically encrypted before being written to the file system, and automatically decrypted as it is read. You can also choose to enforce encryption of data in-transit on all connections to your file systems for compliance needs. Amazon FSx automatically encrypts data-in-transit using SMB Kerberos session keys, when accessed from compute instances that support SMB protocol 3.0 or newer. This includes all Windows versions starting from Windows Server 2012 and Windows 8, and all Linux clients with Samba client version 4.2 or newer.
- **Identity-based authentication:** Amazon FSx supports identity-based authentication over SMB through Microsoft Active Directory (AD). When creating your Amazon FSx file system, you join it to your Microsoft AD -- either an AWS Managed Microsoft AD or your self-managed Microsoft AD. Your users can then use their existing AD-based user identities to authenticate themselves and access the Amazon FSx file system, and to control access to individual files and folders.
- **Access control and monitoring:** Amazon FSx supports Windows Access Control Lists (ACLs) for fine-grained file and folder access control. For network-level access control, you can use Amazon Virtual Private Cloud (Amazon VPC) security groups to control access to your Amazon FSx resources. Amazon FSx is integrated with AWS Identity and Access Management (IAM) to control the actions that your AWS IAM users and groups can take on specific Amazon FSx resources. Amazon FSx integrates with AWS CloudTrail to monitor and log administration actions. Amazon FSx also offers user storage quotas to monitor and control user-level storage consumption.
- **Network isolation:** You access your Amazon FSx file system from your Amazon VPCs. You can configure firewall settings and control network access to your Amazon FSx file systems using Amazon VPC Security Groups and VPC Network ACLs.
- **File access auditing:** Amazon FSx supports auditing end-user access to your files, folders, and file shares using Windows event logs. Logs are published to Amazon CloudWatch Logs or streamed to Amazon Kinesis Data Firehose, enabling you to view and query logs on CloudWatch Logs, archive logs in Amazon S3, trigger Lambda functions to take reactive actions, or perform post-processing on AWS Partner solutions such as Splunk and Datadog.
- **Highly available and durable:** To ensure high availability and durability, Amazon FSx automatically replicates your data within an Availability Zone (AZ) it resides in (which you specify during creation) to protect it from component failure, continuously monitors for hardware failures, and automatically replaces infrastructure components in the event of a failure. Amazon FSx offers single AZ and multi-AZ deployment options for your Windows file-based workloads.
- **Multi-AZ deployments:** Amazon FSx offers a multiple availability (AZ) deployment option, designed to provide continuous availability to data, even in the event that an AZ is unavailable. Multi-AZ file systems include an active and standby file server in separate

AZs, and any changes written to disk in your file system are synchronously replicated across AZs to the standby. During planned maintenance, or in the event of a failure of the active file server or its AZ, Amazon FSx automatically fails over to the standby so you can resume file system operations without a loss of availability to your data.

- **Support for High Availability Microsoft SQL Server deployments:** High Availability (HA) Microsoft SQL Server is typically deployed across multiple database nodes in a Windows Server Failover Cluster (WSFC), with each node having access to shared file storage. With support for Continuously Available (CA) file shares, Amazon FSx enables you to provide highly-available shared file storage for these clusters.
- **Automated daily backups:** To help ensure that your data is protected, Amazon FSx automatically takes highly durable, file-system consistent daily backups to S3. Amazon FSx uses the Volume Shadow Copy Service (VSS) to make your backups file system-consistent. You can take additional backups of your file system at any point.
- **Storage options:** Amazon FSx provides two types of storage – Hard Disk Drives (HDD) and Solid State Drives (SSD) – enabling you to optimize cost and performance to meet your workload needs. HDD storage is designed for a broad spectrum of workloads, including home directories, user and departmental shares, and content management systems. SSD storage is designed for the highest-performance and most latency-sensitive workloads, including databases, media processing workloads, and data analytics applications.
- **User quotas:** Amazon FSx offers user quotas to monitor and control user-level storage consumption on your file systems for use cases such as cost allocation across teams and limiting storage consumption on a user-level.

41.1.2. Benefits

- **Built on Windows Server:** Amazon FSx is built on Windows Server, providing a rich set of administrative features that include end-user file restore, user quotas, and Access Control Lists (ACLs). With Windows Server's native support for the SMB protocol, Windows-based applications have access to fully-compatible shared file storage. And since SMB file shares can also be accessed from Linux and MacOS, any application or user can access the storage regardless of operating system. To control user access, Amazon FSx integrates with your on-premises Microsoft Active Directory as well as with AWS Microsoft Managed AD.
- **Broadly accessible:** By supporting the SMB protocol, Amazon FSx can connect your file system to Amazon EC2, Amazon ECS, VMware Cloud on AWS, Amazon WorkSpaces, and Amazon AppStream 2.0 instances. Amazon FSx supports all Windows versions starting from Windows Server 2008 and Windows 7, and current versions of Linux and MacOS. Amazon FSx also supports on-premises access via AWS Direct Connect or AWS VPN, and access from multiple VPCs, accounts, and regions using VPC Peering or AWS Transit Gateway. Amazon FSx File Gateway provides efficient, low-latency on-premises access with a local cache for frequently accessed file data.
- **Fully managed:** Because Amazon FSx is a fully managed service, it makes it simple to launch and scale reliable, performant, and secure shared file storage in the cloud. In minutes, you can easily create Amazon FSx file systems that span multiple AZs by using the AWS Management Console, AWS CLI, or AWS SDK. Amazon FSx sets up and provisions file servers and storage volumes, replicates data, manages failover and

failback, and eliminates much of the need for administrative overhead. Amazon FSx also takes care of Windows Server software updates.

- **Simple and seamless migration with AWS DataSync:** You can easily move your self-managed file systems to fully managed Windows storage on Amazon FSx in minutes with AWS DataSync. Integration with AWS DataSync automates and accelerates copying data over the internet or AWS Direct Connect, and copies your files together with file attributes and metadata.
- **Easy file-level restores (Microsoft Windows shadow copies):** To enable end-users to easily undo changes and compare file versions, Amazon FSx supports restoring individual files and folders to previous versions using Windows shadow copies.
- **Centralized backup and compliance with AWS Backup:** To meet enterprise compliance and data protection requirements, Amazon FSx is integrated with AWS Backup allowing you to create scheduled, policy-driven backup plans for your Amazon FSx file systems.
- **Pricing:** You pay only for the resources you use, with no minimum commitments, licensing costs, or up-front fees. You are billed hourly for your Amazon FSx file systems, based on your configured storage capacity (priced per GB-month) and throughput capacity (priced per MBps-month). You are billed hourly for your backup storage (priced per GB-month). For more details, see the Amazon FSx pricing page.
- **Data deduplication:** You can enable data deduplication and compression to automatically reduce costs associated with redundant data by storing duplicated portions of your dataset only once. Typical savings average 50-60% for general purpose file shares, 30-50% savings for user documents, and 70-80% savings for software development data sets.

41.2. Backup/Restore and Disaster Recovery

To ensure high availability and durability, Amazon FSx automatically replicates your data within the Availability Zone (AZ) it resides in to protect it from component failure, continuously monitors for hardware failures, and automatically replaces infrastructure components in the event of a failure. Automatic, file-system consistent daily backups are done to S3 and you can take additional backups of your file system at any point. Multi-AZ deployments can be implemented using Microsoft DFS Replication and Namespace Features.

41.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

41.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/fsx/
- **Service quotas:** https://docs.aws.amazon.com/fsx/latest/WindowsGuide/limits.html
- **Service FAQs:** https://aws.amazon.com/fsx/windows/faqs/

41.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/fsx/ and the following links for comprehensive technical documentation regarding this service.

- Amazon FSx for Windows File Server User Guide: Describes key concepts of Amazon FSx for Windows File Server and provides instructions for launching an FSx for Windows File Server file system and using the features of FSx for Windows File Server.

42. Amazon GuardDuty

42.1. Service Overview

Amazon GuardDuty offers threat detection that enables you to continuously monitor and protect your AWS accounts and workloads. Amazon GuardDuty analyzes continuous streams of meta-data generated from your account and network activity found in AWS CloudTrail Events, Amazon VPC Flow Logs, and DNS Logs. It also uses integrated threat intelligence such as known malicious IP addresses, anomaly detection, and machine learning to identify threats more accurately. Amazon GuardDuty operates completely independently from your resources, so there is no risk of performance or availability impacts to your workloads.

42.1.1. Features

- **Accurate, account-level threat detection:** Amazon GuardDuty gives you accurate threat detection of compromised accounts, which can be difficult to detect quickly if you are not continuously monitoring factors in near real-time. GuardDuty can detect signs of account compromise, such as AWS resource access from an unusual geo-location at an atypical time of day.
- **Continuous monitoring across AWS accounts without added cost and complexity:** Amazon GuardDuty continuously monitors and analyzes your AWS account and workload event data found in AWS CloudTrail, VPC Flow Logs, and DNS Logs. There is no additional security software or infrastructure to deploy and maintain. By associating your AWS accounts together, you can aggregate threat detection instead of working on an account-by-account basis
- **Threat detections developed and optimized for the cloud:** Amazon GuardDuty gives you access to built-in detection techniques developed and optimized for the cloud. AWS Security continuously maintains and improves these detection algorithms.
- **Threat severity levels for efficient prioritization:** Amazon GuardDuty provides three severity levels (Low, Medium, and High) to help customers prioritize their response to potential threats. A “Low” severity level indicates suspicious or malicious activity that was blocked before it compromised your resource. A “Medium” severity level indicates suspicious activity. A “High” severity level indicates that the resource in question (e.g. an Amazon EC2 instance or a set of IAM user credentials) is compromised and is actively being used for unauthorized purposes.
- **Threat response and remediation automation:** Amazon GuardDuty offers HTTPS APIs, command-line interface (CLI) tools, and Amazon CloudWatch Events to support automated security responses to security findings.
- **Highly available threat detection:** Amazon GuardDuty is designed to automatically manage resource utilization based on the overall activity levels within your AWS accounts, workloads, and data stored in Amazon S3. GuardDuty adds detection capacity only when necessary, and reduces utilization when capacity is no longer needed.

- **One-click deployment with no additional software or infrastructure to deploy and manage**: With one click in the AWS Management Console or a single API call, you can enable Amazon GuardDuty on a single account. With a few more clicks in the console, you can enable GuardDuty across multiple accounts. Amazon GuardDuty supports multiple accounts through AWS Organizations integration as well as natively within GuardDuty.

42.1.2. Benefits

- **Visibility:** Achieve organization-wide visibility into possible threats with only a few clicks.
- **Threat detection:** Expose threats quickly with AWS threat intelligence, behavioral models, and third-party security feeds.
- **Threat mitigation:** Mitigate threats early by triggering automated responses.
- **Stop unauthorized activity:** Guard against use of compromised credentials, unusual data access in Amazon Simple Storage Service (S3), API calls from known malicious IP addresses, and more.
- **Enable continuous monitoring and analysis:** Gain insight into security events with findings that provide context, metadata, and details on impacted resources.
- **Simplify forensics:** Quickly determine the root cause of suspicious activities using Amazon GuardDuty's console integration with Amazon Detective.

42.2. Backup/Restore and Disaster Recovery

This requirement is not applicable for this service. For additional information beyond what is described herein, please refer to http://aws.amazon.com/documentation/.

42.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

42.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/guardduty/
- **Service quotas:** https://docs.aws.amazon.com/guardduty/latest/ug/guardduty_limits.html
- **Service FAQs:** https://aws.amazon.com/guardduty/faqs/

42.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/guardduty/ https://docs.aws.amazon.com/guardduty/and the following links for comprehensive technical documentation regarding this service.

- User Guide: Walks through how to set up Amazon GuardDuty and evaluate the security of your AWS environment.
- API Reference: Describes all of the API operations for Amazon GuardDuty.

43. Amazon Healthlake

43.1. Service Overview

Amazon HealthLake is a HIPAA-eligible service enabling healthcare and life sciences companies to securely store and transform their data into a consistent and queryable fashion. Using the HealthLake APIs, healthcare organizations can easily copy health data, such as medical reports or patient notes, from on-premises systems to a secure data lake in the cloud, and analyse it at petabyte scale. HealthLake uses machine learning (ML) models to automatically understand and extract meaningful medical information from the raw data, such as medications, procedures, and diagnoses. HealthLake organizes and indexes all the information and stores it in the Fast Healthcare Interoperability Resources (FHIR) industry standard format to provide a complete view of each patient's medical history. Organizations can build ML models with Amazon SageMaker and use advanced Amazon QuickSight analytics to understand relationships, identify trends, and make predictions from the newly normalized and structured data.

43.1.1. Features

- **Quickly & easily ingest health data:** With the Amazon HealthLake import API you can easily migrate FHIR files from Amazon S3 to the Amazon HealthLake Data Store including clinical notes, lab reports, insurance claims, and more. HealthLake supports data in the FHIR R4 industry standard. If your data is not in this format, you can work with an AWS partner to convert your health data FHIR.
- **Store health data in a secure, compliant, & auditable manner:** Data Store helps index all information so it can be easily queried. The Data Store creates a complete view of each patient’s medical history in chronological order and facilitates information exchange using the V4 FHIR specification. The Data Store is always running to keep your index up-to-date, offering you the ability to query the information anytime using the standard FHIR Operations with durable primary storage and index scaling.
- **Transform unstructured medical data using NLP:** Integrated medical natural language processing (NLP) transforms all raw medical text data in the Data Store to understand and extract meaningful information from unstructured healthcare data. With integrated medical NLP, you can automatically extract entities (e.g., medical procedures, medications), entity relationships (e.g., a medication and its dosage), entity traits (e.g., positive or negative test result, time of procedure), and Protected Health Information (PHI) data from your medical text.
- **Powerful query & search capabilities:** Amazon HealthLake supports FHIR Create/Read/Update/Delete (CRUD) and FHIR Search operations. You can query records by performing a Create Operation for adding new patients and their information, like medications.
- **Identify trends & make predictions:** Amazon HealthLake supports the bulk export of FHIR data from the HealthLake Data Store to an S3 bucket. With Amazon QuickSight, you can create dashboards on the exported and normalized data to quickly explore patient trends. You can also build, train, and deploy your own predictive analytics using machine learning models with Amazon SageMaker.

43.1.2. Benefits

- **Analyze unstructured data:** Extract meaning from unstructured data with natural language processing (NLP) for easy search and querying.

- **Predictive health:** Make predictions with health data using Amazon SageMaker machine learning (ML) models and Amazon QuickSight analytics.
- **Interoperability:** Support interoperable standards such as the Fast Healthcare Interoperability Resources (FHIR) format.
- **Holistic view of patient:** Create a complete and chronological view of patient health data, including prescriptions, procedures, and diagnoses.
- **Manage population health:** Analyze population health trends, predict outcomes, and manage costs with advanced analytics tools and ML models.
- **Improve care quality:** Identify opportunities to close gaps in care and apply preventative treatment with a complete timeline view of patient medical histories.
- **Optimize hospital efficiency:** Apply advanced analytics and ML to newly structured data to optimize appointment scheduling, reduce unnecessary procedures, and predict hospital bed availability.

43.2. Backup/Restore and Disaster Recovery

This requirement is not applicable for this service. For additional information beyond what is described herein, please refer to http://aws.amazon.com/documentation/.

43.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

43.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/healthlake/index.html
- **Service quotas:** https://docs.aws.amazon.com/healthlake/latest/devguide/quotas.html
- **Service FAQs:** https://aws.amazon.com/healthlake/faqs/

43.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/healthlake/index.html and the following links for comprehensive technical documentation regarding this service.

- Amazon HealthLake Developer Guide: Provides a conceptual overview of Amazon HealthLake and detailed instructions for using the various features.
- API Reference: Provides information about API operations and data types. Includes request and response JSON and parameter descriptions, along with valid values and other useful information for developers.

44. Amazon Honeycode

44.1. Service Overview

Amazon Honeycode, which is available in beta, is a fully managed service that allows customers to build powerful mobile and web applications quickly—with no programming required. With Amazon Honeycode, customers can use a simple visual application builder to create highly interactive web and mobile applications to track and manage things like process approvals, event scheduling, customer relationship management, user surveys, to-do lists, and

content and inventory tracking. Customers can build applications that range in complexity from a task-tracking application for a small team to a project management system that manages a complex workflow for multiple teams or departments.

44.1.1. Features

Build custom apps without programming: Amazon Honeycode gives you the power to build apps that improve how your team works. Build one or many apps - the only limit is your creativity.

Integrate existing workflows: You can easily integrate Honeycode with popular SaaS applications, AWS services, and other tools by using Zapier or Amazon AppFlow.

Automate manual steps: Set up your custom Honeycode app to automatically notify the team when there's an update or remind people when it's their turn to take action.

44.1.2. Benefits

Stay in the loop: As your business changes, you can adapt your custom app built with Honeycode. Any updates made to your app or its data are instantly shared to your team.

Personalize for teams: You can configure your custom Honeycode app so that each team member sees only the data they need to see and nothing more.

Make it mobile: You can build your custom Honeycode app for web browsers and mobile devices so your team can work from anywhere.

44.2. Backup/Restore and Disaster Recovery

This requirement is not applicable for this service. For additional information beyond what is described herein, please refer to http://aws.amazon.com/documentation/.

44.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

44.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/honeycode/?id=docs_gateway
- **Service quotas:** https://docs.aws.amazon.com/general/latest/gr/honeycode.html
- **Service FAQs:** https://docs.aws.amazon.com/honeycode/latest/UserGuide/faqs.html

44.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/honeycode/?id=docs_gateway and the following links for comprehensive technical documentation regarding this service.

- User Guide: Provides conceptual overview and describes key concepts of Amazon Honeycode with detailed instructions for using its various features.
- Administrator's Guide: Provides information for IT Professionals for integrating Honeycode with Single Sign-On systems.
- API Reference: Describes the Honeycode APIs in detail. These API operations allow programmatic interaction with the Honeycode app.

- Honeycode Help and Community: Help and Community

45. Amazon Inspector

45.1. Service Overview

Amazon Inspector is a vulnerability management service that continually scans AWS workloads for software vulnerabilities and unintended network exposure. With a few clicks in the AWS Management Console, Amazon Inspector can be used across all accounts in your organization. Once started, Amazon Inspector automatically discovers running Amazon Elastic Compute Cloud (EC2) instances and container images residing in Amazon Elastic Container Registry (ECR), at any scale, and immediately starts assessing them for known vulnerabilities.

Amazon Inspector calculates a highly contextualized risk score for each finding by correlating common vulnerabilities and exposures (CVE) information with factors such as network access and exploitability. This score is used to prioritize the most critical vulnerabilities to improve remediation response efficiency. All findings are aggregated in a newly designed Amazon Inspector console and pushed to AWS Security Hub and Amazon EventBridge to automate workflows. Vulnerabilities found in container images are also sent to Amazon ECR for resource owners to view and remediate. With Amazon Inspector, even small security teams and developers can ensure infrastructure workload security and compliance across your AWS workloads.

45.1.1. Features

- **Vulnerability management for Amazon EC2 and container workloads:** Amazon Inspector is a comprehensive vulnerability management tool that functions across multiple resources, including Amazon Elastic Compute Cloud (EC2) and container workloads. It identifies different types of vulnerabilities, including software vulnerabilities and unintended network exposure, that can be used to compromise workloads, repurpose resources for malicious use, or exfiltrate data.
- **Simplified one-click enabling and integration with AWS Organizations:** Start Amazon Inspector across multiple accounts with one click in the Amazon Inspector console or a single API call. Amazon Inspector allows you to assign an Inspector Delegated Administrator (DA) account for your organization, which can start and configure all member accounts as well as consolidate all findings.
- **Automated discovery and continual vulnerability scanning:** Once started, Amazon Inspector automatically discovers all EC2 instances and container images residing in Amazon Elastic Container Registry (ECR) that are identified for scanning, and then immediately starts scanning them for software vulnerabilities and unintended network exposure. All workloads are continually rescanned when a new common vulnerabilities and exposures (CVE) is published, or when there are changes in the workloads, such as installation of new software in an EC2 instance.
- **AWS Systems Manager Agent:** Amazon Inspector uses the widely deployed AWS Systems Manager (SSM) Agent to collect the software inventory and configurations from your Amazon EC2 instances. The collected application inventory and configurations are used to assess workloads for vulnerabilities.
- **Inspector risk score for findings:** Amazon Inspector generates a highly contextualized Inspector risk score for each finding by correlating CVE information with environmental factors such as network reachability results and exploitability data. This helps prioritize the findings and highlights the most critical findings and vulnerable resources. The

Inspector score calculation (and which factors influenced the score) can be viewed in the Inspector Score tab within the Findings Details side panel.

- **Suppression of findings:** Amazon Inspector supports suppression of findings based on criteria you define. You can create these suppression rules to suppress findings that your organization deems an acceptable risk.
- **Automatic closure of remediated findings:** Amazon Inspector automatically detects if a vulnerability has been patched or remediated. Once detected, Amazon Inspector automatically changes the state of the finding to “Closed” without manual intervention.
- **Detailed coverage monitoring:** Amazon Inspector offers an aggregated, near real-time view of the environment coverage across an organization so you can avoid gaps in coverage. It provides metrics and detailed information on accounts using Amazon Inspector, as well as EC2 instances, ECR repositories, and container images that are actively being scanned by Amazon Inspector. Additionally, Amazon Inspector highlights the resources not being actively monitored and provides guidance on how to include them.
- **Integration with Security Hub and EventBridge:** All findings are aggregated in the Amazon Inspector console, routed to AWS Security Hub, and pushed through Amazon EventBridge to automate workflows such as ticketing.

45.1.2. Benefits

- **Quickly discover vulnerabilities:** Automatically discover and quickly route vulnerability findings in near real time to the appropriate teams so they can take immediate action.
- **Prioritize patch remediation:** Use up-to-date common vulnerabilities and exposures (CVE) information combined with factors such as network accessibility to create context-based risk scores that help you prioritize and address vulnerable resources.
- **Meet compliance requirements:** Support compliance requirements and best practices for NIST CSF, PCI DSS, and other regulations with Amazon Inspector scans.
- **Identify zero-day vulnerabilities sooner:** Accelerate MTTR by using over 50 sources for vulnerability intelligence to help identify zero-day vulnerabilities quickly.

45.2. Backup/Restore and Disaster Recovery

This requirement is not applicable for this service. For additional information beyond what is described herein, please refer to http://aws.amazon.com/documentation/.

45.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

45.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/inspector/
- **Service quotas:** https://docs.aws.amazon.com/general/latest/gr/inspector.html
- **Service FAQs:** https://aws.amazon.com/inspector/faqs/

45.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/inspector/ and the following links for comprehensive technical documentation regarding this service.

- User Guide: Walks through how to set up Amazon Inspector and evaluate your security configuration.
- API Reference: Describes all the API operations for Amazon Inspector in detail. Also provides sample requests, responses, and errors for the supported web services protocols.

Amazon Inspector section of the AWS CLI Reference: Describes the AWS CLI commands that you can use to administer Amazon Inspector. Provides syntax, options, and usage examples for each command.

46. Amazon IVS

46.1. Service Overview

Amazon Interactive Video Service (Amazon IVS) is a managed live streaming solution that is quick and easy to set up, and ideal for creating interactive video experiences. Send your live streams to Amazon IVS using streaming software and the service does everything you need to make low-latency live video available to any viewer around the world, letting you focus on building interactive experiences alongside the live video. You can easily customize and enhance the audience experience through the Amazon IVS player SDK and timed metadata APIs, allowing you to build a more valuable relationship with your viewers on your own websites and applications.

46.1.1. Features

- **Set Up New Live Video Streams in Minutes:** Amazon Interactive Video Service provides all the components needed for a low-latency live video streaming solution. You send a live video stream from an encoder or encoding software to the ingest point with the provided stream key. Then, use the playback URL with the player SDK and watch the live streams on websites, and iOS and Android applications.
- **Optimized for Live Video Streaming:** Get video to Amazon IVS, and distribute live streams across a global network optimized for low latency. To improve quality of service when sending video, Amazon IVS detects the optimal network paths from the streamer's location and intelligently selects the best endpoint to receive the input video stream. Live streams are delivered to audiences in seconds via a content pipeline designed end-to-end for live video.
- **Timed Metadata API with Low-Latency Live Video:** Use a simple REST API to inject metadata into a stream and an event-based interface within the Amazon IVS player SDK to retrieve the metadata for clients to build graphics, polls, and other synchronized components such as live sports scores and e-commerce functionality. Live video streams can achieve latency of less than three seconds from ingest to playback, making applications like polling your audience and voting engaging and interactive.
- **Multi-Platform Broadcast SDK:** Use iOS and Android applications to send live video to Amazon IVS with the mobile broadcast SDK. Applications using the SDK can use the device camera, microphone, screen sources, or use custom audio and video sources. Streams sent from devices will use optimized encoder settings and network congestion handling.

- **Multi-Platform Player SDK:** Provide a low-latency experience for web, iOS, and Android with the Amazon IVS player SDK (the web SDK integrates with Video.js). The player SDK is designed for Amazon IVS live video streams and includes support for chunked streaming and an adaptive bitrate switching algorithm. The switching algorithm allows for optimal performance and low latency without trade-offs in quality of service or video quality. You can also restrict video playlists for your streams by channel and viewer using playback authorization.
- **Record Live Streams for Use as Video-On-Demand:** You can configure Amazon IVS to record live video to an Amazon Simple Storage Service (Amazon S3) bucket. Video streams are saved as video files, and can be used to create video-on-demand content with AWS Elemental MediaConvert, or streamed directly as VOD.
- **Playback on Different Devices and Networks:** With a standard channel, Amazon IVS converts the incoming live video stream to a range of video resolutions and bitrates. This adaptive bitrate (ABR) stream provides an improved quality of experience and quality of service across different devices and network conditions.
- **Frustration-Free Scale:** Cover one, or hundreds of events at a moment's notice. Built on the same live streaming technology that powers Twitch, Amazon IVS channels take just seconds to start streaming live video, and the service scales to deliver streams to millions of concurrent viewers.
- **Cost-Effective:** Amazon IVS provides simple, pay-as-you-go pricing based on hours of video sent to the service (input), and hours of video delivered to your audiences (output).

46.1.2. Benefits

- **Simple-to-use live video streaming:** Create and configure live streams with just a few clicks, and start streaming in seconds. Send video to Amazon IVS, and the service manages everything needed to get your live stream to your viewers.
- **Build an engaging audience experience:** Amazon IVS streams are designed to provide low-latency live video "out of the box." Latency that can be less than three seconds means you can build engaging and interactive experiences alongside live video.
- **Optimized for live streaming:** Built on the same live streaming technology that powers Twitch, Amazon IVS provides an enhanced quality of service and experience by receiving and delivering streams over managed global infrastructure that is optimized for live video ingest, processing, and delivery.
- **Easy integration into websites and apps:** Get live streams into your iOS, Android, and web apps, quickly and easily with the Amazon IVS player and broadcast SDKs. The SDKs give your streamers and audiences a consistent experience and low-latency live streams across different platforms, without compromising video quality or increasing buffering.

46.2. Backup/Restore and Disaster Recovery

This requirement is not applicable for this service. For additional information beyond what is described herein, please refer to http://aws.amazon.com/documentation/.

46.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

46.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/ivs/index.html
- **Service quotas:** https://docs.aws.amazon.com/ivs/latest/userguide/service-quotas.html
- **Service FAQs:** https://aws.amazon.com/ivs/faqs/

46.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/ivs/index.html and the following links for comprehensive technical documentation regarding this service.

- User Guide: Introduces you to and helps you get started with Amazon IVS. Provides instructions on using various features with the console, API, and command-line interface. Includes information about using the Amazon IVS Player on various platforms.
- API Reference: Describes in detail all Amazon IVS API operations. Includes sample request and response formats and error information.
- Code Samples and Demos: Explore Amazon IVS resources. Discover demos, code samples, and blog posts.
- Broadcast: SDK for Android Reference: Describes in detail all operations for the Amazon IVS Broadcast SDK for Android. Includes sample request and response formats and error information.
- Broadcast: SDK for iOS Reference: Describes in detail all operations for the Amazon IVS Broadcast SDK for iOS. Includes sample request and response formats and error information.
- Player: SDK for Android Reference: Describes in detail all operations for the Amazon IVS Player SDK for Android. Includes sample request and response formats and error information.
- Player: SDK for iOS Reference: Describes in detail all operations for the Amazon IVS Player SDK for iOS. Includes sample request and response formats and error information.
- Player: SDK for Web Reference: Describes in detail all operations for the Amazon IVS Player SDK for Web. Includes sample request and response formats and error information.

47. Amazon Kendra

47.1. Service Overview

Amazon Kendra is an intelligent search service powered by machine learning (ML). Kendra reimagines enterprise search for your websites and applications so your employees and customers can easily find the content they're looking for, even when it's scattered across multiple locations and content repositories within your organization.

Using Amazon Kendra, you can stop searching through troves of unstructured data and discover the right answers to your questions when you need them. Amazon Kendra is a fully managed service, so there are no servers to provision and no ML models to build, train, or deploy.

47.1.1. Features

- **Intelligent search:** Amazon Kendra uses ML to deliver more-relevant answers from unstructured data. Amazon Kendra also supports FAQ matching and extracts answers from curated FAQs using a specialized model that pinpoints the closest question and returns the corresponding answer. To complement extracted answers and FAQ matching, Amazon Kendra uses a deep learning semantic search model for accurate document ranking. Overall, this provides a richer search experience that presents specific answers, as well as related content to explore if you need more information.
- **Incremental learning:** Amazon Kendra uses ML to continuously optimize search results based on end-user search patterns and feedback. To determine the most relevant document for this question, Amazon Kendra will learn from the user interactions and feedback to promote preferred documents to the top of the list. Amazon Kendra applies incremental learning techniques automatically without the need for ML expertise.
- **Tuning and accuracy:** You can fine-tune search results and boost specific answers and documents in the results based on specific business objectives. To extend Amazon Kendra's understanding of your specific business vocabulary, you can provide your own custom synonyms. Amazon Kendra uses these to automatically expand queries to include content and answers that match the extended vocabulary.
- **Connectors:** Using connectors is quick and easy—just add data sources to your Amazon Kendra index and select the connector type. Connectors can be scheduled to automatically sync your index with your data source, so you're always securely searching through the most up-to-date content.
- **Domain optimization:** Amazon Kendra uses deep learning models to understand natural language queries and document content and structures for a wide range of internal use cases, including HR, operations, support, and R&D. Amazon Kendra is also optimized to understand complex language from domains such as IT, financial services, insurance, pharmaceuticals, industrial manufacturing, oil and gas, legal, media and entertainment, travel and hospitality, health, news, telecommunications, mining, food and beverage, and automotive.
- **Experience Builder:** You can now deploy a fully functional and customizable search experience with Amazon Kendra in a few clicks, without any coding or ML experience. Experience Builder delivers an intuitive visual workflow to quickly build, customize, and launch your Amazon Kendra-powered search application securely on the cloud.
- **Search Analytics Dashboard:** Amazon Kendra Search Analytics Dashboard provides a snapshot of how your users interact with your search application and how effective your search results are. The analytics data can be viewed in a visual dashboard in the console, or you can build your own dashboards by accessing the Search Analytics data through an API. It empowers customers to dive deep into search trends and user behavior to identify insights, and also helps to bring clarity to potential areas of improvement.
- **Custom Document Enrichment:** With Amazon Kendra Custom Document Enrichment capabilities, you can build a custom ingestion pipeline that can pre-process documents before they get indexed into Amazon Kendra. The enrichment is performed by simple rules that can be configured in the console or by invoking functions from AWS Lambda. These functions can optionally call other AWS AI Services such as Amazon Comprehend, Amazon Transcribe, or Amazon Textract.

- **Query autocompletion:** Amazon Kendra includes the functionality to autocomplete an end user's search query. Query autocompletion not only helps the user reduce typing by about 25 percent, but it also helps by guiding them toward more-precise and commonly asked questions.

47.1.2. Benefits

- **Find relevant answers, quickly:** Say goodbye to sifting through long lists of links and scanning documents hoping that one has the information you need. Unlike conventional search technology, natural language search capabilities return the answers you're looking for quickly and accurately, no matter where the information lives within your organization.
- **Centralize access to knowledge:** Using Amazon Kendra, you can easily aggregate content from content repositories such as Microsoft SharePoint, Amazon S3, ServiceNow, Salesforce, and Amazon RDS into a centralized index that lets you quickly search all of your enterprise data and find the most accurate answer.
- **Fine-tune search results:** Amazon Kendra's deep learning models come pretrained across 14 industry domains, allowing it to extract more accurate answers across a wide range of business use cases. You can also fine-tune search results by manually adjusting the importance of data sources, authors, or freshness, or by using custom tags.
- **Deploy with just a few clicks:** Setup is quick, giving you faster access to Amazon Kendra's intelligent search capabilities compared to setup times for conventional search solutions. With just a few clicks, you can easily configure an index, connect relevant data sources, and deploy a fully functional and customizable search interface without any coding or ML experience.
- **Accelerate research and development:** Scientists and developers leading new research and innovation need access to information from prior work that's buried within their corporate data stores. With faster, more accurate search, they spend less time searching and more time innovating.
- **Minimize regulatory and compliance risks:** Use ML to quickly identify and interpret regulatory policies published across hundreds of diverse websites so you can improve policy enforcement and compliance processes.
- **Improve customer interactions:** Whether it's through Q&A chatbots, agent-assist, or customer web search, Amazon Kendra better understands what your customers are asking and provides more-relevant answers and intuitive experiences.
- **Increase employee productivity:** By unifying and indexing content from diverse, disparate, and multi-structure information silos across the organization, enterprises can build and maintain a single active knowledge catalog for all employees. With this centralized view, users can quickly search and access the most relevant information across any knowledge source to make more informed decisions.

47.2. Backup/Restore and Disaster Recovery

This requirement is not applicable for this service. For additional information beyond what is described herein, please refer to http://aws.amazon.com/documentation/.

47.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace

47.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/kendra/index.html
- **Service quotas:** https://docs.aws.amazon.com/kendra/latest/dg/quotas.html
- **Service FAQs:** https://aws.amazon.com/kendra/faqs/

47.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/kendra/index.html and the following links for comprehensive technical documentation regarding this service.

- Developer Guide: Provides a conceptual overview of Amazon Kendra. Learn about Amazon Kendra features and get started indexing your documents using the console, AWS CLI, and API.

48. Amazon Keyspaces (for Apache Cassandra)

48.1. Service Overview

Amazon Keyspaces (for Apache Cassandra) is a scalable, highly available, and managed Apache Cassandra–compatible database service. With Amazon Keyspaces, you can run your Cassandra workloads on AWS using the same Cassandra application code and developer tools that you use today. You don't have to provision, patch, or manage servers, and you don't have to install, maintain, or operate software. Amazon Keyspaces is serverless, so you pay for only the resources you use and the service can automatically scale tables up and down in response to application traffic. You can build applications that serve thousands of requests per second with virtually unlimited throughput and storage. Data is encrypted by default and Amazon Keyspaces enables you to back up your table data continuously using point-in-time recovery. Amazon Keyspaces gives you the performance, elasticity, and enterprise features you need to operate business-critical Cassandra workloads at scale.

48.1.1. Features

- **Compatible with Cassandra Query Language (CQL):** Amazon Keyspaces is compatible with the open-source Cassandra CQL API, so you can migrate your existing Cassandra tables to Amazon Keyspaces while continuing to use your existing application code.
- **Support for existing Apache Cassandra 2.0–licensed drivers and developer tools:** You can use existing Apache Cassandra 2.0–licensed drivers and developer tools with Amazon Keyspaces. Open-source Cassandra drivers are available for Java, Python, Ruby, .NET, Node.js, PHP, C++, and Perl.
- **On-demand capacity mode:** With on-demand capacity mode, you do not have to overprovision throughput for unexpected peak workloads. Capacity is managed automatically, and you pay for only the resources you use.
- **Provisioned capacity mode:** Provisioned capacity mode helps you optimize the price of throughput if you have predictable application traffic. Just specify the number of reads and writes per second in advance you expect your application to perform. You can use auto scaling to adjust your table's capacity automatically in response to changes in application traffic to maintain performance without overprovisioning capacity.

- **Fully managed Time to Live (TTL):** Use Time to Live (TTL) to set expiration times on rows and attributes in your Keyspaces tables, and automatically delete the records after they expire. Keyspaces TTL is fully managed, so you don't need to manage or provision additional read/write capacity. You also do not need to manage tombstones or low-level system operations such as compaction. Keyspaces deletes expired data automatically and transparently.
- **Consistent performance at any scale:** Amazon Keyspaces provides consistent single-digit-millisecond read and write performance at any scale, so you can build applications with low latency to provide a smooth user experience.
- **Elastic scaling with virtually unlimited throughput:** Amazon Keyspaces tables scale in response to actual application traffic, with virtually unlimited throughput and storage. There is no limit on the size of tables or number of rows per table.
- **Performance monitoring:** Amazon Keyspaces is integrated with Amazon CloudWatch. CloudWatch collects and processes data from Amazon Keyspaces into readable metrics, providing you with visibility into how your application is performing.
- **Fully managed and highly available data storage:** Amazon Keyspaces provides fully managed and highly available data storage. Your table data is replicated automatically three times across multiple AWS Availability Zones for durability.
- **Point-in-time recovery:** Point-in-time recovery (PITR) helps protect your Amazon Keyspaces tables from accidental write or delete operations. PITR provides continuous backups of your Amazon Keyspaces table data, and you can restore that table to any second in the preceding 35 days. You can enable PITR or initiate backup-and-restore operations with a single click in the AWS Management Console or a single API call.

48.1.2. Benefits

- **Compatible with Apache Cassandra:** Amazon Keyspaces enables you to use the Cassandra Query Language (CQL) API code, Cassandra drivers, and developer tools that you already use. Updating applications to use Amazon Keyspaces is as easy as changing the Cassandra hostname to the Amazon Keyspaces service endpoint.
- **No servers to manage:** You don’t need to provision, patch, or manage servers, so you can focus on building better applications. Tables can scale up and down automatically, and you can optimize the cost of reads and writes based on your application’s traffic patterns by choosing either on-demand or provisioned capacity mode.
- **Performance at scale:** Consistent, single-digit-millisecond response times at any scale. Build applications with virtually unlimited throughput and storage that can serve thousands of requests per second without the need for capacity planning. You can monitor performance by using Amazon CloudWatch to help keep your applications running smoothly.
- **Highly available and secure:** Amazon Keyspaces offers a 99.99% availability SLA within an AWS Region. Tables are encrypted by default and replicated three times in multiple AWS Availability Zones for high availability. You can create continuous table backups with hundreds of terabytes of data with no performance impact to your application, and recover to any point in time in the preceding 35 days.
- **Build applications that require low latency:** Process data at high speeds for applications that require single-digit-millisecond latency, such as industrial equipment maintenance, trade monitoring, fleet management, and route optimization.

- **Build applications using open-source technologies:** Build applications on AWS by using open-source Cassandra APIs and drivers that are available for a wide range of programming languages such as Java, Python, Ruby, .NET, Node.js, PHP, C++, and Perl.
- **Move your Cassandra workloads to the cloud:** Managing Cassandra tables yourself can be time consuming and expensive. With Amazon Keyspaces, you can set up, secure, and scale Cassandra tables in the AWS Cloud without managing additional infrastructure.
- **Data store for applications:** Use Amazon Keyspaces to store information about devices for Internet of Things (IoT) applications or player profiles for games. You also can use Amazon Keyspaces to store large volumes of time-series data, such as entries in a log file or the message history for a chat application.

48.2. Backup/Restore and Disaster Recovery

PITR provides you with continuous backups of your Amazon Keyspaces table data to help you protect against accidental writes or deletes. When you enable PITR, Amazon Keyspaces backs up your table data automatically with per-second granularity. You can restore your table data to any second in time in the preceding 35 days. You can enable PITR with a single click in the AWS Management Console or with a simple Cassandra Query Language (CQL) API call.

48.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

48.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/keyspaces/index.html
- **Service quotas:** https://docs.aws.amazon.com/keyspaces/latest/devguide/quotas.html
- **Service FAQs:** https://aws.amazon.com/keyspaces/faqs/

48.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/keyspaces/index.html and the following links for comprehensive technical documentation regarding this service.

- Developer Guide: Provides a conceptual overview of Amazon Keyspaces and includes detailed development instructions for using the various features.

49. Amazon Kinesis Data Firehose

49.1. Service Overview

Amazon Kinesis Data Firehose is the easiest way to load streaming data into data stores and analytics tools. Kinesis Data Firehose is a fully managed service that makes it easy to capture, transform, and load massive volumes of streaming data from hundreds of thousands of sources into Amazon S3, Amazon Redshift, Amazon OpenSearch Service (successor to Amazon Elasticsearch Service), Kinesis Data Analytics, generic HTTP endpoints, and service providers like Datadog, New Relic, MongoDB, and Splunk, enabling near real-time analytics and insights.

49.1.1. Features

- **Easy launch and configuration:** You can launch Amazon Kinesis Data Firehose and create a delivery stream to load data into Amazon S3, Amazon Redshift, Amazon OpenSearch Service, HTTP endpoints, Datadog, New Relic, MongoDB, or Splunk with just a few clicks in the AWS Management Console. You can send data to the delivery stream by calling the Firehose API, or running the Linux agent we provide on the data source. Kinesis Data Firehose then continuously loads the data into the specified destinations.
- **Load new data in near real time:** You can specify a batch size or batch interval to control how quickly data is uploaded to destinations. For example, you can set the batch interval to 60 seconds if you want to receive new data within 60 seconds of sending it to your delivery stream. Additionally, you can specify if data should be compressed. The service supports common compression algorithms including GZip, Hadoop-Compatible Snappy, Zip, and Snappy. Batching and compressing data before uploading enables you to control how quickly you receive new data at the destinations.
- **Elastic scaling to handle varying data throughput:** Once launched, your delivery streams automatically scale up and down to handle gigabytes per second or more of input data rate, and maintain data latency at levels you specify for the stream, within the limits. No intervention or maintenance is needed.
- **Apache Parquet or ORC format conversion:** Kinesis Data Firehose supports Columnar data formats such as Apache Parquet and Apache ORC are optimized for cost-effective storage and analytics using services such as Amazon Athena, Amazon Redshift Spectrum, Amazon EMR, and other Hadoop based tools. Kinesis Data Firehose can convert the format of incoming data from JSON to Parquet or ORC formats before storing the data in Amazon S3, so you can save storage and analytics costs.
- **Deliver partitioned data to S3:** Dynamically partition your streaming data before delivery to S3 using static or dynamically defined keys like “customer_id” or “transaction_id”. Kinesis Data Firehose groups data by these keys and delivers into key-unique S3 prefixes, making it easier for you to perform high performance, cost efficient analytics in S3 using Athena, EMR, and Redshift Spectrum.
- **Integrated data transformations:** You can configure Amazon Kinesis Data Firehose to prepare your streaming data before it is loaded to data stores. Simply select an AWS Lambda function from the Amazon Kinesis Data Firehose delivery stream configuration tab in the AWS Management console. Amazon Kinesis Data Firehose will automatically apply that function to every input data record and load the transformed data to destinations. Amazon Kinesis Data Firehose provides pre-built Lambda blueprints for converting common data sources such as Apache logs and system logs to JSON and CSV formats. You can use these pre-built blueprints without any change, or customize them further, or write your own custom functions. You can also configure Amazon Kinesis Data Firehose to automatically retry failed jobs and back up the raw streaming data.
- **Support for multiple data destinations:** Amazon Kinesis Data Firehose currently supports Amazon S3, Amazon Redshift, Amazon OpenSearch Service, HTTP endpoints, Datadog, New Relic, MongoDB, and Splunk as destinations. You can specify the destination Amazon S3 bucket, the Amazon Redshift table, the Amazon OpenSearch Service domain, generic HTTP endpoints, or a service provider where the data should be loaded.

- **Optional automatic encryption:** Amazon Kinesis Data Firehose provides you the option to have your data automatically encrypted after it is uploaded to the destination. As part of the delivery stream configuration, you can specify an AWS Key Management System (KMS) encryption key.
- **Metrics for monitoring performance:** Amazon Kinesis Data Firehose exposes several metrics through the console, as well as Amazon CloudWatch, including volume of data submitted, volume of data uploaded to destination, time from source to destination, the delivery stream limits, throttled records number and upload success rate. You can use these metrics to monitor the health of your delivery streams, take any necessary actions such as modifying destinations, setting alarms when getting closer to the limits, and ensure that the service is ingesting data and loading it to destinations.
- **Pay-as-you-go pricing:** With Amazon Kinesis Data Firehose, you pay only for the volume of data you transmit through the service, and if applicable, for data format conversion. You also pay for Amazon VPC delivery and data transfer when applicable. There are no minimum fees or upfront commitments. You don't need staff to operate, scale, and maintain infrastructure or custom applications to capture and load streaming data.

49.1.2. Benefits

- **Easily capture, transform, and load streaming data:** Create a delivery stream, select your destination, and start streaming real-time data with just a few clicks.
- **Automatically provisioning and scaling:** You can automatically provision and scale compute, memory, and network resources without ongoing administration.
- **Data transformation:** Transform raw streaming data into formats like Apache Parquet, and dynamically partition streaming data without building your own processing pipelines.
- **Integrated with AWS:** Connect with 30+ fully integrated AWS services and streaming destinations such as Amazon Simple Storage Service (S3) and Amazon Redshift.

49.2. Backup/Restore and Disaster Recovery

Kinesis Data Firehose uses Amazon S3 to backup all or failed only data that it attempts to deliver to your chosen destination. You can specify the S3 backup settings for your Kinesis Data Firehose delivery stream if you made one of the following choices:

- If you set Amazon S3 as the destination for your Kinesis Data Firehose delivery stream and you choose to specify an AWS Lambda function to transform data records or if you choose to convert data record formats for your delivery stream.
- If you set Amazon Redshift as the destination for your Kinesis Data Firehose delivery stream and you choose to specify an AWS Lambda function to transform data records.
- If you set any of the following services as the destination for your Kinesis Data Firehose delivery stream: Amazon OpenSearch Service, Datadog, Dynatrace, HTTP Endpoint, LogicMonitor, MongoDB Cloud, New Relic, Splunk, or Sumo Logic.

49.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

49.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/kinesis/index.html
- **Service quotas:** https://docs.aws.amazon.com/firehose/latest/dev/limits.html
- **Service FAQs:** https://aws.amazon.com/kinesis/data-firehose/faqs/?nc=sn&loc=5

49.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/kinesis/index.html and the following links for comprehensive technical documentation regarding this service.

- Developer Guide: Provides a conceptual overview of Kinesis Data Firehose and includes detailed instructions for using the service.
- API Reference: Describes all the API operations for Kinesis Data Firehose in detail. Also provides sample requests, responses, and errors for the supported web services protocols.

50. Amazon Kinesis Data Streams

50.1. Service Overview

Amazon Kinesis Data Streams is a massively scalable, durable, and low-cost streaming data service. Kinesis Data Streams can continuously capture gigabytes of data per second from hundreds of thousands of sources, such as website clickstreams, database event streams, financial transactions, social media feeds, IT logs, and location-tracking events. The collected data is available in milliseconds to allow real-time analytics use cases, such as real-time dashboards, real-time anomaly detection, dynamic pricing.

50.1.1. Features

- **Serverless:** There are no servers to manage with Amazon Kinesis Data Streams. The on-demand mode further removes the need to provision or manage throughput by automatically scaling capacity when there is an increase in workload traffic. You can get started with Kinesis Data Streams with a few clicks from the AWS Management Console.
- **Highly available and durable:** Synchronously replicate your streaming data across three Availability Zones (AZs) in an AWS Region, and store that data for up to 365 days to provide multiple layers of data loss protection.
- **Low latency:** Make your streaming data available to multiple real-time analytics applications, to Amazon Kinesis Data Analytics, or to AWS Lambda within 70 milliseconds of being collected.
- **Dedicated throughput per consumer:** You can attach up to 20 consumers to your Kinesis data stream, each with its own dedicated read throughput.
- **Choose between on-demand and provisioned capacity mode:** You can choose between on-demand mode for automated capacity management, and provisioned mode for granular control over scaling capacity up and down as needed.
- **Secure and compliant:** Encrypt sensitive data within Kinesis Data Streams to meet your regulatory and compliance needs, and securely access your data via Amazon

Virtual Private Cloud (VPC). Data can be secured at rest using server-side encryption and AWS Key Management Service (KMS) keys.

- **Integrated with other AWS services:** Use Kinesis Data Streams integrations with other AWS services, such as Amazon DynamoDB, Amazon QLDB, Amazon Aurora, AWS Database Migration Service, Amazon Cloudwatch, AWS Lambda, Amazon Kinesis Data Analytics, and Amazon Kinesis Data Firehose to build complete applications quickly.

50.1.2. Benefits

- **Managed service:** With Amazon Kinesis Data Streams, there are no servers to manage. The on-demand mode eliminates the need to provision or manage capacity required for running applications.
- **Stream gigabytes/second:** Adjust your capacity to stream gigabytes per second of data with Kinesis Data Streams. Get automatic provisioning and scaling with the on-demand mode.
- **Cost-Effective:** Pay only for what you use with Kinesis Data Streams, starting as low as $0.015 per hour. With the on-demand mode, you don't need to worry about over-provisioning.
- **Integrated with AWS Services:** Use built-in integrations with other AWS services to create analytics, serverless, and application integration solutions on AWS quickly.

50.2. Backup/Restore and Disaster Recovery

This requirement is not applicable for this service. For additional information beyond what is described herein, please refer to http://aws.amazon.com/documentation/.

50.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

50.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/kinesis/index.html
- **Service quotas:** https://docs.aws.amazon.com/streams/latest/dev/service-sizes-and-limits.html
- **Service FAQs:** https://aws.amazon.com/kinesis/data-streams/faqs/

50.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/kinesis/index.html and the following links for comprehensive technical documentation regarding this service.

Developer Guide: Provides a conceptual overview of Kinesis Data Streams and includes detailed development instructions for using the various features.

API Reference: Describes all the API operations for Kinesis Data Streams in detail. Also provides sample requests, responses, and errors for the supported web services protocols.

51. Amazon Kinesis Video Streams

51.1. Service Overview

Amazon Kinesis Video Streams makes it easy to securely stream video from connected devices to AWS for analytics, machine learning (ML), playback, and other processing. Kinesis Video Streams automatically provisions and elastically scales all the infrastructure needed to ingest streaming video data from millions of devices. It durably stores, encrypts, and indexes video data in your streams, and allows you to access your data through easy-to-use APIs. Kinesis Video Streams enables you to playback video for live and on-demand viewing, and quickly build applications that take advantage of computer vision and video analytics through integration with Amazon Rekognition Video, and libraries for ML frameworks such as Apache MxNet, TensorFlow, and OpenCV. Kinesis Video Streams also supports WebRTC, an open-source project that enables real-time media streaming and interaction between web browsers, mobile applications, and connected devices via simple APIs. Typical uses include video chat and peer-to-peer media streaming.

51.1.1. Features

- **SDKs to securely stream data from devices:** Amazon Kinesis Video Streams provides SDKs in C++ and Java that you can build and configure for your connected devices. These SDKs manage receiving data from the device's media source and securely transmitting it to a Kinesis video stream on a frame-by-frame basis in real-time. The SDK is also available as a GStreamer-plugin for constructing custom media-data flows.
- **Live and on-demand video playback with HTTP Live Streaming (HLS):** Amazon Kinesis Video Streams enables playback of the ingested video using a fully-managed HTTP Live Streaming (HLS) capability. As devices stream video into Kinesis Video Streams, you can do live playback and replay archived video on any browser or mobile platform.
- **Built-in integration with Amazon Rekognition Video:** Amazon Rekognition Video allows you to specify any of your Amazon Kinesis video streams as an input. This enables you to automatically detect and recognize faces in streaming video. Using this built-in integration, you can quickly build computer vision applications for use cases like security monitoring.
- **Real-time APIs:** Amazon Kinesis Video Streams offers easy-to-use APIs that allow you to retrieve the data from your streams on a frame-by-frame basis for building real-time applications.
- **Support for low-latency two-way media streaming with WebRTC:** Amazon Kinesis Video Streams supports WebRTC for low-latency, peer-to-peer, two-way media streaming. WebRTC is an open-source project that enables real-time media streaming and interaction between web browsers, mobile applications, and connected devices via simple APIs. Kinesis Video Streams includes managed end-points for WebRTC signalling that allows applications to securely connect with each other for peer-to-peer live media streaming. Next, it includes managed end-points for TURN that enables media relay via the cloud when applications cannot stream peer-to-peer media. It also includes managed end-points for STUN that enables applications to discover their public IP address when they are located behind a NAT or a firewall. Additionally, it provides easy to use SDKs to enable camera IoT devices with WebRTC capabilities. Finally, it provides client SDKs for Android, iOS, and for Web applications to integrate Kinesis

Video Streams WebRTC signalling, TURN, and STUN capabilities with any WebRTC compliant mobile or web player.

- **Durable storage:** Amazon Kinesis Video Streams uses Amazon S3 as the underlying data store, which means your data is stored durably and reliably. You can set and control retention periods on a per-stream basis, allowing you to cost-effectively store the data in your streams for a limited time period or indefinitely. You can change the stream retention period at any point.
- **Pay per use:** With Amazon Kinesis Video Streams, you pay only for the volume of data you ingest, store, and consume through the service. There are no upfront costs or minimum fees, and you need not worry about paying for idle video streams.
- **Automatic indexing for retrieval:** Amazon Kinesis Video Streams automatically indexes the data you store in your video streams based on timestamps generated by the device, or timestamps generated by Kinesis Video Streams when it receives the video. You can combine stream-level tags with timestamps to easily search and retrieve specific video fragments for playback, analytics, and other processing.
- **Video stream parser library:** Amazon Kinesis Video Streams offers a stream parser library that you can use within your applications to easily retrieve frame-level objects, extract and collect metadata attached to fragments, merge consecutive fragments, and more. It allows you to readily integrate popular ML frameworks such as Apache MxNet, TensorFlow, and OpenCV.
- **Automatic data encryption in transit and at rest:** The Amazon Kinesis Video Streams SDK encrypts the frames and fragments generated by the device's hardware for secure streaming using Transport Layer Security (TLS), a protocol that provides privacy and data integrity between two communicating applications. Amazon Kinesis Video Streams automatically encrypts the data you put into your video streams using AWS Key Management Service (KMS), helping you protect your data at rest. Data is encrypted before it is written to the Kinesis Video Streams storage, and it is decrypted after it is retrieved from storage. As a result, your data is always encrypted at rest within the stream.

51.1.2. Benefits

- **Stream video from millions of devices:** Amazon Kinesis Video Streams provides SDKs that make it easy for devices to securely stream media to AWS for playback, storage, analytics, machine learning, and other processing. Kinesis Video Streams can ingest data from edge devices, smartphones, security cameras, and other data sources such as RADARs, LIDARs, drones, satellites, dash cams, and depth-sensors.
- **Build real-time vision and video-enabled apps:** Easily build applications with real-time computer vision capabilities through integration with Amazon Rekognition Video, and with real-time video analytics capabilities using popular open source machine learning frameworks.
- **Playback live and recorded video streams:** Easily stream live and recorded media from your Kinesis video streams to your browser or mobile application using the Kinesis Video Streams HTTP Live Streaming (HLS) capability.
- **Build apps with two-way, real-time media streaming:** Amazon Kinesis Video Streams supports the open-source project WebRTC for two-way, real-time media streaming between web browsers, mobile applications, and connected devices. With support for WebRTC, you can use simple APIs to build rich applications like video chat and peer-to-

peer data sharing with ultra-low latency and two-way communication between your applications and connected devices.

- **Secure:** Amazon Kinesis Video Streams allows you to control access to your streams using AWS Identity and Access Management (IAM). It helps you protect your data by automatically encrypting the data at rest using AWS Key Management Service (KMS) and in transit using the industry-standard Transport Layer Security (TLS) protocol.
- **Durable, searchable storage:** Amazon Kinesis Video Streams uses Amazon S3 as the underlying data store, which means your data is stored durably and reliably. Kinesis Video Streams enables you to quickly search and retrieve video fragments based on device and service generated timestamps.
- **No infrastructure to manage:** Amazon Kinesis Video Streams manages all the infrastructure for you. You don't have to worry about configuration, software updates, failures, or scaling infrastructure as the number of streams and consuming applications grows. Kinesis Video Streams handles all the administration and maintenance required to manage your streams, so you can focus your time on building innovative applications.

51.2. Backup/Restore and Disaster Recovery

This requirement is not applicable for this service. For additional information beyond what is described herein, please refer to http://aws.amazon.com/documentation/.

51.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

51.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/kinesis/index.html
- **Service quotas:** https://docs.aws.amazon.com/general/latest/gr/akv.html
- **Service FAQs:** https://aws.amazon.com/kinesis/video-streams/faqs/?nc=sn&loc=5

51.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/kinesis/index.html and the following links for comprehensive technical documentation regarding this service.

- Kinesis Video Streams Developer Guide: Provides a conceptual overview of Kinesis Video Streams, includes detailed instructions for using the various features, and provides a complete API reference for developers.
- Kinesis Video Streams with WebRTC Developer Guide: Provides a conceptual overview of Kinesis Video Streams with WebRTC capability and includes detailed instructions for using the various features.
- API Reference: Describes all the API operations for Kinesis Video Streams in detail. Also provides sample requests, responses, and errors for the supported web services protocols.

52. Amazon Lex

52.1. Service Overview

Powered by the same technology as Alexa, Amazon Lex provides you with the tools to tackle challenging deep learning problems, such as speech recognition and language understanding, through an easy-to-use fully managed service. Amazon Lex integrates with AWS Lambda which you can use to easily trigger functions for execution of your back-end business logic for data retrieval and updates. Once built, your bot can be deployed directly to chat platforms, mobile clients, and IoT devices. You can also use the reports provided to track metrics for your bot. Amazon Lex provides a scalable, secure, easy to use, end-to-end solution to build, publish and monitor your bots.

52.1.1. Features

- **High quality speech recognition and natural language understanding:** Amazon Lex provides automatic speech recognition and natural language understanding technologies to create a Speech Language Understanding system. Amazon Lex uses the same proven technology that powers Alexa. Amazon Lex is able to learn the multiple ways users can express their intent based on a few sample utterances provided by the developer. The speech language understanding system takes natural language speech and text input, understands the intent behind the input, and fulfills the user intent by invoking the appropriate response.
- **Context management:** As the conversation develops, being able to classify utterances accurately requires managing context across multi-turn conversations. Amazon Lex supports context management natively, so you can manage the context directly without the need for custom code. As initial prerequisite intents are filled, you can create “contexts” to invoke related intents. This simplifies bot design and expedites the creation of conversational experiences.
- **8 kHz telephony audio support:** The Amazon Lex speech recognition engine has been trained on telephony audio (8 kHz sampling rate), providing increased speech recognition accuracy for telephony use-cases. When building a conversational bot with Amazon Lex, the 8 kHz support allows for higher fidelity with telephone speech interactions, such as through a contact center application or helpdesk.
- **Multi-turn dialog:** Amazon Lex bots provide the ability for multi-turn conversations. Once an intent has been identified, users will be prompted for information that is required for the intent to be fulfilled (for example, if “Book hotel” is the intent, the user is prompted for the location, check-in date, number of nights, etc.). Amazon Lex gives you an easy way to build multi-turn conversations for your chatbots. You simply list the slots/parameters you want to collect from your bot users, as well as the corresponding prompts, and Amazon Lex takes care of orchestrating the dialogue by prompting for the appropriate slot.
- **Powerful Lifecycle Management Capabilities:** Amazon Lex lets you apply versioning to the Intents, Slot Types, and Bots that you create. Versioning and rollback mechanisms enables you to easily maintain code as you test and deploy in a multi-developer environment. You can create multiple aliases for each Amazon Lex bot and associate different versions to each such as “production,” “development,” and “test”. This allows you to continue making improvements and changes to the bot and release them as new versions under one alias. This removes the need to update all the clients when a new version of the bot is deployed.

- **One-click deployment to multiple platforms:** Amazon Lex allows you to easily publish your bot to chat services directly from the Amazon Lex console, reducing multi-platform development efforts. Rich formatting capabilities provide an intuitive user experience tailored to chat platforms like Facebook Messenger, Slack, and Twilio SMS.
- **Enhanced console experience:** The Lex V2 console experience makes it easier to build, deploy and manage conversational experiences. With Lex V2, you can add a new language to a bot at any time and manage all the languages through the lifecycle of design, test and deployment as a single resource. A simplified information architecture lets you efficiently manage your bot versions. Capabilities such as a 'Conversation Flow', saving of partially configured bots and bulk upload of utterances simplify the process and give you more flexibility.
- **Streaming conversations:** Natural conversations are punctuated with pauses and interruptions. For example, a caller may ask to pause the conversation or hold the line while looking up the necessary information before answering a question to retrieve credit card details when providing bill payments. With streaming conversation APIs, you can pause a conversation and handle interruptions directly as you configure the bot. You can quickly enhance the conversational capability of virtual contact center agents or smart assistants.
- **Service Integrations:** Amazon Lex is integrated with services like Amazon Kendra, Amazon Polly, AWS Lambda and Amazon Connect.

52.1.2. Benefits

- **AI Capability:** Easily add AI that understands intent, maintains context, and automates simple tasks across many languages.
- **Omnichannel:** Design and deploy omnichannel conversational AI in one click, without worrying about hardware or infrastructure.
- **Integrated with AWS Services:** Connect seamlessly with other AWS services to query data, execute business logic, monitor performance, and more.
- **Cost Effective:** Pay only for speech and text requests with no upfront costs or minimum fees.

52.2. Backup/Restore and Disaster Recovery

This requirement is not applicable for this service. For additional information beyond what is described herein, please refer to http://aws.amazon.com/documentation/.

52.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

52.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/lex/index.html
- **Service quotas:** https://docs.aws.amazon.com/lex/latest/dg/gl-limits.html
- **Service FAQs:** https://aws.amazon.com/lex/faqs/?nc=sn&loc=6

52.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/lex/index.html and the following links for comprehensive technical documentation regarding this service.

- V2 Developer Guide: Provides a conceptual overview of Amazon Lex V2. Includes detailed instructions for using various features, such as streaming conversations, and provides a complete reference for the V2 APIs.
- V1 Developer Guide: Provides a conceptual overview of Amazon Lex V1, includes detailed instructions for using the various features, and provides a complete V1 API reference for developers.

53. Amazon Lightsail

53.1. Service Overview

Amazon Lightsail provides easy-to-use cloud resources to get your web application or websites up and running in just a few clicks. Lightsail offers simplified services such as instances, containers, databases, storage, and more. With Lightsail, you can easily spin up websites or applications using pre-configured blueprints like WordPress, Prestashop, or LAMP. You can use Lightsail features to simply host static content, connect your content to an audience around the globe, or get your Windows Business server up and running. The Lightsail console guides you through the configuration process, and in many cases, has components already configured.

53.1.1. Features

- **Instances:** Lightsail offers virtual servers (instances) that are easy to set up and backed by the power and reliability of AWS. You can launch your website, web application, or project in minutes, and manage your instance from the intuitive Lightsail console or API.
- **Containers:** Lightsail offers an easy way to run containers in the cloud. With Lightsail Container Service, customers can now run and securely access containerized applications from the internet with just a few steps.
- **Simplified load balancers:** Lightsail's simplified load balancing routes web traffic across your instances so your websites and applications can accommodate variations in traffic, protected against outages, and deliver a seamless visitor experience.
- **Managed databases:** Lightsail offers a fully configured MySQL or PostgreSQL databases plan that includes memory, processing, storage, and transfer allowance. With Lightsail managed databases, you can easily scale your databases independently of your virtual servers, improve application availability, or run standalone databases in the cloud. You can also deploy multi-tiered applications, all within Lightsail, by creating multiple instances connected to a central managed database, and a load balancer that directs traffic to the instances.
- **Block and object storage:** Amazon Lightsail offers both block and object storage. You can scale your storage quickly and easily with highly available SSD-backed storage for your Linux or Windows virtual server. For Lightsail Object Storage, you can easily host static content on the cloud.
- **CDN distributions:** Lightsail enables content delivery network (CDN) distributions, which are built on the same infrastructure as Amazon CloudFront. This allows you to easily distribute your content to a global audience by setting up proxy servers across the

world, so that your users across the globe can access your website geographically closer to them, thus reducing latency.

- **Lightsail virtual servers:** Lightsail offers virtual servers (instances) that are easy to set up and backed by the power and reliability of AWS. You can launch your website, web application, or project in minutes, and manage your instance from the intuitive Lightsail console or API.
- **Upgrade to EC2:** As your cloud ideas expand, you can easily move to EC2 with a simple, guided experience. With this feature, Lightsail offers you the comfort of knowing that as you grow your website or application we can help you scale in a way that fits your needs.
- **Access to AWS services:** Amazon Lightsail uses a focused set of features like instances, managed databases and load balancers to make it easier to get started. But that doesn't mean you're limited to those options –you can integrate your Lightsail project with some of the 90+ other services in AWS through Amazon VPC peering.

53.1.2. Benefits

- **Easy to use:** Create a website or application in just a few clicks. Automatically configure networking, access, and security environments.
- **Scalable:** Easily scale as you grow—or migrate your resources to other AWS services, such as Amazon EC2.
- **Secure:** Leverage the security and reliability of the world's leading cloud platform.

53.2. Backup/Restore and Disaster Recovery

This requirement is not applicable for this service. For additional information beyond what is described herein, please refer to http://aws.amazon.com/documentation/.

53.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

53.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/lightsail/
- **Service quotas:** https://docs.aws.amazon.com/general/latest/gr/lightsail.html
- **Service FAQs:** https://aws.amazon.com/lightsail/faq/

53.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/lightsail/ and the following links for comprehensive technical documentation regarding this service.

- Developer Guide: Describes how to get started with Lightsail to create a development environment or an application.
- API Reference: Describes all the API operations for Lightsail in detail. Also provides sample requests, responses, and errors for the supported web services protocols.

54. Amazon Location Service

54.1. Service Overview

Amazon Location Service makes it easy for developers to add location functionality to applications without compromising data security and user privacy. Location data is a vital ingredient in today's applications, enabling capabilities ranging from asset tracking to location-based marketing. However, developers face significant barriers when integrating location functionality into their applications. This includes cost, privacy and security compromises, and tedious and slow integration work.

Amazon Location Service provides affordable data, tracking and geofencing capabilities, and native integrations with AWS services, so you can create sophisticated location-enabled applications quickly, without the high cost of custom development. You retain control of your location data with Amazon Location, and you can combine proprietary data with data from the service. Amazon Location provides cost-effective location-based services (LBS) using high-quality data from global, trusted providers Esri and HERE.

54.1.1. Features

- **Maps:** Maps help you visualize location information and are the foundations of many location-based service capabilities. Amazon Location Service provides map tiles of different styles sourced from global location data providers Esri and HERE. The map tiles from HERE and Esri are trusted by millions of customers worldwide, and have been continuously fine-tuned over the decades for a wide range of customer applications.
- **Places:** Amazon Location Service Places enables your application to offer point-of-interest search functionality, convert addresses into geographic coordinates in latitude and longitude (geocoding), and convert a coordinate into a street address (reverse geocoding). Amazon Location Service sources high-quality geospatial data from Esri and HERE, so you can use them to improve the quality and the versatility of your address database.
- **Routes:** With Amazon Location Routes, your application can request the travel time, distance, and directions between a departure point and one or more destinations, with specific travel restrictions such as Truck mode, vehicle dimension, and avoidances. Amazon Location Service provides routing data sourced from global location data providers Esri and HERE. This enables your application to obtain accurate estimates of travel time based on up-to-date roadway information and live traffic information.
- **Tracking:** With Amazon Location Service Trackers, you can instantly retrieve current and historical location of devices running your tracking-enabled application. This gives you the ability to use the location history of your devices to optimize operations such as inventory placements, manufacturing sequencing, or delivery dispatch.
- **Geofencing:** Amazon Location Service Geofences gives your application the ability to detect and act when a tracked device enters or exits a geographical boundary you define as a geofence. When a breach of the geofence is detected, Amazon Location Service automatically sends the entry or exit event to Amazon EventBridge, which can then trigger downstream actions, such as sending a notification to a restaurant that a delivery driver is nearby.
- **Data Control:** With Amazon Location Service, you retain control of your organization's data. Amazon Location Service anonymizes all queries sent to data providers by

removing customer metadata and account information. Additionally, sensitive tracking and geofencing location information, such as facility, asset, or personnel locations is always encrypted at rest and in transit.

- **Data Rights:** With Amazon Location Service, you retain all rights to your organization's data. Amazon Location Service licensing terms do not grant Amazon or third parties rights to sell your data or use it for advertising, for instance when you display your data on a map, perform a search, or request a route.
- **Secure Access:** Amazon Location Service integrates with proven AWS security services, including AWS Identity and Access Management (IAM) and Amazon Cognito, so you can move to production faster by using existing identity management and authentication tools to help keep your application secure both for your administrators and your application users.
- **Integrated Monitoring and Management:** Amazon Location Service is integrated with AWS CloudFormation, Amazon CloudWatch, AWS CloudTrail, and Amazon EventBridge, so you can efficiently provision and manage resources, monitor health metrics, and automatically act upon events.
- **Developer Tools:** Amazon Location Service offers a variety of tools for developers to build location-enabled applications. These include the standard AWS SDKs, front-end mobile and web SDKs, and sample code to combine it with open source libraries such as MapLibre.

54.1.2. Benefits

- **Privacy and security:** With Amazon Location Service, you retain control of your organization's data. Amazon Location Service anonymizes all queries sent to data providers by removing customer metadata and account information. Sensitive tracking and geofencing location information, such as facility, asset, and personnel locations, is processed and retained only in your account. This helps you shield sensitive information from third parties, protect user privacy, and reduce your application's security risks. With Amazon Location Service, neither Amazon nor third parties have rights to sell your data or use it for advertising.
- **High-quality and cost-effective:** Amazon Location Service provides high-quality geospatial data from established, global providers Esri and HERE. Application developers trust Esri and HERE's data to route millions of vehicles and power hundreds of thousands of mobile applications and websites today. Now you can affordably build applications using their data through Amazon Location Service, along with tracking and geofencing capabilities that you don't have to build in-house.
- **Simple access to location data:** With the fully managed Amazon Location Service, integrating geospatial information into your application is as easy as calling the Amazon Location Service application programming interface (API). The API is consistent across LBS data providers, so you don't need to learn and integrate multiple APIs to get the data you want for different use cases or geographies. To help you explore and try geolocation functions, the Amazon Location Service console includes a visual, interactive learning tool. You can get started quickly with the Amazon Location Service back-end and front-end software development kits (SDKs), sample code for tasks such as data visualization, and solution guides for asset tracking, geomarketing, and delivery.

- **Shorter time to production:** Amazon Location Service helps you move applications from experimentation to production faster by providing integrations with Amazon CloudFormation, AWS CloudTrail, Amazon CloudWatch, Amazon EventBridge, AWS Identity and Access Management (IAM), AWS Key Management Service (KMS), and includes tagging functionality. Native integrations with AWS services provide you with built-in features from health and monitoring of your application, using existing security controls, to inclusion in event-driven architectures. With Amazon Location Service you can apply your organization's cloud best practices and avoid establishing new standards for additional vendors.

54.2. Backup/Restore and Disaster Recovery

This requirement is not applicable for this service. For additional information beyond what is described herein, please refer to http://aws.amazon.com/documentation/.

54.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

54.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/location/
- **Service quotas:** https://docs.aws.amazon.com/location/latest/developerguide/location-quotas.html
- **Service FAQs:** https://aws.amazon.com/location/faqs/

54.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/location/ and the following links for comprehensive technical documentation regarding this service.

- Developer Guide: Provides a conceptual overview and practical examples to help you understand the features provided by Amazon Location and how to use them.
- Sample apps for tutorials: The GitHub repository for Amazon Location Service sample apps described in the developer guide tutorials.

55. Amazon Lookout for Equipment

55.1. Service Overview

Successfully implementing predictive maintenance requires using the specific data collected from all of your machine sensors, under your unique operating conditions and then applying machine learning (ML) to enable highly accurate predictions. However, implementing an ML solution for your equipment can be difficult and time-consuming.

Amazon Lookout for Equipment analyses the data from the sensors on your equipment (e.g. pressure in a generator, flow rate of a compressor, revolutions per minute of fans), to automatically train a machine learning model based on just your data, for your equipment – with no ML expertise required. Lookout for Equipment uses your unique ML model to analyse incoming sensor data in real-time and accurately identify early warning signs that could lead to

machine failures. This means you can detect equipment abnormalities with speed and precision, quickly diagnose issues, take action to reduce expensive downtime, and reduce false alerts.

55.1.1. Features

- **Sensor and data quality evaluation:** Time series data coming from sensors on industrial equipment can be highly erratic and the quality/usability of each sensor is difficult to determine. Amazon Lookout for Equipment will derive key statistics on ingested data from each sensor, grade the overall data quality and give a justification for its grade. This output advises a user on which sensors are preferred inputs.
- **Automated machine learning:** Amazon Lookout for Equipment will automatically leverage data from up to 300 sensors at once, as well as maintenance history, in order to search through up to 28,000 possible algorithm combinations and determine the optimal multi-variate model that best learns the normal behavior of the specified equipment.
- **Model diagnostics:** For each detected abnormal behavior, Amazon Lookout for Equipment will understand the behavior and indicate to a user which sensors are impacting the issue and what is happening in each of those sensors. Customers can use this information to diagnose the problem and take corrective action.
- **Continuous monitoring:** Easily deploy the developed model on real time data by setting up an inference scheduler. The scheduler will run inferencing on newly generated sensor data at intervals as low as once per minute and as high as once per hour.

55.1.2. Benefits

- **High accuracy results:** Amazon Lookout for Equipment uses data from your existing sensors to build a custom ML model specific to your equipment. It handles data from up to 300 sensors in one model, along with historical logs, to give you accurate alerts when your equipment behaves abnormally.
- **Respond to issues faster and with precision:** Amazon Lookout for Equipment automatically monitors your equipment and identifies any anomalies compared to healthy operation. Lookout for Equipment can then pinpoint the sensor or sensors indicating anomalies, enabling you to respond quickly.
- **Accelerate issue resolution:** You can use data from Amazon Lookout for Equipment to set up automatic actions to be taken when anomalies are detected, such as filing a trouble ticket, or sending an automatic alarm that notifies you immediately of any issues. The data from Lookout for Equipment can be integrated into your existing monitoring software or you can use AWS IoT SiteWise to collect, store, organize and monitor data. Lookout for Equipment can also work with data from common industrial data systems such as OSIsoft.
- **Improve accuracy of alerts over time:** Amazon Lookout for Equipment continuously improves model performance and accuracy of alerts by incorporating human review feedback on the anomalies detected and learning expected operational usage trends.

55.2. Backup/Restore and Disaster Recovery

This requirement is not applicable for this service. For additional information beyond what is described herein, please refer to http://aws.amazon.com/documentation/.

55.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

55.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/lookout-for-equipment/index.html
- **Service quotas:** https://docs.aws.amazon.com/lookout-for-equipment/latest/ug/guidelines-and-limits.html
- **Service FAQs:** https://aws.amazon.com/lookout-for-equipment/faqs/

55.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/lookout-for-equipment/index.html and the following link for comprehensive technical documentation regarding this service.

- User Guide: Provides a conceptual overview of Amazon Lookout for Equipment, offers detailed instructions for using various features, and includes a complete API reference for developers.

56. Amazon Lookout for Metrics

56.1. Service Overview

Amazon Lookout for Metrics uses machine learning (ML) to automatically detect and diagnose anomalies (i.e. outliers from the norm) in business and operational data, such as a sudden dip in sales revenue or customer acquisition rates. In a couple of clicks, you can connect Amazon Lookout for Metrics to popular data stores like Amazon S3, Amazon Redshift, and Amazon Relational Database Service (RDS), as well as third-party SaaS applications, such as Salesforce, ServiceNow, Zendesk, and Marketo, and start monitoring metrics that are important to your business. Amazon Lookout for Metrics automatically inspects and prepares the data from these sources to detect anomalies with greater speed and accuracy than traditional methods used for anomaly detection. You can also provide feedback on detected anomalies to tune the results and improve accuracy over time. Amazon Lookout for Metrics makes it easy to diagnose detected anomalies by grouping together anomalies that are related to the same event and sending an alert that includes a summary of the potential root cause. It also ranks anomalies in order of severity so that you can prioritize your attention to what matters the most to your business.

56.1.1. Features

- **ML-powered anomaly detection:** Amazon Lookout for Metrics monitors metrics and detects anomalies with high accuracy using ML technology. It uses specialized ML models to detect anomalies based on the characteristics of your data. You don’t need ML experience to use Amazon Lookout for Metrics.
- **Anomaly Grouping and Ranking:** Amazon Lookout for Metrics automatically groups anomalies that might be related to the same event and ranks them in the order of severity, so that you can focus on what matters the most at any given time.

- **Tunable results:** You can provide feedback on the relevance of detected anomalies. This feedback helps Amazon Lookout for Metrics tune the results for your metrics. The accuracy of the results continues to increase over time as you provide more feedback.
- **Data source compatibility:** Amazon Lookout for Metrics seamlessly connects to popular data sources, including Amazon Simple Storage Service (Amazon S3), Amazon Redshift, Amazon Relational Database Service (Amazon RDS), Amazon CloudWatch, Salesforce, Marketo, Dynatrace, Singular, Zendesk, Servicenow, Infor Nexus, Trendmicro, Veeva, Google Analytics, and Amplitude.
- **Custom, automated alerts:** After Amazon Lookout for Metrics creates an anomaly detection model, or detector, you can attach alerts to it using supported output connectors such as Amazon Simple Notification Service (Amazon SNS), AWS Lambda functions, Datadog, PagerDuty, Webhooks, and Slack. You can create custom alerts to notify you when Amazon Lookout for Metrics detects an anomaly of a specified severity level.

56.1.2. Benefits

- **Detect anomalies with better accuracy:** Amazon Lookout for Metrics uses ML to accurately detect anomalies in business metrics and reduce false positives. You can provide feedback on detected anomalies to tune the results and continuously improve accuracy and control.
- **Diagnose root cause faster:** Amazon Lookout for Metrics makes it easier to diagnose the root cause of important anomalies by automatically grouping related anomalies together, summarizing potential root causes, and ranking them in the order of severity.
- **Easily integrate popular data sources and SaaS applications:** Amazon Lookout for Metrics connects seamlessly to popular AWS databases and provides pre-built connectors to third-party SaaS applications. You can get started monitoring metrics and detecting anomalies in just a few clicks.
- **Automate customized alerts and actions:** You can easily connect Amazon Lookout for Metrics with event and notification services like Amazon Simple Notification Service and AWS Lambda to automate customized alerts and actions when anomalies are detected, such as filing a trouble ticket.

56.2. Backup/Restore and Disaster Recovery

This requirement is not applicable for this service. For additional information beyond what is described herein, please refer to http://aws.amazon.com/documentation/.

56.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

56.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/lookout-for-metrics/index.html
- **Service quotas:** https://docs.aws.amazon.com/lookoutmetrics/latest/dev/quotas.html
- **Service FAQs:** https://aws.amazon.com/lookout-for-metrics/faqs/

56.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/lookout-for-metrics/index.html and the following links for comprehensive technical documentation regarding this service.

- Amazon Lookout for Metrics Developer Guide: Provides a conceptual overview of Amazon Lookout for Metrics and detailed instructions for using its features.
- Amazon Lookout for Metrics API Reference: Describes all the API operations for Amazon Lookout for Metrics in detail.

57. Amazon Lookout for Vision

57.1. Service Overview

You can use Amazon Lookout for Vision to identify missing components in products, damage to vehicles or structures, irregularities in production lines, minuscule defects in silicon wafers, and other similar problems. It uses machine learning (ML) to see and understand images from any camera as a person would, but with an even higher degree of accuracy and at a much larger scale.

57.1.1. Features

- **Dashboard view:** The Amazon Lookout for Vision console provides a holistic view across all your production lines with an easy-to-use dashboard. The dashboard shows the projects by most defects, recent defects, and highest anomaly ratio, which enables you to quickly identify the production lines and processes that need immediate attention.
- **Simplified labelling:** The Amazon Lookout for Vision console provides a visual interface to label your images quickly and simply, by applying a normal or anomaly label to the entire image with a click of a button.
- **Quick evaluation:** Evaluate your anomaly detection model's performance on your test dataset. If you do not provide your own test dataset, Amazon Lookout for Vision can automatically create a test dataset for you to evaluate your model's performance.
- **Trial anomaly detection tasks and feedback:** You can instantly run test detection tasks on additional images to get normal or anomaly predictions using your model. You can track your predictions, correct any mistakes, and provide the feedback to retrain newer models to improve anomaly detection accuracy.
- **Using your trained models at the edge:** You can use your trained Amazon Lookout for Vision models on a hardware device of your choice.
- **Manufacturing line integration:** You can integrate Amazon Lookout for Vision with your manufacturing lines and implement automated visual inspection workflows for your use cases with just a few clicks in the Amazon Lookout for Vision console and a few API parameters.

57.1.2. Benefits

- **Easy to use:** Easily create a machine learning (ML) model to spot anomalies from your live process line with as few as 30 images.
- **Real time:** Identify visual anomalies in real time to reduce and prevent defects and improve production quality.

- **Efficient:** Prevent unplanned downtime and reduce operational costs by using visual inspection data to spot potential issues and take corrective action.

57.2. Backup/Restore and Disaster Recovery

This requirement is not applicable for this service. For additional information beyond what is described herein, please refer to http://aws.amazon.com/documentation/.

57.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

57.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/lookout-for-vision/
- **Service quotas:** https://docs.aws.amazon.com/lookout-for-vision/latest/developer-guide/limits.html
- **Service FAQs:** https://aws.amazon.com/lookout-for-vision/faqs/

57.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/lookout-for-vision/ and the following links for comprehensive technical documentation regarding this service.

- Developer Guide: Describes all Amazon Lookout for Vision concepts and provides instructions on using the various features with the console, command line interface, and the AWS SDK.
- API Reference: Describes all the API operations for Amazon Lookout for Vision in detail.

58. Amazon Macie

58.1. Service Overview

Amazon Macie is a fully managed data security and data privacy service that uses machine learning and pattern matching to discover and protect your sensitive data in AWS.

As organizations manage growing volumes of data, identifying and protecting their sensitive data at scale can become increasingly complex, expensive, and time-consuming. Amazon Macie automates the discovery of sensitive data at scale and lowers the cost of protecting your data. Macie automatically provides an inventory of Amazon S3 buckets including a list of unencrypted buckets, publicly accessible buckets, and buckets shared with AWS accounts outside those you have defined in AWS Organizations. Then, Macie applies machine learning and pattern matching techniques to the buckets you select to identify and alert you to sensitive data, such as personally identifiable information (PII).

Macie's alerts, or findings, can be searched and filtered in the AWS Management Console and sent to Amazon EventBridge, formerly called Amazon CloudWatch Events, for easy integration with existing workflow or event management systems, or to be used in combination with AWS services, such as AWS Step Functions to take automated remediation actions. This can help you meet regulations, such as the Health Insurance Portability and Accountability Act (HIPAA) and General Data Privacy Regulation (GDPR). You can get started with Amazon Macie by

leveraging the 30-day free trial for bucket evaluation. The trial includes 30-days of Amazon S3 bucket inventory and bucket-level security and access control assessment at no cost. Note that sensitive data discovery is not included in the 30-day free trial for bucket evaluation.

58.1.1. Features

- **Ongoing evaluation of your Amazon S3 environment:** Amazon Macie continually evaluates your Amazon S3 environment and provides an S3 resource summary across all of your accounts. You can search, filter, and sort buckets by metadata variables, such as bucket names, tags, and security controls like encryption status or public accessibility. For any unencrypted buckets, publicly accessible buckets, or buckets shared with AWS accounts outside those you have defined in AWS Organizations, you can be alerted in order to take action.
- **Scalable on-demand and automated sensitive data discovery jobs:** Amazon Macie allows you to run one-time, daily, weekly, or monthly sensitive data discovery jobs for all, or a subset of objects in an Amazon S3 bucket. For sensitive data discovery jobs, Amazon Macie automatically tracks changes to the bucket and only evaluates new or modified objects over time.
- **Fully managed sensitive data types:** Amazon Macie maintains a growing list of sensitive data types that include common personally identifiable information (PII) and other sensitive data types as defined by data privacy regulations, such as GDPR, PCI-DSS, and HIPAA. These data types use various data detection techniques including machine learning and are continually added to and improved upon over time.
- **Custom-defined sensitive data types:** Amazon Macie provides you the ability to add custom-defined data types using regular expressions to enable Macie to discover proprietary or unique sensitive data for your business.
- **Detailed and actionable security and sensitive data discovery findings:** Macie reduces alert volume and speeds up triage by consolidating findings by object or bucket. Based on severity level, Macie findings are prioritized and each finding includes details, such as the sensitive data type, tags, public accessibility, and encryption status. Findings are retained for 30-days and are available in the AWS Management Console or through the API. The full sensitive data discovery details are automatically written to a customer-owned S3 bucket for long-term retention.
- **One-click deployment with no upfront data source integration:** With one-click in the AWS Management Console or a single API call, you can enable Amazon Macie in a single account. With a few more clicks in the console, you can enable Macie across multiple accounts. Once enabled, Macie generates an ongoing Amazon S3 resource summary across accounts that includes bucket and object counts as well as the bucket-level security and access controls.
- **Multi-account support and integration with AWS Organizations:** In the multi-account configuration, a single Macie administrator account can manage all member accounts, including the creation and administration of sensitive data discovery jobs across accounts. Amazon Macie supports multiple accounts through AWS Organizations integration as well as natively within Macie. Security and sensitive data discovery findings are aggregated in the Macie administrator account and sent to Amazon CloudWatch Events. Now using one account, you can integrate with event management, workflow, and ticketing systems or use Macie findings with AWS Step Functions to automate remediation actions.

58.1.2. Benefits

- **Discover your sensitive data at scale:** Amazon Macie uses machine learning and pattern matching to cost efficiently discover sensitive data at scale. Macie automatically detects a large and growing list of sensitive data types, including personal identifiable information (PII) such as names, addresses, and credit card numbers. The service also allows you to define your own custom sensitive data types so you can discover and protect the sensitive data that may be unique to your business or use case.
- **Visibility of your data security posture:** Amazon Macie gives you constant visibility of the data security and data privacy of your data stored in Amazon S3. Macie automatically and continually evaluates all of your S3 buckets and alerts you to any unencrypted buckets, publicly accessible buckets, or buckets shared with AWS accounts outside those you have defined in the AWS Organizations. Macie provides native multi-account support so you can view your data security posture across your entire S3 environment from a single Macie administrator account.
- **Easy to setup and manage:** Getting started with Amazon Macie is fast and easy with one-click in the AWS Management Console or a single API call. Macie provides multi-account support using AWS Organizations, so you can enable Macie across all of your accounts with a few clicks. Macie maintains a fully-managed set of sensitive data types, so there is no custom configuration required.

58.2. Backup/Restore and Disaster Recovery

This requirement is not applicable for this service. For additional information beyond what is described herein, please refer to http://aws.amazon.com/documentation/.

58.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

58.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/macie/
- **Service quotas:** https://docs.aws.amazon.com/macie/latest/user/macie-quotas.html
- **Service FAQs:** https://aws.amazon.com/macie/faq/

58.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/macie/ and the following links for comprehensive technical documentation regarding this service.

- User Guide: Describes key concepts and provides detailed instructions for using Amazon Macie.
- API Reference: Describes the API operations for Amazon Macie.
- Amazon Macie in the AWS CLI Reference: Describes the commands for Amazon Macie.

59. Amazon Managed Blockchain

59.1. Service Overview

Amazon Managed Blockchain is a fully managed service that allows you to join public networks or set up and manage scalable private networks with just a few clicks. Amazon Managed Blockchain eliminates the overhead required to create the network or join a public network, and automatically scales to meet the demands of thousands of applications running millions of transactions. Once your network is up and running, Managed Blockchain makes it easy to manage and maintain your blockchain network. It manages your certificates and lets you easily invite new members to join the network.

59.1.1. Features

- **Setting up a network in a few clicks:** Getting started with Amazon Managed Blockchain is easy—you can launch a blockchain network in minutes without additional configuration. Then configure your network membership and launch blockchain peer nodes using the AWS Management Console. You can invite other AWS accounts to join your blockchain network, or you can create additional members in your AWS account to build a simulated network for testing.
- **Adding new members with voting:** When building permissioned blockchain networks, enabling existing members to vote on the addition (or removal) of new members can require custom development and permissions management. To make this easier, Amazon Managed Blockchain provides a voting API that enables members in a blockchain network to quickly vote on proposals for adding or removing new members.
- **Join a public network in minutes:** Joining a public network is easy – you can choose the public network that you want to join and then provision a peer node using the AWS Management Console. Amazon Managed Blockchain provides secure networking, fast and reliable syncs to the blockchain network, durable elastic storage for ledger data, encryption at rest and transport, and secure access to open-source APIs.
- **Choice of Frameworks:** With Amazon Managed Blockchain you can choose between two popular blockchain frameworks, Hyperledger Fabric and Ethereum, so you can choose the framework that best fits your needs.
- **Support for Hyperledger Fabric:** Hyperledger Fabric is an open source blockchain framework from the Linux Foundation that enables you to write blockchain applications and offers access control and permissions for data on the blockchain. With it, you can easily create a private blockchain network and limit the transactions that each party can see.
- **Support for Ethereum:** Ethereum is a decentralized blockchain framework that establishes a peer-to-peer network that securely executes and verifies application code, called smart contracts. Smart contracts allow participants to conduct verified transactions without a trusted central authority. Transaction records are immutable, verifiable, and securely distributed across the network, giving participants full ownership and visibility into this data. Transactions are sent and received by Ethereum accounts that are created by users. A sender must sign transactions and spend Ether, Ethereum's native cryptocurrency, as a cost for processing transactions on the network.
- **Easy to scale:** After creating an Amazon Managed Blockchain network, you can easily invite other entities to join your network. After accepting the invitation and setting up a

membership, each new member of your blockchain network configures peer nodes that provide compute, storage, and memory to execute decentralized applications and maintain a copy of the ledger. If you need to scale an application, adding peer nodes can help process transactions more quickly. Managed Blockchain provides APIs that let you quickly create new nodes to meet the changing demands of your application. Also, Managed Blockchain provides a selection of instance families--bc.t3, bc.m5, and bc.c5--with varying combinations of CPU and memory so you can choose the appropriate mix of resources to support your workload.

- **Backed by AWS Key Management Service:** Amazon Managed Blockchain uses AWS Key Management Service (KMS) technology to secure Hyperledger Fabric's certificate authority, a component that manages user identities and issues enrolment certificates for securely communicating within the blockchain network. With Managed Blockchain, you don’t have to worry about setting up your own security device, such as a hardware security module (HSM), for this purpose.
- **Secure Interactions with VPC endpoints:** You can securely interact with your Hyperledger Fabric components managed by Amazon Managed Blockchain through Amazon VPC (Virtual Private Cloud) endpoints. Additionally, you can safely interact with blockchain peer nodes from other members in your network through this endpoint to endorse transactions.
- **Augmented ordering service with Amazon QLDB technology:** Hyperledger Fabric’s default ordering service can use Apache Kafka to support the communication of transactions across the network. While Kafka meets the needs of providing a messaging platform that can deliver transactions sequentially across the network, it is not optimized to store a complete history of transactional data, making it hard to recover historical transactions in case of a failure. Amazon Managed Blockchain's ordering service is built using Amazon QLDB technology, which has an immutable change log and maintains the complete history of all uncommitted transactions in the blockchain network, making the ordering service more durable.

59.1.2. Benefits

- **Fully managed:** With Amazon Managed Blockchain, you can quickly create blockchain networks that span multiple AWS accounts, enabling a group of members to execute transactions and share data without a central authority. Unlike self-hosting your blockchain infrastructure, Amazon Managed Blockchain eliminates the need for manually provisioning hardware, configuring software, and setting up networking and security components. With Managed Blockchain’s voting API, network participants can vote to add or remove members. Once a new member is added, Managed Blockchain lets that member launch and configure multiple blockchain peer nodes to process transaction requests and store a copy of the ledger. Managed Blockchain also monitors the network and automatically replaces poorly performing nodes.
- **Choice of Hyperledger Fabric or Ethereum:** Amazon Managed Blockchain supports two popular blockchain frameworks, Hyperledger Fabric and Ethereum. Hyperledger Fabric is well-suited for applications that require stringent privacy and permission controls with a known set of members, for example, a financial application where certain trade-related data is only shared with select banks. Ethereum is well suited for highly distributed blockchain networks where transparency of data for all members is important, for example, a customer loyalty blockchain network that allows any retailer in the network to independently verify a user's activity across all members to redeem benefits.

Alternatively, Ethereum can also be used for joining a public Ethereum blockchain network.

- **Scalable and Secure:** Amazon Managed Blockchain can easily scale your blockchain network as the usage of applications on the network grows over time. When a network member requires additional capacity for creating and validating transactions, the member can quickly add a new peer node using Managed Blockchain's APIs. Managed Blockchain provides a selection of instance types that comprise varying combinations of CPU and memory to give you the flexibility to choose the appropriate mix of resources for your workload. Additionally, Managed Blockchain secures your network's certificates with AWS Key Management Service (KMS) technology, eliminating the need for you to set up your own secure key storage.
- **Reliability:** Amazon Managed Blockchain improves the reliability of the "ordering service," a component in the Hyperledger Fabric framework that ensures delivery of transactions across the blockchain network. Hyperledger Fabric's default ordering service does not store a complete history of transactions, making it hard to keep track of and recover transaction history when needed. Managed Blockchain's ordering service is built using Amazon QLDB technology and has an immutable change log that accurately maintains the complete history of all transactions in the blockchain network, ensuring that you durably save this data.

59.2. Backup/Restore and Disaster Recovery

Managed Blockchain can replicate an immutable copy of your blockchain network activity into Amazon Quantum Ledger Database (QLDB), a fully managed ledger database. Please refer to https://docs.aws.amazon.com/managed-blockchain/latest/managementguide/what-is-managed-blockchain.html for more information.

59.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

59.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/managed-blockchain/
- **Service quotas:** https://docs.aws.amazon.com/general/latest/gr/managedblockchain.html
- **Service FAQs:** https://aws.amazon.com/managed-blockchain/faqs/

59.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/managed-blockchain/ and the following links for comprehensive technical documentation regarding this service.

- Hyperledger Fabric Developer Guide: Provides conceptual, instructional, and reference information for creating a network, joining a network, managing resources, and developing chaincode for Hyperledger Fabric on Managed Blockchain.
- Ethereum Developer Guide: Provides conceptual, instructional, and reference information for creating a node on an Ethereum network using Managed Blockchain.

60. Amazon Managed Service for Grafana

60.1. Service Overview

Amazon Managed Grafana is a fully managed service for open source Grafana developed in collaboration with Grafana Labs. Grafana is a popular open source analytics platform that enables you to query, visualize, alert on and understand your metrics no matter where they are stored.

With Amazon Managed Grafana, you can analyse your metrics, logs, and traces without having to provision servers, configure and update software, or do the heavy lifting involved in securing and scaling Grafana in production. You can create, explore, and share observability dashboards with your team, and spend less time managing your Grafana infrastructure and more time improving the health, performance, and availability of your applications. Connect Amazon Managed Grafana to multiple data sources in your observability stack, including AWS data sources like Amazon Managed Service for Prometheus, Amazon CloudWatch, and Amazon Elasticsearch Service, third-party ISVs like Datadog and Splunk, and self-managed data sources like InfluxDB. Amazon Managed Grafana natively integrates with AWS services so you can securely add, query, visualize, and analyse your AWS data across multiple accounts and Regions with a few clicks in the AWS Console. N.B. Grafana Enterprise is excluded from this offering.

60.1.1. Features

- **Visualize and correlate data across multiple data sources:** Amazon Managed Grafana connects to multiple data sources, enabling you to visualize, analyze, and correlate your metrics, logs, and traces in a unified dashboard. Amazon Managed Grafana securely and natively integrates with AWS services such as Amazon Managed Service for Prometheus, making it simple to query your AWS data across multiple accounts and multiple Regions in a single console. For example, you can create a dashboard that correlates container metrics from Amazon Managed Service for Prometheus, AWS services metrics from Amazon CloudWatch, and logs from Amazon Elasticsearch Service to monitor the health and performance of your applications running in containers. In the same console, you can layer and visualize data from self-managed data sources like Graphite, and third-party ISVs like Datadog and Splunk in the same dashboard.
- **Get started easily with pre-built panels and dashboards:** Amazon Managed Grafana makes it easy to construct the right queries and customize the display properties so that you can create the dashboard you need. With multiple pre-built dashboards for various data sources, you can instantly start visualizing and analyzing your application data without having to build dashboards from scratch.
- **Set up alerts to identify issues quickly:** By quickly identifying unintended changes in your system, you can minimize disruptions to your services. With Amazon Managed Grafana, you can configure alerts to identify problems in your system moments after they occur. You define the alert rule, how often it should be evaluated, the conditions that must be met for the alert to trigger, and how the alert notification should be delivered.
- **Share dashboards easily with user authentication and authorization:** With Amazon Managed Grafana, you can easily share interactive dashboards with specific users or across teams within your organization. With AWS SSO and SAML 2.0 integration with Identity Providers, you can leverage your existing corporate directory services to grant

user access and authentication to your Grafana workspaces. You can assign user Read/Write or Read-Only roles by giving them Administrator, Editor, or Viewer privileges. You can also create Teams to restrict dashboard and data source access to the right users. Amazon Managed Grafana integrates with popular corporate directory services including Microsoft Active Directory, Azure Active Directory, Okta, Ping Identity, OneLogin, and CyberArk. With the Grafana Team Sync feature, Amazon Managed Grafana keeps track of all synchronized users in teams giving you flexibility to combine group memberships from your directory services with Grafana teams.

- **Troubleshoot and collaborate with your team:** You can create multiple Grafana Teams to easily grant data source access permissions and share dashboards to groups of users. New team members added later will also inherit access permissions to shared resources without having to manually grant permissions one dashboard at a time. Users can view and edit dashboards in real time, track dashboard version changes, and easily share dashboards with other users in the same Team so that everyone is viewing the same data while troubleshooting operational issues. Users can also easily share dashboards with other teams or external entities by creating dashboard snapshots that can be publicly accessed.
- **Security and authentication:** Amazon Managed Grafana tightly integrates with multiple AWS services to meet your corporate security and compliance requirements. Access to Amazon Managed Grafana is authenticated through AWS SSO or your existing Identity Provider via SAML 2.0, enabling re-use of existing trust relationships between AWS and your corporate user directories. You can track changes made to Grafana workspaces for compliance and audit tracking using audit logs provided by AWS CloudTrail. Amazon Managed Grafana also natively integrates with multiple AWS data sources including Amazon Elasticsearch Service, Amazon CloudWatch, AWS X-Ray, AWS IoT SiteWise, Amazon Timestream, and Amazon Managed Service for Prometheus, so you don't have to manually manage IAM credentials and permissions for each data source. Amazon Managed Grafana also discovers the resources in your account across multiple Regions and across your Organizational Units, and automatically provisions the right IAM policies to access your data.
- **No servers to manage:** With a few clicks in the Amazon Managed Grafana console, you can instantly create one or many workspaces to visualize and analyze your metrics, logs, and traces without having to build, package, or deploy any hardware or infrastructure. Amazon Managed Grafana automatically provisions, configures, and manages the operations of your Grafana workspaces, with automatic version upgrades to ensure that your Grafana workspaces are always up-to-date with the latest features. The service auto scales to meet your dynamic usage demands.
- **Automatic recovery and patching:** Amazon Managed Grafana workspaces are highly available with multi-AZ replication. Amazon Managed Grafana also continuously monitors the health of your Grafana workspaces and replaces unhealthy nodes, without impacting your access to Grafana workspaces. Amazon Managed Grafana manages the availability of your compute and database nodes so that you don't have to start, stop, or reboot any infrastructure resources.
- **Encryption and security:** Amazon Managed Grafana encrypts your data at rest without special configuration, third-party tools, or additional cost. Amazon Managed Grafana also encrypts data in-transit via TLS.

60.1.2. Benefits

- **Enjoy the power of Grafana at scale:** Amazon Managed Grafana builds, packages, and deploys workspaces for you, and manages their provisioning, setup, scaling, and maintenance, so you don't have to. You can then create Grafana dashboards and visualizations in each workspace to analyze your metrics, logs, and traces. With Amazon Managed Grafana, you create Grafana workspaces where you can define user access and policy controls for data sources that you specify.
- **Visualize, analyze, and correlate securely across multiple data sources:** Amazon Managed Grafana natively integrates with AWS data sources that collect operational data, discovering the resources in your AWS account or across your Organizational Units, and automatically provisions the right AWS Identity and Access Management (IAM) policies to access your data. AWS data sources include Amazon CloudWatch, Amazon Elasticsearch Service, AWS X-Ray, AWS IoT SiteWise, Amazon Timestream, and Amazon Managed Service for Prometheus. You can query these data sources across multiple AWS accounts and Regions. Amazon Managed Grafana also supports popular third-party data sources such as Graphite, InfluxDB, and more.
- **Secure access to data and dashboards:** Amazon Managed Grafana integrates with multiple AWS security services and supports AWS Single Sign-On as well as Security Assertion Markup Language (SAML) 2.0 to meet your corporate security and compliance requirements. When setting up a workspace in Amazon Managed Grafana, you can grant users in your corporate directory access to specific dashboards and data sources, and control their read/write access without having to manage multiple user identity pools. You can track changes made to workspaces for compliance and audit logging, and third-party auditors can assess security and compliance as part of multiple AWS compliance programs, such as AWS CloudTrail logs.
- **Migrate from self-managed Grafana, easily:** No need to start from scratch when migrating from your existing Grafana environment. Use Amazon Managed Grafana APIs to easily create new Grafana workspaces within minutes. And you can use Grafana APIs to programmatically create, delete, or import your Grafana dashboards – you don't have to re-create dashboard definitions that you've already built and perfected.
- **Get started quickly with pre-built dashboards:** After selecting your data source, you can choose from a variety of pre-built visualizations to quickly start analyzing metrics, logs, and traces without having to build a dashboard from scratch. Amazon Managed Grafana offers a variety of visualizations across a broad set of data sources.
- **Access third-party Enterprise data source plugins:** With one click, you can upgrade to Grafana Enterprise license and gain access to a variety of their third-party Enterprise data source plugins, including AppDynamics, Atlassian Jira, Datadog, Dynatrace, Gitlab, Honeycomb, MongoDB, NewRelic, Oracle Database, Salesforce, SAP HANA, ServiceNow, VMware Tanzu Observability by Wavefront, and Snowflake.

60.2. Backup/Restore and Disaster Recovery

This requirement is not applicable for this service. For additional information beyond what is described herein, please refer to http://aws.amazon.com/documentation/.

60.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

60.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/grafana/index.html
- **Service quotas:** https://docs.aws.amazon.com/grafana/latest/userguide/AMG_quotas.html
- **Service FAQs:** https://aws.amazon.com/grafana/faqs/?nc=sn&loc=5&refid=ps_a134p000006gb2oaau&trkcampaign=acq_paid_search_brand

60.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/grafana/index.html and the following links for comprehensive technical documentation regarding this service.

- Amazon Managed Grafana User Guide: Provides a conceptual overview of Amazon Managed Grafana and includes detailed instructions for using the various features.
- Amazon Managed Grafana API Reference: Describes all the API operations for Amazon Managed Grafana in detail. Also provides sample requests, responses, and errors for the supported web services protocols.
- Amazon Managed Grafana section of the AWS CLI Reference: Describes the commands in the AWS CLI that you can use to manage Amazon Managed Grafana. Provides the syntax and options for each command.

61. Amazon Managed Service for Prometheus

61.1. Service Overview

Amazon Managed Service for Prometheus is a Prometheus-compatible monitoring and alerting service that makes it easy to monitor containerized applications and infrastructure at scale. The Cloud Native Computing Foundation's Prometheus project is a popular open source monitoring and alerting solution optimized for container environments. With Amazon Managed Service for Prometheus, you can use the open source Prometheus query language (PromQL) to monitor and alert on the performance of containerized workloads, without having to scale and operate the underlying infrastructure. Amazon Managed Service for Prometheus automatically scales the ingestion, storage, alerting, and querying of operational metrics as workloads grow or shrink, and is integrated with AWS security services to enable fast and secure access to data. The service is integrated with Amazon Elastic Kubernetes Service (Amazon EKS), Amazon Elastic Container Service (Amazon ECS), and AWS Distro for OpenTelemetry.

61.1.1. Features

- **Easy to deploy and manage**
  - Setup and configuration: Getting started with Amazon Managed Service for Prometheus is easy. You can create an Amazon Managed Service for Prometheus workspace, which is a Prometheus instance, with a few clicks in the AWS console. Each Amazon Managed Service for Prometheus workspace is automatically deployed across multiple Availability Zones, and is immediately ready to ingest and query metrics. You can quickly enable metric collection in multiple ways. You can configure AWS Distro for OpenTelemetry to collect metrics from a Prometheus-instrumented application, and send the metrics to Amazon Managed Service for

Prometheus. You can also ingest metrics from Prometheus servers in your Amazon EKS clusters, and in self-managed Kubernetes clusters running on Amazon EC2.

  - No servers to manage:With a few clicks in the Amazon Managed Service for Prometheus console, you can instantly create one or many workspaces to monitor the performance of containerized workloads without having to build, package, or deploy any hardware or infrastructure. Amazon Managed Service for Prometheus automatically scales the ingestion, storage, and querying of operational metrics as workloads grow or shrink, and is integrated with AWS security services to enable fast and secure access to data.

- **Cost-effective:** Pay only for what you use: With Amazon Managed Service for Prometheus, there are no upfront fees or commitments. You pay only for what you use based on metrics ingested, stored, and queried.
- **Highly secure, scalable, and available**
  - Security: Amazon Managed Service for Prometheus offers enterprise-ready security and compliance. Amazon Managed Service for Prometheus includes built-in support for AWS Identity and Access Management (IAM), and fine-grained access control for ingesting and exporting metrics from AWS services. Amazon Managed Service for Prometheus also integrates with AWS CloudTrail, so you can get a record of actions taken by a user, a role, or an AWS service in Amazon Managed Service for Prometheus. CloudTrail captures all API calls for Amazon Managed Service for Prometheus as events, which you can set up to be continuously delivered to an Amazon S3 bucket. If you are using Amazon Managed Service for Prometheus and Amazon Managed Grafana together, they seamlessly and securely connect using IAM authentication and private VPC endpoint connectivity. With AWS PrivateLink, you can connect your VPCs to Amazon Managed Service for Prometheus and other services in AWS in a secure and scalable manner. Network traffic that uses AWS PrivateLink doesn't traverse the public internet, reducing the exposure to threat vectors such as brute force and distributed denial-of-service attacks. Amazon Managed Service for Prometheus supports the latest API versions and will be automatically updated with the latest Prometheus feature set and patched to address any critical security vulnerabilities.
  - Scalability: Amazon Managed Service for Prometheus is specifically architected to handle the high cardinality monitoring data with a large volume of tags and dimensions that is generated by container-based applications. Amazon Managed Service for Prometheus manages the operational complexity of elastically scaling the ingestion, storage, and querying of metrics.
  - Availability: Amazon Managed Service for Prometheus is highly available and deployed in multiple AWS Regions and across Availability Zones, which are connected with low-latency, high-throughput, and highly redundant networking. With Availability Zones, you can design and operate applications and databases that automatically fail over between zones without interruption. Availability Zones are more highly available, fault tolerant, and scalable than traditional single or multiple data center infrastructures. AWS delivers the highest network availability of any cloud provider.
- **Ingest and Collect**: Amazon Managed Service for Prometheus includes a remote write-compatible API that can ingest metrics from OpenTelemetry, Prometheus libraries, and existing Prometheus servers. Metrics can be ingested from any clusters running on AWS and hybrid environments, with on-demand scaling to meet your growing needs. Existing

metric collectors such as the OpenTelemetry collector and the Prometheus server can be used to securely remote write Prometheus metrics to Amazon Managed Service for Prometheus from over 150+ third party Exporters such as Java/JMX, Apache Kafka, and Redis. A full list of Prometheus third party exporters can be found in the Prometheus documentation. The Prometheus server is one of many components of the Prometheus open source monitoring and alerting solution. The server can be used for service discovery of third-party Exporters to scrape and ingest millions of samples per second. The Prometheus server is often deployed in Kubernetes clusters to collect metrics on infrastructure and containerized applications. However, to manage the operational complexity of elastically scaling the ingestion, storage, and querying of metrics, Amazon Managed Service for Prometheus automatically adjusts as your container workloads scale up and down to deliver cost-effective performance metrics and consistent query response times. You can continue to use your Prometheus server to collect metrics, and securely remote write them to Amazon Managed Service for Prometheus. To learn more about writing your own client to use remote write, see Building a Prometheus remote write exporter for the OpenTelemetry Go SDK. AWS Distro for OpenTelemetry is an enterprise-ready AWS supported distribution of OpenTelemetry that makes it easy to collect and send application metrics and traces to multiple AWS and third-party monitoring services, without having to manually instrument your application for each solution. The Collector securely integrates with AWS Identity and Access Management (IAM) for authentication and fine-grained access control to securely remote write these metrics to Amazon Managed Service for Prometheus.

- **Monitor and Alert**: Amazon Managed Service for Prometheus includes a query-compatible HTTP API that allows you to query metrics, metric labels, metric metadata, and time series metrics. Tools such as Grafana, an open source interactive visualization tool for time series data, are commonly used to query and visualize metrics from Prometheus. The Grafana Prometheus data source plugin can easily be configured to query metrics from Amazon Managed Service for Prometheus. You can also use Amazon Managed Grafana, a fully managed AWS service that makes it easy to use Grafana to monitor operational data with interactive data visualizations in a single console across multiple data sources, without needing to deploy, manage, and operate Grafana servers. Amazon Managed Service for Prometheus also supports Prometheus alerting and recording rules that can be imported from your existing Prometheus server. Recording rules allow you to precompute frequently needed or computationally expensive PromQL queries, and save the results as new time series metrics. Alerting rules allow you to define alert conditions using PromQL, and send notifications to Amazon Simple Notification Service (SNS). Alert management features such as inhibition, grouping, and routing are also compatible with the Prometheus solution, so you can import existing Prometheus alert configurations using the Amazon Managed Service for Prometheus APIs. Once imported, PromQL queries defined in the alerts will be continuously evaluated against your Prometheus workspace, and can be integrated with SNS for notification. An Amazon Managed Service for Prometheus workspace is a logical and isolated Prometheus server dedicated to Prometheus resources such as metrics, recording rules, and alerting rules, where you ingest, store, and query your Prometheus metrics.
- **Analyze**: Prometheus provides a flexible query language called PromQL (Prometheus Query Language) to filter, aggregate, and alarm on metrics and quickly gain performance visibility without any code changes. The result of an expression can be

consumed by external systems via the HTTP API and visualization tools such as Grafana, using the Prometheus data source plugin. This allows you to do simple time series selection, subqueries, functions, and operators – dramatically improving the troubleshooting experience and reducing MTTD (mean time to detection).

- **Enterprise-Ready**: Amazon Managed Service for Prometheus integrates with AWS security services to meet your compliance and security needs. AWS account users can control user access and permissions to individual Amazon Managed Service for Prometheus workspaces using AWS Identity and Access Management (IAM). All queries sent to the service are authorized by IAM. Amazon Managed Grafana seamlessly and securely connects to Amazon Managed Service for Prometheus using IAM authentication and private VPC endpoint connectivity. Support for AWS PrivateLink secures access to our APIs without accessing the public internet to manage your workspaces and the ingestion and querying of Prometheus metrics. Amazon Managed Service for Prometheus will always support the latest API versions and will be automatically updated with the latest Prometheus feature set and patched to address any critical security vulnerabilities. You can enable AWS CloudTrail integration to have full audit visibility into changes such as workspaces created, deleted, and updated, and users accessing their Amazon Managed Service for Prometheus workspaces. Amazon Managed Service for Prometheus is a fully managed Prometheus-compatible service that will maintain always up-to-date HTTP API compatibility. The service allows you to self-manage (create, describe, list, and delete) Prometheus workspaces in all supported AWS Regions using the AWS Console, AWS CLI, and AWS SDKs.

61.1.2. Benefits

- **The Prometheus you already know:** Use the familiar, flexible Prometheus query language (PromQL) to filter, aggregate, and alarm on metrics, and quickly gain performance visibility for large volumes of metrics labels. Amazon Managed Service for Prometheus supports all metric types: gauge, counter, summary, and histogram.
- **Prometheus at scale. Easily:** Amazon Managed Service for Prometheus automatically scales as your ingestion and query needs grow, handling millions of unique time series metrics from large container deployments while maintaining consistent query response times. Amazon Managed Service for Prometheus offers multi-AZ replication within an AWS Region.
- **Prometheus with AWS-level security:** Amazon Managed Service for Prometheus integrates with AWS Identity and Access Management for authentication and fine-grained permissions for users and groups. AWS PrivateLink provides easy and secure access to services hosted on AWS, keeping your network traffic within the AWS network. AWS Organizations integration allows for policy control, and API calls are logged to AWS CloudTrail.
- **Monitor and alert on containers running on AWS and on-premises:** You can quickly set up Amazon Managed Service for Prometheus to collect and query metrics from AWS container services including Amazon EKS, Amazon ECS, and AWS Fargate. Amazon Managed Service for Prometheus also includes APIs that enable you to securely ingest, alert on, and query metrics from all your self-managed Kubernetes clusters, on AWS and on-premises.
- **Reduce operational costs:** Using Amazon Managed Service for Prometheus , you can remove the undifferentiated heavy lifting of running open source Prometheus at scale.

This helps significantly reduce the operational costs of configuring, upgrading, and scaling standalone Prometheus servers.

- **Maximize impact with AWS observability services:** Amazon Managed Service for Prometheus also works with AWS Distro for OpenTelemetry as a collection agent for Prometheus metrics, and with Amazon Managed Grafana to create rich, powerful data visualizations.

61.2. Backup/Restore and Disaster Recovery

This requirement is not applicable for this service. For additional information beyond what is described herein, please refer to http://aws.amazon.com/documentation/.

61.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

61.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/prometheus/index.html
- **Service quotas:** https://docs.aws.amazon.com/prometheus/latest/userguide/AMP_quotas.html
- **Service FAQs:** https://aws.amazon.com/prometheus/faqs/

61.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/prometheus/index.html and the following links for comprehensive technical documentation regarding this service.

- Amazon Managed Service for Prometheus User Guide: Provides a conceptual overview of Amazon Managed Service for Prometheus and includes detailed instructions for using the various features.
- Amazon Managed Service for Prometheus API Reference: Describes the API operations for Amazon Managed Service for Prometheus in detail. Also provides sample requests, responses, and errors for the supported web services protocols.

62. Amazon Managed Streaming for Apache Kafka

62.1. Service Overview

Amazon Managed Streaming for Apache Kafka (MSK) offers fully managed Apache Kafka. This means Amazon MSK provisions your servers, configures your Apache Kafka clusters, replaces servers when they fail, orchestrates server patches and upgrades, architects clusters for high availability, ensures data is durably stored and secured, sets up monitoring and alarms, and runs scaling to support load changes. With a managed service, you can spend your time developing and running streaming event applications.

Amazon MSK provides open-source, highly secure Apache Kafka clusters distributed across multiple Availability Zones (AZs), giving you resilient, highly available streaming storage. Amazon MSK is highly configurable, observable, and scalable, allowing for the flexibility and control needed for various use cases.

Application development is simpler with Amazon MSK because of tight integrations with other AWS services. Amazon MSK integrates with AWS Identity and Access Management (IAM) and AWS Certificate Manager for security, AWS Glue Schema Registry for schema governance, Amazon Kinesis Data Analytics and AWS Lambda for stream processing, and more. Amazon MSK provides the integration backbone for modern messaging and event-driven applications at the centre of data ingest and processing services, as well as microservice application architectures.

62.1.1. Features

- **Fully managed:** With a few clicks in the console, you can create a fully managed Apache Kafka cluster that follows Apache Kafka's deployment best practices, or create your own cluster using a custom configuration. Once you create your desired configuration, Amazon MSK automatically provisions, configures, and manages your Apache Kafka cluster operations and Apache ZooKeeper nodes.
- **Apache ZooKeeper included:** Apache ZooKeeper is required to run Apache Kafka, coordinate cluster tasks, and maintain state for resources interacting with the cluster. Amazon MSK manages the Apache ZooKeeper nodes for you. Each Amazon MSK cluster includes the appropriate number of Apache ZooKeeper nodes for your Apache Kafka cluster at no additional cost.
- **High availability is default:** All clusters are distributed across multiple AZs (three is the default), are supported by Amazon MSK's service-level agreement (not applicable to MSK Serverless while in preview), and are supported by automated systems that detect and respond to issues within cluster infrastructure and Apache Kafka software. If a component fails, Amazon MSK automatically replaces it without downtime to your applications. Amazon MSK manages the availability of your Apache ZooKeeper nodes so you don't need to start, stop, or directly access the nodes yourself. It also automatically deploys software patches as needed to keep your cluster up to date and running smoothly.
- **Data replication:** Amazon MSK uses multi-AZ replication for high availability. Data replication is included at no additional cost.
- **Private connectivity:** Your Apache Kafka clusters run in an Amazon Virtual Private Cloud (VPC) managed by Amazon MSK. Your clusters are available to your own Amazon VPCs, subnets, and security groups based on the configuration you specify. You have complete control of your network configuration and IP addresses.
- **Granular access control:** IAM Access Control is a no-cost security option that simplifies cluster authentication and Apache Kafka API authorization using IAM roles or user policies to control access. Using IAM Access Control, you no longer need to build and run one-off access management systems to control client authentication and authorization for Apache Kafka. Your clusters are secured using least-privileged permissions by default. For provisioned clusters, you also can use Simple Authentication and Security Layer (SASL)/Salted Challenge Response Authentication Mechanism (SCRAM) or mutual Transport Layer Security (TLS) authentication with Apache Kafka access control lists (ACLs) to control client access.
- **Encryption at rest and in transit:** Amazon MSK encrypts your data at rest without special configuration or third-party tools. For provisioned clusters, all data at rest can be encrypted using AWS Key Management Service (KMS) key by default or your own key. You can also encrypt data in transit via TLS between brokers and between clients and

brokers on your cluster. For serverless clusters, all data at rest is encrypted by default using service-managed keys, and all data in transit is encrypted by default via TLS.

- **Deeply integrated:** No other provider offers the breadth and depth of AWS integrations in Amazon MSK. These integrations include:
    - AWS IAM for Apache Kafka and service-level API access control
    - Amazon Kinesis Data Analytics for running fully managed Apache Flink applications to process streaming data within Apache Kafka
    - Amazon Kinesis Data Analytics Studio to run interactive Streaming SQL and long-running SQL jobs using Apache FlinkSQL
    - AWS Glue Schema Registry to centrally control and evolve schemas
    - AWS IoT Core for IoT event streaming into MSK
    - AWS Database Migration Service (AWS DMS) for change data capture and analytics
    - Amazon Virtual Private Cloud (Amazon VPC) for private client connectivity and network isolation
    - AWS Key Management Service (AWS KMS) for at-rest encryption
    - AWS Certificate Manager Private Certificate Authority for mutual TLS client authentication
    - AWS Secrets Manager for secure storage and management of SASL/SCRAM secrets
    - AWS CloudFormation to deploy Amazon MSK in code
    - Amazon CloudWatch for cluster-, broker-, topic-, consumer-, and partition-level metrics
- **Run with native Apache Kafka:** Amazon MSK deploys native versions of Apache Kafka so applications and tools built for Apache Kafka just work with Amazon MSK out of the box, with no application code changes.
- **Streamlined version availability:** Amazon MSK typically makes newer versions of Apache Kafka available within seven days of public availability.
- **Seamless version upgrades:** You can upgrade Apache Kafka versions on provisioned clusters in just a few clicks, allowing you to decide when to take advantage of features and bug fixes present in new Apache Kafka versions. Amazon MSK automates the deployment of version upgrades on running clusters to maintain client I/O availability for customers following best practices. For serverless clusters, Apache Kafka versions are upgraded automatically by Amazon MSK.
- **Lowest cost:** Amazon MSK lets you get started for less than $2.50 per day. Customers typically pay between $0.05 and $0.07 per GB ingested, all-in, which can be as low as 1/13th the cost of other managed providers. Visit the Amazon MSK Pricing page to learn more about pricing.
- **Broker scaling (provisioned clusters only):** You can scale your Amazon MSK clusters by changing the size or family of your Apache Kafka brokers in minutes with no downtime. Changing the size or family is a popular way to scale Amazon MSK clusters

because it gives you the flexibility to adjust cluster compute capacity for changes in your workloads. This method can be preferred because it doesn't require partition reassignment, which can impact Apache Kafka availability.

- **Cluster scaling (serverless clusters only):** Amazon MSK automatically scales compute and storage resources of your clusters in response to your application's throughput needs.
- **Automatic partition management:** Amazon MSK integrates with Cruise Control, a popular open-source tool for Apache Kafka that automatically manages partition assignment on your behalf. For serverless clusters, Amazon MSK automatically manages partition assignments for you.
- **Automatic storage scaling (provisioned clusters only):** You can seamlessly scale up the amount of storage provisioned per broker to match storage requirement changes using the AWS Management Console or AWS Command Line Interface (AWS CLI). You can also create an auto scaling policy to automatically expand your storage to meet growing streaming requirements.
- **Configurable:** Amazon MSK deploys a best practice cluster configuration for Apache Kafka by default. For provisioned clusters, you have the ability to tune more than 30 different cluster configurations while supporting all dynamic and topic-level configurations. For more information, see Custom MSK Configurations in the documentation.
- **Easy observability of streaming performance with CloudWatch metrics by default:** You can visualize and monitor important metrics using Amazon CloudWatch to understand and maintain streaming application performance.
- **Export JMX and Node metrics to a Prometheus server with Open Monitoring (provisioned clusters only):** Open Monitoring with Prometheus lets you monitor Amazon MSK using solutions such as Datadog, Lenses, New Relic, Sumo Logic, or a Prometheus server, and easily migrate your existing monitoring dashboards to Amazon MSK. For more information, see Open Monitoring with Prometheus in the documentation.

62.1.2. Benefits

- **Fully managed:** Eliminate operational overhead, including the provisioning, configuration, and maintenance of highly available Apache Kafka and Kafka Connect clusters.
- **Flexible:** Use applications and tools built for Apache Kafka out of the box (no code changes required), and scale cluster capacity automatically.
- **Easy to use:** Easily deploy secure, compliant, and production-ready applications using native AWS integrations.
- **Low cost:** Keep costs low with Amazon MSK. With pay-as-you-go pricing, it is offered as low as 1/13 the cost of other providers.

62.2. Backup/Restore and Disaster Recovery

This requirement is not applicable for this service. For additional information beyond what is described herein, please refer to http://aws.amazon.com/documentation/.

62.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

62.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/msk/
- **Service quotas:** https://docs.aws.amazon.com/msk/latest/developerguide/limits.html
- **Service FAQs:** https://aws.amazon.com/msk/faqs/

62.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/msk/ and the following links for comprehensive technical documentation regarding this service.

- Developer Guide: Provides a conceptual overview of Amazon MSK and includes detailed instructions for using the service.

63. Amazon Managed Workflows for Apache Airflow

63.1. Service Overview

Amazon Managed Workflows for Apache Airflow (MWAA) is a managed orchestration service for Apache Airflow that makes it easier to set up and operate end-to-end data pipelines in the cloud at scale. Apache Airflow is an open-source tool used to programmatically author, schedule, and monitor sequences of processes and tasks referred to as “workflows.” With Managed Workflows, you can use Airflow and Python to create workflows without having to manage the underlying infrastructure for scalability, availability, and security. Managed Workflows automatically scales its workflow execution capacity to meet your needs, and is integrated with AWS security services to help provide you with fast and secure access to data.

63.1.1. Features

- **Easy Airflow deployment:** Managed Workflows leverage the same open source Apache Airflow product you know, just made easier. You can deploy Managed Workflows from AWS Management Console, CLI, AWS CloudFormation, or AWS SDK - and leverage the same Airflow user experience you’re familiar with.
- **Automatic scaling:** With Managed Workflows, there’s seamless worker scaling with no configuration required. Worker monitoring is built in - when workers are over-burdened, additional workers are provisioned automatically, and then decommissioned when no longer needed.
- **Built-in security:** Managed Workflows keep your data secure using Amazon’s Virtual Private Cloud (VPC), and data is automatically encrypted using AWS Key Management Service (KMS), so your workflow environment is secure by default.
- **Workflow monitoring in AWS or on-premises:** Managed Workflows automatically sends Apache Airflow system metrics and logs to Amazon Cloudwatch, making it easier for you to view task execution delays and workflow errors across one or more environments without third party tools.

- **Low operational costs:** Managed Workflows remove the operational load of running open source Apache Airflow at scale so you can reduce operational costs and engineering overhead while running a data pipeline orchestration at nearly any scale.
- **Plug-in integration:** Managed Workflows connect to the AWS resources required for your workflows including Athena, Batch, Cloudwatch, DynamoDB, DataSync, EMR, ECS/Fargate, EKS, Firehose, Glue, Lambda, Redshift, SQS, SNS, Sagemaker, and S3. You can use Managed Workflows to connect to your own on-premises resources.

63.1.2. Benefits

- **Deploy Airflow rapidly at scale:** Get started in minutes from the AWS Management Console, CLI, AWS CloudFormation, or AWS SDK. Create an account and begin deploying Directed Acyclic Graphs (DAGs) to your Airflow environment immediately without reliance on development resources or provisioning infrastructure.
- **Run Airflow with built-in security:** With Managed Workflows, your data is secure by default as workloads run in your own isolated and secure cloud environment using Amazon's Virtual Private Cloud (VPC), and data is automatically encrypted using AWS Key Management Service (KMS). You can control role-based authentication and authorization for Apache Airflow's user interface via AWS Identity and Access Management (IAM), providing users Single Sign-ON (SSO) access for scheduling and viewing workflow executions.
- **Reduce operational costs:** Managed Workflows is a managed service, removing the heavy lifting of running open source Apache Airflow at scale. With Managed Workflows, you can reduce operational costs and engineering overhead while meeting the on-demand monitoring needs of end to end data pipeline orchestration.
- **Use a pre-existing plugin or use your own:** Connect to any AWS or on-premises resources required for your workflows including Athena, Batch, Cloudwatch, DynamoDB, DataSync, EMR, ECS/Fargate, EKS, Firehose, Glue, Lambda, Redshift, SQS, SNS, Sagemaker, and S3.

63.2. Backup/Restore and Disaster Recovery

This requirement is not applicable for this service. For additional information beyond what is described herein, please refer to http://aws.amazon.com/documentation/.

63.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

63.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/mwaa/
- **Service quotas:** https://docs.aws.amazon.com/mwaa/latest/userguide/mwaa-quotas.html
- **Service FAQs:** https://aws.amazon.com/managed-workflows-for-apache-airflow/faqs/

63.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/mwaa/ and the following links for comprehensive technical documentation regarding this service.

- User Guide: Describes how to build and manage an Apache Airflow pipeline on an Amazon Managed Workflows for Apache Airflow (Amazon MWAA) environment.

64. Amazon Monitron

64.1. Service Overview

Amazon Monitron gateway device securely transfers data to AWS, and the service automatically analyzes your data for abnormal machine conditions using machine learning. You can start monitoring equipment health in minutes through the Amazon Monitron mobile and web apps, and enable predictive maintenance with the same technology used to monitor equipment in Amazon Fulfillment Centers.

64.1.1. Features

- **Simple device set-up with the Amazon Monitron Mobile App:** Set up Amazon Monitron Sensors quickly and easily by tapping your phone on the sensor utilizing near-field communication (NFC) technology. Set up your Gateway by following a few simple steps in the app. You can quickly install and start using these devices to monitor your equipment without any development work.
- **ISO and ML-based analytics:** Amazon Monitron automatically detects abnormal machine operating states by analyzing vibration and temperature signals using the ISO 20816 standards for vibration, and ML-enabled models.
- **Timely notifications in the Amazon Monitron app:** Amazon Monitron sends push notifications when the service detects abnormal machine patterns from the vibration and temperature settings. You can also review and track these abnormal machine states within the app.
- **Alert feedback:** With just a few taps in the mobile and web apps, you can enter feedback on the alerts received, such as failure mode, failure cause, and action taken. Amazon Monitron learns from that feedback and continually improves over time.

64.1.2. Benefits

- **Machine detection**: Detect machine issues before they occur with machine learning (ML), and take action.
- **Monitoring**: Start monitoring equipment in minutes with easy installation and automatic, secure analysis through the Amazon Monitron end-to-end system.
- **Improve system accuracy**: Improve system accuracy continuously as Amazon Monitron learns from technician feedback entered in the mobile and web apps.

64.2. Backup/Restore and Disaster Recovery

This requirement is not applicable for this service. For additional information beyond what is described herein, please refer to http://aws.amazon.com/documentation/.

64.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

64.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/Monitron/index.html
- **Service quotas:** https://docs.aws.amazon.com/Monitron/latest/admin-guide/quotas.html
- **Service FAQs:** https://aws.amazon.com/monitron/faqs/?refid=ps_a134p000006gb2oaau&trkcampaign=acq_paid_search_brand

64.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/Monitron/index.html and the following links for comprehensive technical documentation regarding this service.

- Getting Started Guide: Describes how IT Managers create projects, how admins create assets (machines) and set up sensors, and how technicians start monitoring equipment.
- IT Manager's Guide: Describes how the IT Manager uses the Amazon Monitron console to set up projects, admins, and users. Also explains how to add admin users and use AWS Single Sign-On to create a user directory and connect to it.
- User Guide: Describes how admins and technicians use the Amazon Monitron mobile and web apps to monitor equipment for potential failures.

65. Amazon MQ

65.1. Service Overview

Amazon MQ is a managed message broker service for Apache ActiveMQ and RabbitMQ that makes it easy to set up and operate message brokers on AWS. Amazon MQ reduces your operational responsibilities by managing the provisioning, setup, and maintenance of message brokers for you. Because Amazon MQ connects to your current applications with industry-standard APIs and protocols, you can easily migrate to AWS without having to rewrite code.

65.1.1. Features

- **Managed Service:** With Amazon MQ, you can use the AWS Management Console, AWS CloudFormation, the Command Line Interface (CLI), or simple API calls to launch a production-ready message broker in minutes. Amazon MQ manages administrative tasks such as hardware provisioning, broker setup, software upgrades, and failure detection and recovery.
- **Security:** Amazon MQ provides encryption of your messages at rest and in transit. It's easy to ensure that your messages are securely stored in an encrypted format. Connections to the broker use SSL, and access can be restricted to a private endpoint within your Amazon VPC, which allows you to isolate your broker in your own virtual network.

  Amazon MQ is integrated with AWS Identity and Access Management (IAM) and provides you the ability to control the actions that your IAM users and groups can take on specific Amazon MQ brokers. Authentication from applications to the broker itself is provided using username and password-based authentication, and optionally using LDAP (Lightweight Directory Access Protocol) for ActiveMQ brokers.

- **Monitoring:** Amazon MQ is integrated with Amazon CloudWatch and AWS CloudTrail. With CloudWatch you can monitor metrics on your brokers, queues, and topics. For example, you can monitor the depth of your queues and generate alarms if messages aren't getting through. Using CloudTrail, you can log, continuously monitor, and retain Amazon MQ API calls.
- **Broker Instance Types:** Amazon MQ currently supports seven broker instance types: mq.t2.micro, mq.t3.micro, mq.m4.large, mq.m5.large, mq.m5.xlarge, mq.m5.2xlarge, and mq.m5.4xlarge, which provide varying combinations of CPU, memory, and network performance. The mq.t3.micro instances are designed for initial product evaluation and the mq.m5.large instance for default production usage. Amazon MQ also supports both single-instance brokers, suitable for evaluation and testing, and replicated highly available deployment modes recommended for production.
- **Pay-as-you-go Pricing:** Amazon MQ provides cost-efficient and flexible capacity, and there is no minimum fee. You pay for the number of hours your broker instance runs and the storage you use monthly. It's easy and inexpensive to create new brokers for additional capacity. For more details see Amazon MQ Pricing.

65.1.2. Benefits

- **Migrate quickly:** Connecting your current applications to Amazon MQ is easy because it uses industry-standard APIs and protocols for messaging, including JMS, NMS, AMQP 1.0 and 0-9-1, STOMP, MQTT, and WebSocket. This enables you to move from any message broker that uses these standards to Amazon MQ by simply updating the endpoints of your applications to connect to Amazon MQ.
- **Offload operational responsibilities:** Amazon MQ manages the administration and maintenance of message brokers and automatically provisions infrastructure for high availability. There is no need to provision hardware or install and maintain software and Amazon MQ automatically manages tasks such as software upgrades, security updates, and failure detection and recovery.
- **Durable messaging made easy:** Amazon MQ is automatically provisioned for high availability and message durability when you connect your message brokers. Amazon MQ stores messages redundantly across multiple Availability Zones (AZ) within an AWS region and will continue to be available if a component or AZ fails.

65.2. Backup/Restore and Disaster Recovery

This requirement is not applicable for this service. For additional information beyond what is described herein, please refer to http://aws.amazon.com/documentation/.

65.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

65.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/amazon-mq/
- **Service quotas:** https://docs.aws.amazon.com/amazon-mq/latest/developer-guide/amazon-mq-limits.html

- **Service FAQs:** https://aws.amazon.com/amazon-mq/faqs/

65.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/amazon-mq/ and the following links for comprehensive technical documentation regarding this service.

- Developer Guide: Provides a conceptual overview of Amazon MQ and includes detailed instructions for creating and managing brokers and migrating from on-premises brokers.
- Migration Guide: Describes migrating commercial, on-premises message brokers to Amazon MQ.

66. Amazon Neptune

66.1. Service Overview

With Amazon Neptune, you can create sophisticated, interactive graph applications that can query billions of relationships in milliseconds. SQL queries for highly connected data are complex and hard to tune for performance. Instead, Amazon Neptune allows you to use the popular graph query languages Apache TinkerPop Gremlin and W3C's SPARQL to execute powerful queries that are easy to write and perform well on connected data. This significantly reduces code complexity, and allows you to more quickly create applications that process relationships.

66.1.1. Features

- **High Throughput, Low Latency for Graph Queries:** Amazon Neptune is a purpose-built, high-performance graph database engine. Neptune efficiently stores and navigates graph data, and uses a scale-up, in-memory optimized architecture to allow for fast query evaluation over large graphs. With Neptune, you can use either Gremlin or SPARQL to execute powerful queries that are easy to write and perform well.
- **Storage that Automatically Scales:** Amazon Neptune will automatically grow the size of your database volume as your database storage needs grow. Your volume will grow in increments of 10 GB up to a maximum of 64 TB. You don't need to provision excess storage for your database to handle future growth.
- **Low Latency Read Replicas:** Increase read throughput to support high volume application requests by creating up to 15 database read replicas. Amazon Neptune replicas share the same underlying storage as the source instance, lowering costs and avoiding the need to perform writes at the replica nodes.
- **Instance Monitoring and Repair:** The health of your Amazon Neptune database and its underlying EC2 instance is continuously monitored. If the instance powering your database fails, the database and associated processes are automatically restarted. Neptune recovery does not require the potentially lengthy replay of database redo logs, so your instance restart times are typically 30 seconds or less. It also isolates the database buffer cache from database processes, allowing the cache to survive a database restart.
- **Multi-AZ Deployments with Read Replicas:** On instance failure, Amazon Neptune automates failover to one of up to 15 Neptune replicas you have created in any of three Availability Zones. If no Neptune replicas have been provisioned, in the case of a failure, Neptune will attempt to create a new database instance for you automatically.

- **Fault-tolerant and Self-healing Storage:** Each 10GB chunk of your database volume is replicated six ways, across three Availability Zones. Amazon Neptune uses fault-tolerant storage that transparently handles the loss of up to two copies of data without affecting database write availability and up to three copies without affecting read availability. Neptune's storage is also self-healing; data blocks and disks are continuously scanned for errors and replaced automatically.
- **Automatic, Continuous, Incremental Backups and Point-in-time Restore:** Amazon Neptune's backup capability enables point-in-time recovery for your instance. This allows you to restore your database to any second during your retention period, up until the last five minutes. Your automatic backup retention period can be configured up to thirty-five days. Automated backups are stored in Amazon S3, which is designed for 99.999999999% durability. Neptune backups are automatic, incremental, and continuous and have no impact on database performance.
- **Database Snapshots:** Database Snapshots are user-initiated backups of your instance stored in Amazon S3 that will be kept until you explicitly delete them. They leverage the automated incremental snapshots to reduce the time and storage required. You can create a new instance from a Database Snapshot whenever you desire.
- **Encryption:** Amazon Neptune allows you to encrypt your databases using keys you create and control through AWS Key Management Service (KMS). On a database instance running with Neptune encryption, data stored at rest in the underlying storage is encrypted, as are the automated backups, snapshots, and replicas in the same cluster.
- **Automatic Software Patching:** Amazon Neptune will keep your database up-to-date with the latest patches. You can control if and when your instance is patched via Database Engine Version Management.
- **Database Event Notifications:** Amazon Neptune can notify you via email or SMS of important database events like automated failover. You can use the AWS Management Console to subscribe to different database events associated with your Amazon Neptune databases.
- **Fast Database Cloning:** Amazon Neptune supports quick, efficient cloning operations, where entire multi-terabyte database clusters can be cloned in minutes. You can clone an Amazon Neptune database with just a few clicks in the Management Console, without impacting the production environment. The clone is distributed and replicated across 3 Availability Zones.
- **Property Graph Bulk Loading:** Amazon Neptune supports fast, parallel bulk loading for Property Graph data that is stored in S3. You can use a REST interface to specify the S3 location for the data. It uses a CSV delimited format to load data into the Nodes and Edges. See the Neptune Property Graph bulk loading documentation for more details.
- **RDF Bulk Loading:** Amazon Neptune supports fast, parallel bulk loading for RDF data that is stored in S3. You can use a REST interface to specify the S3 location for the data. The N-Triples (NT), N-Quads (NQ), RDF/XML, and Turtle RDF 1.1 serializations are supported. See the Neptune RDF bulk loading documentation for more details.

66.1.2. Benefits

- **High performance:** Build and run identity, knowledge, fraud graph, and other applications with performance, and execute more than 100,000 queries per second.

- **Supported APIs:** Deploy high performance graph applications using popular open-source APIs such as Gremlin, openCypher and SPARQL, and easily migrate existing applications.
- **Managed:** Operate graph databases without worrying about hardware provisioning, software patching, setup, configuration, or backups; and pay no upfront licensing costs.

66.2. Backup/Restore and Disaster Recovery

Amazon Neptune's backup capability enables point-in-time recovery for your instance. This allows you to restore your database to any second during your retention period, up until the last five minutes. Your automatic backup retention period can be configured up to thirty-five days. Automated backups are stored in Amazon S3, which is designed for 99.999999999% durability. Neptune backups are automatic, incremental, and continuous and have no impact on database performance.

66.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

66.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/neptune/index.html
- **Service quotas:** https://docs.aws.amazon.com/general/latest/gr/neptune.html
- **Service FAQs:** https://aws.amazon.com/neptune/faqs/

66.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/neptune/index.html and the following links for comprehensive technical documentation regarding this service.

- User Guide: Provides a conceptual overview of Amazon Neptune, detailed instructions for using the various features, and a guide to both Gremlin and SPARQL connections for developers.

67. Amazon OpenSearch Service

67.1. Service Overview

With Amazon OpenSearch Service, choose from a selection of open source engine options. You can deploy and run the latest versions of OpenSearch, as well as 19 versions of ALv2 Elasticsearch (7.10 and earlier). The service also includes visualization capabilities with OpenSearch Dashboards and Kibana (7.10 and earlier).

67.1.1. Features

- **Setup and configuration:** Getting started with Amazon OpenSearch Service is easy. You can set up and configure your Amazon OpenSearch Service cluster using the AWS Management Console or a single API call through the AWS Command Line Interface (CLI). You can specify the number of instances, instance types, storage options, and modify or delete existing clusters at any time.

- **In-place upgrades:** Amazon OpenSearch Service makes it easy to upgrade your OpenSearch and Elasticsearch clusters (up to version 7.10) to newer versions without any downtime, using in-place version upgrades. In-place upgrades eliminates the hassle of taking a manual snapshot, restoring it to a cluster running the newer version, and updating all your endpoint references.
- **Event monitoring and alerting:** Amazon OpenSearch Service provides built-in event monitoring and alerting, enabling you to monitor the data stored in your cluster and automatically send notifications based on pre-configured thresholds. Built using the OpenSearch alerting plugin, this feature lets you configure and manage alerts using your Kibana or OpenSearch Dashboards interface and the REST API.
- **Support for multiple query languages:** With Amazon OpenSearch Service, there's no need for OpenSearch query domain-specific language (DSL) proficiency. Write SQL queries with OpenSearch SQL or use the OpenSearch Piped Processing Language (PPL), a query language that lets you use pipe (|) syntax, to explore, discover, and query your data. OpenSearch Dashboards also includes a SQL and PPL workbench.
- **Integration with open source tools:** Amazon OpenSearch Service offers built-in OpenSearch Dashboards and Kibana (Elasticsearch version 7.10 and previous) and integrates with Logstash, so you can ingest and visualize your data using the open source tools you prefer.
- **Security:** With Amazon OpenSearch Service, you can securely connect your applications to your managed Elasticsearch (version 7.10 and previous) or OpenSearch environment from your Amazon Virtual Private Cloud (VPC) or via the public Internet, configuring network access using VPC security groups or IP-based access policies.
- **Storage tiering – UltraWarm:** UltraWarm is a warm storage tier that complements Amazon OpenSearch Service's hot storage tier by providing less expensive storage for older and less-frequently accessed data while still providing an interactive querying experience. UltraWarm stores data in Amazon S3 and uses custom, highly-optimized nodes, purpose-built on the AWS Nitro System, to cache, pre-fetch, and query that data quickly.
- **Storage tiering – Cold storage:** Cold storage is the lowest-cost storage option for Amazon OpenSearch Service, which allows you to retain infrequently accessed data in Amazon S3 and only pay for compute when you need it. Cold storage builds on UltraWarm, which provides specialized nodes that store data in Amazon S3 and uses a sophisticated caching solution to provide an interactive experience.

67.1.2. Benefits

- **Leading:** Operate OpenSearch with the leading contributor of the community-driven, open source software.
- **Fast:** Quickly search and analyse your unstructured and semi-structured data to easily find what you need.
- **Efficient:** Eliminate operational overhead and reduce cost with automated provisioning, software installation, patching, storage tiering, and more.

- **Machine learning:** Use machine learning (ML) to detect anomalies in real time, autotune your clusters, and personalize your search results.

67.2. Backup/Restore and Disaster Recovery

This service allows for backup and recovery. The service can back up block volumes. Users control this via a manual or scheduled API call. Users schedule and recover backups through a web interface.

67.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

67.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/opensearch-service/
- **Service quotas:** https://docs.aws.amazon.com/general/latest/gr/opensearch-service.html
- **Service FAQs:** https://aws.amazon.com/opensearch-service/faqs/

67.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/opensearch-service/ and the following links for comprehensive technical documentation regarding this service.

- Developer Guide: Describes how to get started with OpenSearch Service, explains key concepts, and provides step-by-step instructions that show you how to use the features.

68. Amazon Personalize

68.1. Service Overview

Amazon Personalize enables developers to build applications with the same machine learning (ML) technology used by Amazon.com for real-time personalized recommendations – no ML expertise required.

Amazon Personalize makes it easy for developers to build applications capable of delivering a wide array of personalization experiences, including specific product recommendations, personalized product re-ranking, and customized direct marketing. Amazon Personalize is a fully managed machine learning service that goes beyond rigid static rule based recommendation systems and trains, tunes, and deploys custom ML models to deliver highly customized recommendations to customers across industries such as retail and media and entertainment.

Amazon Personalize provisions the necessary infrastructure and manages the entire ML pipeline, including processing the data, identifying features, using the best algorithms, and training, optimizing, and hosting the models. You will receive results via an Application Programming Interface (API) and only pay for what you use, with no minimum fees or upfront commitments. All data is encrypted to be private and secure, and is only used to create recommendations for your users.

68.1.1. Features

- **Use case optimized recommenders for retail and media and entertainment:** Introducing new recommenders that make it faster and easier to deliver high-performing

personalized user experiences. You can choose from use cases like “Frequently Bought Together,” “Because You Watched X,” “Top Picks for You,” and more.

- **User segmentation:** Amazon Personalize now offers intelligent user segmentation so you can run more effective prospecting campaigns through your marketing channels. With our two new recipes, you can automatically segment your users based on their interest in different product categories, brands, and more.
- **Automated machine learning:** Amazon Personalize takes care of machine learning for you. Once you have provided your data via Amazon S3 or via real-time integrations, Amazon Personalize can automatically load and inspect the data, lets you to select the right algorithms, train a model, provide accurate metrics, and generate personalized recommendations.
- **Real-time recommendations:** Make your recommendations relevant by responding to the changing intent of your users in real time.
- **Batch recommendations:** Compute recommendations for very large numbers of users or items in one go, store them, and feed them to batch-oriented workflows such as email systems.
- **New user and new item recommendations:** Effectively generate recommendations even for new users and find relevant new item recommendations for your users.
- **Contextual recommendations:** Improve relevance of recommendations by generating them within a context, for instance device type, time of day, and more.
- **Similar item recommendations:** Improve the discoverability of your catalogue by surfacing similar items to your users.
- **Unlock information in unstructured text:** Unlock the information trapped in product descriptions, reviews, movie synopses, or other unstructured text to generate highly relevant recommendations for users. Provide unstructured text as part of your catalogue, and Amazon Personalize automatically extracts key information to use when generating recommendations.
- **Prioritizing your business goals and what is relevant for your users:** Consider what’s relevant to your users and what is important for your business when generating recommendations. You can define an objective, in addition to relevance, to influence recommendations. This can be used to maximize for streaming minutes, increase revenue lift, or any metric you define as important to your business.
- **Easily integrate with your existing tools**: Amazon Personalize can be easily integrated into websites, mobile apps, or content management and email marketing systems, via a simple inference API call. The service lets you generate user recommendations, similar item recommendations and personalized re-ranking of items. You simply call the Amazon Personalize APIs and the service will output item recommendations or a re-ranked item list in a JSON format, which you can use in your application.
- **GetRecommendations API:** returns a list of relevant items given a userID. A representative usage example would be a content recommendation widget on landing page of a video streaming website that suggests a list of videos based on the user’s past watches. The API can also be used to return a list of similar itemIDs given an input

itemID. A representative use case is to recommend similar movies when a user is on the detail page of a movie.

- **GetPersonalizedRanking API:** re-ranks a list of itemIDs given a userID and a list of itemIDs to be re-ranked. The input list can be from any source, for example from an editorially curated list or from a list of itemIDs resulting from a search query. For example, an ecommerce retailer can use what they know about their customers' previous behaviour and past purchases to show the most relevant results, instead of showing the list of products that directly match the keyword.

68.1.2. Benefits

- **Deliver high quality recommendations, in real-time:** The ML algorithms used by Amazon Personalize create higher quality recommendations that respond to the specific needs, preferences, and changing behaviour of your users, improving engagement and conversion. They are also designed to address complex problems such as creating recommendations for new users, products, and content with no historical data.
- **Easily implement personalized recommendations in days, not months:** With Amazon Personalize, you can implement a customized personalization recommendation system, powered by ML, in just a few clicks without the burden of building, training, and deploying a "do it yourself" ML solution.
- **Personalize every touchpoint along the customer journey:** Amazon Personalize easily integrates into your existing websites, apps, SMS, and email marketing systems to provide a unique customer experience for across all channels and devices eliminating high infrastructure or resource costs. Amazon Personalize provides flexibility for you to use real-time or batch recommendations based on what is most appropriate for your use case, enabling you to deliver a wide variety of personalized experiences to customers at scale.
- **Data privacy and security:** All of your data is encrypted to be private and secure, and is only used to create recommendations for your customers. Data is not shared between customers or with Amazon.com. You can also use one of your own AWS Key Management Service (AWS KMS) keys to gain more control over access to data you encrypt. AWS KMS enables you to maintain control over who can use your customer master keys and gain access to your encrypted data.

68.2. Backup/Restore and Disaster Recovery

This requirement is not applicable for this service. For additional information beyond what is described herein, please refer to http://aws.amazon.com/documentation/.

68.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

68.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/personalize/
- **Service quotas:** https://docs.aws.amazon.com/personalize/latest/dg/limits.html
- **Service FAQs:** https://aws.amazon.com/personalize/faqs/

68.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/personalize/ and the following links for comprehensive technical documentation regarding this service.

- Developer Guide: Provides a conceptual overview of Amazon Personalize. Includes detailed instructions for using the features and provides a complete API reference for developers.

69. Amazon Pinpoint

69.1. Service Overview

Amazon Pinpoint is a flexible and scalable outbound and inbound marketing communications service. You can connect with customers over channels like email, SMS, push, voice or in-app messaging. Amazon Pinpoint is easy to set up, easy to use, and is flexible for all marketing communication scenarios. Segment your campaign audience for the right customer and personalize your messages with the right content. Delivery and campaign metrics in Amazon Pinpoint measure the success of your communications. Amazon Pinpoint can grow with you and scales globally to billions of messages per day across channels.

69.1.1. Features

- **Communication channels:** Amazon Pinpoint email, voice, push notification, and SMS channels offers deliverability and scale to reach hundreds of millions of customers around the globe.
- **Email:** Building a large-scale email solution can be a complex and costly challenge for a business. Organizations have to build infrastructure, set up your network, warm up your IP addresses, and protect your sender reputation. Many third-party email solutions come with contracts, minimum charges, and up-front costs. With Amazon Pinpoint, you can start sending emails in minutes, and you only pay for what you use.
- **SMS:** A majority of mobile phone users read incoming SMS messages almost immediately after receiving them. If you need to be able to provide your customers with urgent or essential information, SMS messaging may be the right solution for you. Businesses in a wide variety of industries can also use two-way SMS to keep their customers informed and engaged. For example, medical practices can send messages to their patients asking them to confirm their appointments.
- **Push notifications:** If you need to send important or time-sensitive information to the users of your mobile apps, you can use mobile push notifications. Push notifications are one of the most affordable ways to reach customers around the world and across a wide variety of devices. Your customers don't even have to be using your app in order to receive mobile push notifications.
- **In-app Messaging:** High-value actions are key to providing a good customer experience and improving business metrics within your product. In-app messaging allows targeted message sends displayed within mobile or web applications to drive these actions.
- **Marketing messages:** Send the right message to the right person at the right time by using Amazon Pinpoint for promotional marketing communication.

- **Audience segmentation:** You can create segments based on either real-time data or static lists. Dynamic segments use real-time customer attributes, which ensures your data is always up to date — including results from previous campaigns.
- **Templates and personalization:** Create content that drives results. Create reusable content templates across all Pinpoint projects. Use attributes like name for basic personalization, or drive more dynamic content by integrating your templates with ML data models using Amazon Personalize. Amazon Personalize integrates real-time personalization and recommendation data natively in Amazon Pinpoint.
- **Campaigns:** Specify schedules for Amazon Pinpoint campaigns, or set them up to execute when your customers perform specific actions in real-time. Send test emails to internal teams for quality assurance prior to starting your campaign, or perform A/B testing to determine the best possible content or time for your customer base.
- **Journeys:** With Amazon Pinpoint journeys, you can create multichannel and multi-step experiences for your customers. When you build a journey, you choose the activities that you want to add to each communications touch-point. These activities can perform a variety of different actions across channels, like sending an SMS to journey participants, then waiting a defined period of for a follow-up email.
- **Transactional messages:** Transactional messages are on-demand messages that you send to specific recipients. You can use the Amazon Pinpoint API and the AWS SDKs to send transactional messages through email, push, SMS, or voice.
- **Mobile and web analytics:** Understanding how your customers use your mobile and web applications is critical to improving your customer communications efforts and your products. Amazon Pinpoint collects usage attributes and metrics to help you identify trends in how customers are interacting with your applications.
- **Transactional and campaign results:** Amazon Pinpoint offers rich analytics related to the performance of your communications. Metrics like open rates, clicks, etc., on your campaigns and transactional messages allow you to understand historical trends and identify areas of improvement.
- **Deliverability:** The deliverability dashboard helps you identify and address issues that could impact the delivery of the emails that you send. Increase the chances that the emails you send arrive in your customers' inboxes, instead of their junk mail folders.

69.1.2. Benefits

- **Get started quickly:** Whether you are a marketer or a developer, Amazon Pinpoint is flexible for marketing, bulk, or transactional communications use cases. Marketers can design, orchestrate, and run campaigns visually through the console. Developers can leverage the Amazon Pinpoint APIs for message sending, scheduling campaigns, or tracking web and mobile activities. Send across channels like email, SMS, or push notifications.
- **Segment and personalize for impact:** Segment your audience for the right group of customers based on existing customer lists, attributes or use Amazon Pinpoint to create segments from mobile and web application data. Personalize the right message content to engage and delight your customers using both static and dynamic attributes. Marketers can also visually create a customer journey that automates multi-step campaigns.

- **Measure your efficiency:** From message delivery results to campaign data like opens and clicks, use metrics to understand the success of your communications. Amazon Pinpoint updates your customer lists with results to reflect learnings for the next campaign. View the campaign metrics natively in the Amazon Pinpoint reports or stream data to nearly any destination.
- **Scale securely with the experts:** Based on the scale and security of AWS, Amazon Pinpoint provides a reliable customer experience that can grow with you. Amazon has relationships with the top email providers, telecoms, and spam advisories to ensure the highest customer delivery rates.

69.2. Backup/Restore and Disaster Recovery

This requirement is not applicable for this service. For additional information beyond what is described herein, please refer to http://aws.amazon.com/documentation/.

69.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

69.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/pinpoint/
- **Service quotas:** https://docs.aws.amazon.com/personalize/latest/dg/limits.html
- **Service FAQs:** https://aws.amazon.com/pinpoint/faqs/

69.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/pinpoint/ and the following links for comprehensive technical documentation regarding this service.

- User Guide: Describes key concepts for Amazon Pinpoint and provides instructions for using the Amazon Pinpoint console.
- Developer Guide: Describes how to integrate Amazon Pinpoint functionality into your app and includes development instructions for its features.

70. Amazon Polly

70.1. Service Overview

Amazon Polly is a service that turns text into lifelike speech, allowing you to create applications that talk, and build entirely new categories of speech-enabled products. Polly's Text-to-Speech (TTS) service uses advanced deep learning technologies to synthesize natural sounding human speech. With dozens of lifelike voices across a broad set of languages, you can build speech-enabled applications that work in many different countries.

In addition to Standard TTS voices, Amazon Polly offers Neural Text-to-Speech (NTTS) voices that deliver advanced improvements in speech quality through a new machine learning approach. Polly’s Neural TTS technology also supports a Newscaster speaking style that is tailored to news narration use cases.

Finally, Amazon Polly Brand Voice can create a custom voice for your organization. This is a custom engagement where you will work with the Amazon Polly team to build an NTTS voice for the exclusive use of your organization.

70.1.1. Features

- **Simple-to-Use API:** Amazon Polly provides an API that enables you to quickly integrate speech synthesis into your application. You simply send the text you want converted into speech to the Amazon Polly API, and Amazon Polly immediately returns the audio stream to your application so your application can begin streaming it directly or store it in a standard audio file format, such as MP3.
- **Wide Selection of Voices and Languages:** Amazon Polly includes dozens of lifelike voices and support for a variety of languages, so you can select the ideal voice and distribute your speech-enabled applications in many countries. In addition to Standard TTS voices, Amazon Polly offers Neural Text-to-Speech (NTTS) voices that improve speech quality for more natural and human-like voices.
- **Synchronize Speech for an Enhanced Visual Experience:** Amazon Polly makes it easy to request an additional stream of metadata that provides information about when particular sentences, words and sounds are being pronounced. Using this metadata stream alongside the synthesized speech audio stream, you can now build your applications with an enhanced visual experience, such as speech-synchronized facial animation or karaoke-style word highlighting.
- **Optimize Your Streaming Audio:** With Amazon Polly, you can stream all kinds of information through your application to users in near real time. You can also choose from various sampling rates to optimize bandwidth and audio quality for your application. Amazon Polly supports MP3, Vorbis, and raw PCM audio stream formats.
- **Adjust Speaking Style, Speech Rate, Pitch, and Loudness:** Amazon Polly supports Speech Synthesis Markup Language (SSML), a W3C standard, XML-based markup language for speech synthesis applications, and supports common SSML tags for phrasing, emphasis, and intonation. Custom Amazon SSML tags provide unique options, such as the ability to make certain voices speak in a Newscaster speaking style. This flexibility helps you create lifelike speech that will attract and hold the attention of your audience.
- **Adjust the Maximum Duration of Speech:** Amazon Polly enables you to automatically adjust the speech rate based on a maximum allotted amount of time you define with a feature called time-driven prosody. This is beneficial for many use cases, especially when it comes to localization.
- **Platform and Programming Language Support:** Amazon Polly supports all the programming languages included in the AWS SDK (Java, Node.js, .NET, PHP, Python, Ruby, Go, and C++) and AWS Mobile SDK (iOS/Android). Polly also supports an HTTP API so you can implement your own access layer.
- **Speech Synthesis via API, Console, or Command Line:** Amazon Polly can be accessed via the Polly API (and various language-specific SDKs), AWS Management Console, and the AWS command-line interface (CLI). You have full control over all the capabilities of Amazon Polly, whether you use the service through the console, the API, or the CLI.

- **Custom Lexicons:** With Amazon Polly's custom lexicons, or vocabularies, you can modify the pronunciation of particular words, such as company names, acronyms, foreign words and neologisms (e.g., "ROTFL", "C'est la vie" when spoken in a non-French voice). To customize these pronunciations, you upload an XML file with lexical entries.
- **Brand Voice:** Brand Voice is a custom engagement where you work with the Amazon Polly team to build an Neural Text-to-Speech (NTTS) voice for the exclusive use of your organization.

70.1.2. Benefits

- **Natural sounding voices:** Amazon Polly provides dozens of languages and a wide selection of natural-sounding male and female voices. Amazon Polly's fluid pronunciation of text enables you to deliver high-quality voice output for a global audience.
- **Store & redistribute speech:** Amazon Polly allows for unlimited replays of generated speech without any additional fees. You can create speech files in standard formats like MP3 and OGG, and serve them from the cloud or locally with apps or devices for offline playback.
- **Real-time streaming:** Delivering lifelike voices and conversational user experiences requires consistently fast response times. When you send text to Amazon Polly's API, it returns the audio to your application as a stream so you can play the voices immediately.
- **Customize & control speech output:** Modify Amazon Polly voices to best suit your needs – Amazon Polly supports lexicons and SSML tags which enable you to control aspects of speech, such as pronunciation, volume, pitch, speed rate, etc.
- **Low cost:** Amazon Polly's pay-as-you-go pricing, low cost per character converted, and unlimited replays make it a cost-effective way to voice your applications.

70.2. Backup/Restore and Disaster Recovery

This requirement is not applicable for this service. For additional information beyond what is described herein, please refer to http://aws.amazon.com/documentation/.

70.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

70.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/polly/
- **Service quotas:** https://docs.aws.amazon.com/polly/latest/dg/limits.html
- **Service FAQs:** https://aws.amazon.com/polly/faqs/

70.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/polly/ and the following links for comprehensive technical documentation regarding this service.

- Developer Guide: Provides a conceptual overview of Amazon Polly, including detailed instructions for using the various features. Provides a complete API reference and example applications for developers.

71. Amazon Quantum Ledger Database (QLDB)

71.1. Service Overview

Amazon QLDB is a fully managed ledger database that provides a transparent, immutable, and cryptographically verifiable transaction log owned by a central trusted authority. Amazon QLDB tracks each and every application data change and maintains a complete and verifiable history of changes over time.

71.1.1. Features

- **Append-only journal:** Amazon QLDB has a built-in immutable journal that stores an accurate and sequenced entry of every data change. The journal is append-only, meaning that data can only be added to a journal and it cannot be overwritten or deleted. This ensures that your stored change history cannot be deleted or modified. Even if you delete the data from your ledger, the change history of that data can still be accessed by reading from the immutable journal.
- **Easy access to change history:** With Amazon QLDB, you can access the entire change history of your application's data. You can query a summary of historical changes (e.g. list of all previous owners of a vehicle), and also specific details related to transaction history (e.g. the time of a vehicle sale and name of the new owner).
- **Change History Digest:** Amazon QLDB uses cryptography to create a concise summary of your change history. This secure summary, commonly known as a digest, is generated using a cryptographic hash function (SHA-256). The digest acts as a proof of your data's change history, allowing you to look back and verify the integrity of your data changes. You can use this digest with QLDB's API to prove the integrity of any transaction (e.g. whether a transaction occurred or not).
- **Easy to Scale:** Amazon QLDB delivers seamless, automatic scaling to meet the demands of your application without the need to provision capacity or configure read and write limits. Also, since QLDB is a database, it provides better performance and scale than blockchain frameworks. QLDB can easily scale up and execute 2-3x as many transactions as common blockchain frameworks.
- **Easy setup:** Getting started with Amazon QLDB is easy as there are no servers to manage or capacity to provision. You can create a new ledger in minutes using the AWS Management Console, AWS Command Line Interface (CLI), an AWS CloudFormation template, or by making calls to the QLDB API.
- **Monitoring and metrics:** Amazon QLDB provides Amazon CloudWatch metrics for your ledgers. With QLDB, you can view key operational metrics for your read and write IOs.
- **PartiQL Support:** Amazon QLDB supports PartiQL, which is a new open standard query language. PartiQL supports SQL-compatible access to QLDB's document-oriented data model that includes semi-structured and nested data while remaining independent of any particular data source. With PartiQL you can easily query, manage, and update your data using familiar SQL operators.

- **Document-oriented data model:** Data models define how data is processed and stored inside a database. Amazon QLDB stores data using a document-oriented data model, which provides you the flexibility to store structured and semi-structured data. QLDB's data model also supports nested data structures, which can simplify your applications.
- **Transactional Consistency and ACID Semantics:** When performing a database operation, Amazon QLDB provides atomicity, consistency, isolation, and durability (ACID) properties. Also, QLDB transactions have full serializability- the highest level of isolation. The ACID properties of transactions make it easy to write correct applications.
- **Streaming Capability:** Amazon QLDB's streaming capability provides a near real-time flow of any changes to your data stored in QLDB via Amazon Kinesis Data Streams. QLDB's stream data always retains the core QLDB characteristics of "complete & verifiable" data storage.
- **Event Driven Architecture:** You can build applications with an event-driven architecture using AWS Lambda. For example, a bank can implement a notification system that sends a text message or an email to a customer when the account balance drops below a certain threshold.
- **Analytics:** You can run analytics jobs on real-time or historical data. For example, an e-commerce website can run ad-hoc analytics to generate hourly aggregated metrics, such as number of t-shirts sold per day of a particular colour, from historic data. Amazon QLDB is able to provide this unique ability to replay historic event data, leveraging the Journal-first architecture of QLDB. You can choose to start a QLDB Stream from any point in time in the past and subsequent changes will be streamed to Amazon Kinesis.
- **Replication to Purpose-Built Data Stores:** You can connect Amazon QLDB to other purpose-built data stores. For example, a bank can provide powerful text search capabilities to find debit and credit transactions in an account, using Amazon Elasticsearch. You can also replicate to other purpose-built data stores that provide a different materialized view, such as graph-based view using Amazon Neptune, enabling them to use the best tool for the job.

71.1.2. Benefits

- **Tracking:** Track and maintain a sequenced history of every application data change using an immutable and transparent journal.
- **Integrity:** Trust the integrity of your data. Built-in cryptographic verification enables third-party validation of data changes.
- **QLDB ACID:** Build correct, event-driven systems with QLDB ACID transactions and support for real-time streaming to Amazon Kinesis.
- **Value:** Start small and pay only for what you use with serverless architecture that provides automatic storage and resource scaling.

71.2. Backup/Restore and Disaster Recovery

This requirement is not applicable for this service. For additional information beyond what is described herein, please refer to http://aws.amazon.com/documentation/.

71.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

71.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/qldb/
- **Service quotas:** https://docs.aws.amazon.com/qldb/latest/developerguide/limits.html
- **Service FAQs:** https://aws.amazon.com/qldb/faqs/

71.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/qldb/ and the following links for comprehensive technical documentation regarding this service.

- Developer Guide: Provides a conceptual overview of Amazon QLDB and includes detailed development instructions for using QLDB features.

72. Amazon QuickSight

72.1. Service Overview

Amazon QuickSight allows everyone in your organization to understand your data by asking questions in natural language, exploring through interactive dashboards, or automatically looking for patterns and outliers powered by machine learning.

72.1.1. Features

- **Serverless auto-scaling:** QuickSight is serverless and can automatically scale to tens of thousands of users without any infrastructure to manage or capacity to plan for. QuickSight gives users and analysts self-service business intelligence (BI), so they can answer their own questions, collaborate, and share insights. With QuickSight, your users can connect to data sources, create/edit datasets, create visual analyses, invite co-workers to collaborate on analyses, and publish dashboards and reports.
- **Broad data source support:** QuickSight allows you to directly connect to and import data from a wide variety of cloud and on-premises data sources. These include SaaS applications such as Salesforce, Square, ServiceNow, Twitter, Github, and JIRA; 3rd party databases such as Teradata, MySQL, Postgres, and SQL Server; native AWS services such as Redshift, Athena, S3, RDS, and Aurora; and private VPC subnets. You can also upload a variety of file types including Excel, CSV, JSON, and Presto.
- **SPICE (super-fast, parallel, in-memory, calculation engine):** With SPICE, QuickSight's in-memory calculation engine you achieve blazing fast performance at scale. SPICE automatically replicates data for high availability allowing thousands of users to simultaneously perform fast, interactive analysis while shielding your underlying data infrastructure, saving you time and resources.
- **Global collaboration and multi-tenancy:** As a native AWS service with customers all over the world, QuickSight has been designed and built as a global product from the beginning. The QuickSight application is localized in 10 major languages including: English, German, Spanish, French, Italian, Portuguese, Japanese, Korean, Simplified

Chinese, and Traditional Chinese. QuickSight is also available across multiple AWS regions including: N. Virginia, Oregon, Ohio, Dublin, Japan, Singapore, and Sydney.

- **Built-in security and compliance:** Quicksight provides a secure platform allowing you to distribute dashboards and insights securely to tens of thousands of users. In addition to the multi-region availability and built-in redundancy, QuickSight allows you to securely manage your users and content via a comprehensive set of security features including role-based access control, active directory integration, CloudTrail auditing, single sign-on (IAM, 3rd party), private VPC subnets, and data backup. QuickSight is also FedRamp, HIPAA, PCI PSS, ISO, and SOC compliant to help you meet any industry-specific or regulatory requirements.
- **Mobile app support:** QuickSight Mobile for iOS and Android enables you to securely get insights from your data from anywhere. Easily favourite, browse and interact with all of your dashboards in an easy to use mobile optimized experience. You can explore your data with drill down and filters, stay ahead of the curve via forecasting, get email alerts when unexpected changes happen in your data, and share those insights with colleagues. QuickSight mobile is available as a FREE download for all QuickSight users from the App Store, and Google Play Store.
- **Pay-per-use pricing:** QuickSight offers a unique, industry first pay-per-session model for dashboard readers, users who consume dashboards others have created. Instead of paying a fixed license cost per month, readers are billed $0.30 for a 30-minute session up to a maximum charge of $5/reader/month for unlimited use. This pricing model allows all of your users to access secure, interactive dashboards and email reports on a pay-per-session basis with no upfront costs or complex capacity planning.

72.1.2. Benefits

- **Self-service BI with QuickSight Q:** With QuickSight Q, anyone in an organization can ask business questions in natural language and receive accurate answers with relevant visualizations that help them gain insights from the data. QuickSight Q uses machine learning to interpret the intent of a question and analyse the correct data to provide accurate answers to business questions quickly.
- **Embedded analytics for all applications:** Amazon QuickSight lets you to quickly embed interactive dashboards and visualizations into your applications without needing to build your own analytics capabilities. Blend analytics seamlessly into your application with QuickSight's Embedded themes, which let you personalize the look and feel (e.g. colours, fonts etc.) of your reports and dashboards. Integrate into your applications user experience by setting defaults and handling errors to match your app's user experience.
- **Advanced analytics with ML Insights:** ML Insights leverages AWS’s proven machine learning (ML) and natural language capabilities to help you gain deeper insights from your data. These powerful, out-of-the-box features make it easy for anyone to discover hidden trends and outliers, identify key business drivers, and perform powerful what-if analysis and forecasting with no technical expertise or ML experience needed.
- **Built for enterprise workloads:** Amazon QuickSight has a serverless architecture that automatically scales to tens of thousands of users without the need to setup, configure, or manage your own servers. It also ensures that your users don’t have to deal with slow dashboards during peak-hours, when multiple BI users are accessing the same dashboards or datasets. And with a pay-per-session pricing, you only pay when your

users access the dashboards or reports, which makes it cost-effective for deployments with lots of users.

72.2. Backup/Restore and Disaster Recovery

This requirement is not applicable for this service. For additional information beyond what is described herein, please refer to http://aws.amazon.com/documentation/.

72.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

72.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/quicksight/
- **Service quotas:** https://docs.aws.amazon.com/general/latest/gr/quicksight.html
- **Service FAQs:** https://aws.amazon.com/quicksight/resources/faqs/

72.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/quicksight/ and the following links for comprehensive technical documentation regarding this service.

- User Guide: Describes all Amazon QuickSight concepts and features, and provides instructions on using these features in the Amazon QuickSight web application.
- Developer Guide: Provides usage examples of API operations for Amazon QuickSight and procedural walkthroughs of common tasks.

73. Amazon Redshift

73.1. Service Overview

Amazon Redshift is a fast, fully managed data warehouse that makes it simple and cost effective to analyze all your data using standard SQL and your existing business intelligence tools. It allows you to run complex analytic queries against petabytes of structured data by using sophisticated query optimization, columnar storage on high-performance local disks, and massively parallel query execution. Most results come back in seconds.

73.1.1. Features

- **Query Editor v2:** Use SQL to make your Amazon Redshift data and data lake more accessible to data analysts, data engineers, and other SQL users with a web-based analyst workbench for data exploration and analysis. Query Editor v2 lets you visualize query results in a single click, create schemas and tables, load data visually, and browse database objects. It also provides an intuitive editor for authoring and sharing SQL queries, analyses, visualizations, and annotations, and securely sharing them with your team.
- **Automated Table Design:** Amazon Redshift monitors user workloads and uses sophisticated algorithms to find ways to improve the physical layout of data to optimize query speeds. Automatic Table Optimization selects the best sort and distribution keys to optimize performance for the cluster's workload. If Amazon Redshift determines that

applying a key will improve cluster performance, tables will be automatically altered without requiring administrator intervention. The additional features Automatic Vacuum Delete, Automatic Table Sort, and Automatic Analyze eliminate the need for manual maintenance and tuning of Redshift clusters to get the best performance for new clusters and production workloads.

- **Query using your own tools:** Amazon Redshift gives you the flexibility to run queries within the console or connect SQL client tools, libraries, or data science tools including Amazon Quicksight, Tableau, PowerBI, QueryBook and Jupyter Notebook.
- **Federated query**: With the new federated query capability in Amazon Redshift, you can reach into your operational relational database. Query live data across one or more Amazon Relational Database Service (RDS) and Aurora PostgreSQL and RDS MySQL and Aurora MySQL databases to get instant visibility into the full business operations without requiring data movement. You can join data from your Redshift data warehouse, data in your data lake, and data in your operational stores to make better data-driven decisions. Amazon Redshift offers sophisticated optimizations to reduce data moved over the network and complements it with its massively parallel data processing for high-performance queries.
- **Query and export data to and from your data lake**: No other cloud data warehouse makes it as easy to both query data and write data back to your data lake in open formats. You can query open file formats such as Parquet, ORC, JSON, Avro, CSV, and more directly in S3 using familiar ANSI SQL. To export data to your data lake, simply use the Amazon Redshift UNLOAD command in your SQL code and specify Parquet as the file format, and Amazon Redshift automatically takes care of data formatting and data movement into S3. This gives you the flexibility to store highly structured, frequently accessed data and semi-structured data in an Amazon Redshift data warehouse, while keeping up to exabytes of structured, semi-structured and unstructured data in Amazon S3. Exporting data from Amazon Redshift back to your data lake lets you analyse the data further with AWS services such as Amazon Athena, Amazon EMR, and Amazon SageMaker.
- **AWS Data Exchange for Amazon Redshift**: Query Amazon Redshift datasets from your own Redshift cluster without extracting, transforming, and loading ETL the data. You can subscribe to Redshift cloud data warehouse products in AWS Data Exchange. As soon as a provider makes an update, the change is visible to subscribers. If you are a data provider, access is automatically granted when a subscription starts and revoked when it ends, invoices are automatically generated when payments are due, and payments are collected through AWS. You can license access to flat files, data in Amazon Redshift, and data delivered through APIs, all with a single subscription.
- **Redshift ML**: Redshift ML makes it easy for data analysts, data scientists, BI professionals, and developers to create, train, and deploy Amazon SageMaker models using SQL. With Redshift ML, you can use SQL statements to create and train Amazon SageMaker models on your data in Amazon Redshift and then use those models for predictions such as churn detection, financial forecasting, personalization, and risk scoring directly in your queries and reports.
- **RA3 instances**: RA3 instances deliver up to 3x better price performance of any cloud data warehouse service. These Amazon Redshift instances maximize speed for performance-intensive workloads that require large amounts of compute capacity, with

the flexibility to pay separately for compute independently of storage by specifying the number of instances you need.

- **Advanced Query Accelerator (AQUA) for Amazon Redshift**: AQUA is a new distributed and hardware-accelerated cache that enables Amazon Redshift to run up to 10x faster than other enterprise cloud data warehouses by automatically boosting certain types of queries. AQUA uses high-speed solid state storage, field-programmable gate arrays (FPGAs), and AWS Nitro to speed queries that scan, filter, and aggregate large datasets. AQUA is included with the Redshift RA3 instance type at no additional cost.
- **Petabyte-scale data warehousing**: With a few clicks in the console or a simple API call, you can easily change the number or type of nodes in your data warehouse, and scale up or down as your needs change. With managed storage, capacity is added automatically to support workloads up to 8 PB of compressed data. You can also run queries against petabytes of data in Amazon S3 without having to load or transform any data with the Amazon Redshift Spectrum feature. You can use S3 as a highly available, secure, and cost-effective data lake to store unlimited data in open data formats. Redshift Spectrum runs queries across thousands of parallelized nodes to deliver fast results, regardless of the complexity of the query or the amount of data.

73.1.2. Benefits

- **End-to-end encryption**: With just a few parameter settings, you can set up Amazon Redshift to use SSL to secure data in transit, and hardware-accelerated AES-256 encryption for data at rest. If you choose to enable encryption of data at rest, all data written to disk will be encrypted as well as any backups. Amazon Redshift takes care of key management by default.
- **Network isolation**: Amazon Redshift lets you configure firewall rules to control network access to your data warehouse cluster. You can run Amazon Redshift inside Amazon Virtual Private Cloud (VPC) to isolate your data warehouse cluster in your own virtual network and connect it to your existing IT infrastructure using an industry-standard encrypted IPsec VPN.
- **Audit and compliance**: Amazon Redshift integrates with AWS CloudTrail to enable you to audit all Redshift API calls. Redshift logs all SQL operations, including connection attempts, queries, and changes to your data warehouse. You can access these logs using SQL queries against system tables, or save the logs to a secure location in Amazon S3. Amazon Redshift is compliant with SOC1, SOC2, SOC3, and PCI DSS Level 1 requirements.
- **Flexible pricing options**: Amazon Redshift is the most cost-effective data warehouse, and you can optimize how you pay. You can start small for just $0.25 per hour with no commitments, and scale out for just $1,000 per terabyte per year. Amazon Redshift is the only cloud data warehouse that offers on-demand pricing with no upfront costs, Reserved Instance pricing that can save you up to 75% by committing to a one- or three-year term, and per-query pricing based on the amount of data scanned in your Amazon S3 data lake. Amazon Redshift's pricing includes built-in security, data compression, backup storage, and data transfer. As the size of data grows, you use managed storage in the RA3 instances to store data cost-effectively at $0.024 per GB per month.
- **Predictable cost, even with unpredictable workloads**: Amazon Redshift allows you to scale with minimal cost impact, as each cluster earns up to one hour of free Concurrency Scaling credits per day. These free credits are sufficient for the

concurrency needs of 97% of customers. This provides you with predictability in your month-to-month cost, even during periods of fluctuating analytical demand.

- **Machine learning to maximize throughput and performance**: Advanced ML capabilities in Amazon Redshift deliver high throughput and performance, even with varying workloads or concurrent user activity. Amazon Redshift uses sophisticated algorithms to predict and classify incoming queries based on their run times and resource requirements to dynamically manage performance and concurrency while also helping you prioritize your business-critical workloads. Short query acceleration (SQA) sends short queries from applications such as dashboards to an express queue for immediate processing rather than being starved behind large queries. Automatic workload management (WLM) uses ML to dynamically manage memory and concurrency, helping maximize query throughput. In addition, you can now easily set the priority of your most important queries, even when hundreds of queries are being submitted. Amazon Redshift is also a self-learning system that observes the user workload, determining the opportunities to improve performance as the usage grows, applying optimizations seamlessly, and making recommendations through Redshift Advisor when an explicit user action is needed to further turbocharge Redshift performance.
- **Result caching**: Amazon Redshift uses result caching to deliver sub-second response times for repeat queries. Dashboard, visualization, and business intelligence tools that run repeat queries experience a significant performance boost. When a query runs, Amazon Redshift searches the cache to see if there is a cached result from a prior run. If a cached result is found and the data has not changed, the cached result is returned immediately instead of re-running the query.
- **AWS services integration**: Native integration with AWS services, database, and machine learning services makes it easier to handle complete analytics workflows without friction. For example, AWS Lake Formation is a service that makes it easy to set up a secure data lake in days. AWS Glue can extract, transform, and load (ETL) data into Amazon Redshift. Amazon Kinesis Data Firehose is the easiest way to capture, transform, and load streaming data into Amazon Redshift for near real-time analytics. You can use Amazon EMR to process data using Hadoop/Spark and load the output into Amazon Redshift for BI and analytics. Amazon QuickSight is the first BI service with pay-per-session pricing that you can use to create reports, visualizations, and dashboards on Redshift data. You can use Amazon Redshift to prepare your data to run machine learning (ML) workloads with Amazon SageMaker. To accelerate migrations to Amazon Redshift, you can use the AWS Schema Conversion tool and the AWS Database Migration Service (DMS). Amazon Redshift is also deeply integrated with Amazon Key Management Service (KMS) and Amazon CloudWatch for security, monitoring, and compliance. You can also use Lambda user-defined functions (UDFs) to invoke a Lambda function from your SQL queries as if you are invoking a UDF in Amazon Redshift. You can write Lambda UDFs to integrate with AWS Partner services and to access other popular AWS services such as Amazon DynamoDB and Amazon SageMaker.
- **Partner console integration**: You can accelerate data onboarding and create valuable business insights in minutes by integrating with select Partner solutions in the Amazon Redshift console. With these solutions you can bring data from applications such as Salesforce, Google Analytics, Facebook Ads, Slack, Jira, Splunk, and Marketo into your Redshift data warehouse in an efficient and streamlined way. It also lets you join these disparate datasets and analyze them together to produce actionable insights.

- **Data Sharing**: Amazon Redshift data sharing allows you to extend the ease of use, performance, and cost benefits of Amazon Redshift in a single cluster to multi-cluster deployments while being able to share data. Data sharing enables instant, granular, and fast data access across Redshift clusters without the need to copy or move it. Data sharing provides live access to data so your users always see the most current and consistent information as it's updated in the data warehouse. You can securely share live data with Redshift clusters in the same or different AWS accounts and across Regions.

73.2. Backup/Restore and Disaster Recovery

This service allows for backup and recovery. The service can back up block volumes. Users control this via a manual or scheduled API call. Users schedule and recover backups through a web interface.

73.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

73.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/redshift/index.html
- **Service quotas:** https://docs.aws.amazon.com/redshift/latest/mgmt/amazon-redshift-limits.html
- **Service FAQs:** https://aws.amazon.com/redshift/faqs/?nc=sn&loc=5&dn=4

73.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/redshift/index.html and the following links for comprehensive technical documentation regarding this service.

- Redshift Getting Started Guide: Introduces you to Amazon Redshift, helps you set up an account, and walks you through a simple example to use Amazon Redshift for the first time. Also provides tips and links to advanced product features and resources.
- Redshift Database Developer Guide: Explains how to design, build, query, and maintain the databases that make up your data warehouse, and includes syntax for Redshift SQL commands and functions.
- Redshift Cluster Management Guide: Shows you how to create and manage Redshift clusters.

74. Amazon Rekognition

74.1. Service Overview

Amazon Rekognition Video is a machine learning powered video analysis service that detects objects, scenes, celebrities, text, activities, and any inappropriate content from your videos stored in Amazon S3. Rekognition Video also provides highly accurate facial analysis and facial search capabilities to detect, analyse, and compare faces, and helps understands the movement of people in your videos.

Each result or detection is paired with a timestamp so that you can easily create an index for detailed video search, or navigate quickly to an interesting part of the video for further analysis. For objects, faces, text, and people, Rekognition Video also returns bounding box coordinates, which is the specific location of the detection in the frame.

Amazon Rekognition Video can also monitor a live stream that you create from Amazon Kinesis Video Streams to detect and search faces from face data that you provide. You can incorporate audio analysis such as closed captioning, profanity filtering and streaming video transcription into your applications by using Amazon Transcribe along with Amazon Rekognition Video.

74.1.1. Features

- **Object, scene, and activity detection:** Amazon Rekognition Video automatically identifies thousands of objects such as vehicles or pets, scenes like a city, beach, or wedding, and activities such as delivering a package or dancing. For each label detected, you get a confidence score. For common objects such as 'Person' or 'Car', you also get object bounding boxes to enable counting and object localization. Amazon Rekognition Video relies on motion in the video to accurately identify complex activities, such as “blowing out a candle” or “extinguishing fire”. Using this rich metadata, you can make your content searchable or serve advertisements that best match the context of the content preceding it.
- **Content moderation:** Amazon Rekognition Video automatically detects inappropriate content such as nudity, violence or weapons in videos, and provides timestamps for each detection. You also get a hierarchical list of labels with confidence scores, describing sub-categories of unsafe content. For example, 'Graphic Female Nudity' is a sub-category of 'Explicit Nudity'. Confidence scores and detailed labels allow you to set up varied business rules to serve the compliance needs of different markets and geographies.
- **Text detection:** Amazon Rekognition Video automatically detect and read text in videos, and provides the detection confidence, location bounding box, as well as the timestamp for each text detection. In addition, you get convenient options to filter out words by regions of interest (ROIs), word bounding box size, and word confidence score. For example, you may only want to detect text in the bottom third region for on-screen graphics or only the top left corner for reading scoreboards in a soccer game.
- **Celebrity recognition:** With Amazon Rekognition Video, you can detect and recognize when and where well known persons appear in a video. The time-coded output includes the name and unique id of the celebrity, and URLs pointing to related content for the celebrity, for example, the celebrity's IMDB link.
- **Face detection and analysis:** Amazon Rekognition Video can detect up to 100 faces in a video frame, and return the bounding box location. For each detected face, you can also get additional attributes such as gender, emotions, estimated age range, and whether the person is smiling, along with timestamps for each detection.
- **Face search:** Amazon Rekognition Video can identify known people in a video by searching against a private repository of face images. You get a similarity score for each match, and timestamps for each instance where the same person is identified during the video. Amazon Rekognition Video can also cluster all unknown people in a video who don’t have any matches in the repository, and return timestamps with unique identifiers for each such person.

- **Person pathing:** With Amazon Rekognition Video you can capture where, when and how each person is moving in your video. Amazon Rekognition also provides a unique index for each person found, allowing you to count the number of people in the video.
- **Live stream video analysis**: Amazon Rekognition Video can analyze your live video streams in real time to detect and search for faces. By providing a stream from Amazon Kinesis Video Streams as an input to Rekognition Video, you can perform face search against a repository of your own images with very low latency.

74.1.2. Benefits

- **Detect inappropriate content:** Quickly and accurately identify unsafe or inappropriate content across image and video assets based on general or business-specific standards and practices.
- **Verify identity online**: Use facial comparison and analysis in your user onboarding and authentication workflows to remotely verify the identity of opted-in users.
- **Streamline media analysis:** Automatically detect key video segments to reduce the time, effort, and costs of video ad insertion, content operations, and content production.
- **Keep workers safe continuously:** Analyze workplace images and video to determine if employees are using PPE according to established guidelines.
- **Analyze Video**: Analyze workplace images and video to determine if employees are using PPE according to established guidelines.
- **Easy to use:** Quickly add pre-trained or customizable computer vision APIs to your applications without building machine learning (ML) models and infrastructure from scratch.
- **Huge capacity:** Analyze millions of images and videos within minutes and augment human visual review tasks with artificial intelligence (AI).
- **Scalable:** Scale up and down based on your business needs with fully managed AI capabilities and pay only for the images and videos you analyze.

74.2. Backup/Restore and Disaster Recovery

This requirement is not applicable for this service. For additional information beyond what is described herein, please refer to http://aws.amazon.com/documentation/.

74.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

74.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/rekognition/index.html
- **Service quotas:** https://docs.aws.amazon.com/rekognition/latest/dg/limits.html
- **Service FAQs:** https://aws.amazon.com/rekognition/faqs/?nc=sn&loc=7

74.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/rekognition/index.html and the following links for comprehensive technical documentation regarding this service.

- Developer Guide: Provides a conceptual overview of Amazon Rekognition, detailed instructions for using the various features, and a complete API reference for developers.
- Custom Labels Developer Guide: Provides information about creating an Amazon Rekognition Custom Labels model that predicts the presence of objects, scenes, and concepts in images. It includes instructions for training a model, evaluating a model, and using the model to make predictions.

75. Amazon Relational Database Service (RDS)

75.1. Service Overview

Amazon Relational Database Service (Amazon RDS) makes it easy to set up, operate, and scale a relational database in the cloud. It provides cost-efficient and resizable capacity while automating time-consuming administration tasks, such as hardware provisioning, database setup, patching, and backups. It frees you to focus on your applications so you can give them the fast performance, high availability, security, and compatibility they need.

Amazon RDS is available on several database instance types - optimized for memory, performance, or I/O - and provides you with six familiar database engines to choose from, including Amazon Aurora, PostgreSQL, MySQL, MariaDB, Oracle Database, and SQL Server. You can use the AWS Database Migration Service to easily migrate or replicate your existing databases to Amazon RDS.

75.1.1. Features

- **Easy to use:** You can use the AWS Management Console, the Amazon RDS Command Line Interface, or simple API calls to access the capabilities of a production-ready relational database in minutes. Amazon RDS database instances are pre-configured with parameters and settings appropriate for the engine and class you have selected. You can launch a database instance and connect your application within minutes. DB Parameter Groups provide granular control and fine-tuning of your database.
- **Automatic software patching:** Amazon RDS will make sure that the relational database software powering your deployment stays up-to-date with the latest patches. You can exert optional control over when and if your database instance is patched.
- **Best practice recommendations:** Amazon RDS provide best practice guidance by analyzing configuration and usage metrics from your database instances. Recommendations cover areas such as database engine versions, storage, instance types, and networking. You can browse the available recommendations and perform a recommended action immediately, schedule it for their next maintenance window, or dismiss it entirely.
- **General Purpose (SSD) Storage:** Amazon RDS General Purpose Storage is an SSD-backed storage option delivers a consistent baseline of 3 IOPS per provisioned GB and provides the ability to burst up to 3,000 IOPS above the baseline. This storage type is suitable for a broad range of database workloads.
- **Provisioned IOPS (SSD) Storage:** Amazon RDS Provisioned IOPS Storage is an SSD-backed storage option designed to deliver fast, predictable, and consistent I/O performance. You specify an IOPS rate when creating a database instance, and Amazon RDS provisions that IOPS rate for the lifetime of the database instance. This storage type is optimized for I/O-intensive transactional (OLTP) database workloads. You can provision up to 40,000 IOPS per database instance, although your actual

realized IOPS may vary based on your database workload, instance type, and database engine choice.

- **Push-button compute scaling**: You can scale the compute and memory resources powering your deployment up or down, up to a maximum of 32 vCPUs and 244 GiB of RAM. Compute scaling operations typically complete in a few minutes.
- **Easy storage scaling:** As your storage requirements grow, you can also provision additional storage. The Amazon Aurora engine will automatically grow the size of your database volume as your database storage needs grow, up to a maximum of 64 TB or a maximum you define. The MySQL, MariaDB, Oracle, and PostgreSQL engines allow you to scale up to 64 TB of storage and SQL Server supports up to 16 TB. Storage scaling is on-the-fly with zero downtime.
- **Read Replicas:** Read Replicas make it easy to elastically scale out beyond the capacity constraints of a single DB instance for read-heavy database workloads. You can create one or more replicas of a given source DB instance and serve high-volume application read traffic from multiple copies of your data, thereby increasing aggregate read throughput. Read replicas are available in Amazon RDS for MySQL, MariaDB, PostgreSQL, and Oracle as well as Amazon Aurora.
- **Automated backups:** The automated backup feature of Amazon RDS enables point-in-time recovery for your database instance. Amazon RDS will backup your database and transaction logs and store both for a user-specified retention period. This allows you to restore your database instance to any second during your retention period, up to the last five minutes. Your automatic backup retention period can be configured to up to thirty-five days.
- **Encryption at rest and in transit:** Amazon RDS allows you to encrypt your databases using keys you manage through AWS Key Management Service (KMS). On a database instance running with Amazon RDS encryption, data stored at rest in the underlying storage is encrypted, as are its automated backups, read replicas, and snapshots. Amazon RDS supports Transparent Data Encryption in SQL Server and Oracle. Transparent Data Encryption in Oracle is integrated with AWS CloudHSM, which allows you to securely generate, store, and manage your cryptographic keys in single-tenant Hardware Security Module (HSM) appliances within the AWS cloud. Amazon RDS supports the use of SSL to secure data in transit.
- **Pay only for what you use:** There is no up-front commitment with Amazon RDS; you simply pay a monthly charge for each database instance that you launch. And, when you're finished with a database instance, you can easily delete it. To see more details, visit the Amazon RDS Instance Types page and the Amazon RDS Pricing page.

75.1.2. Benefits

- **Easy to administer:** Amazon RDS makes it easy to go from project conception to deployment. Use the Amazon RDS Management Console, the AWS RDS Command-Line Interface, or simple API calls to access the capabilities of a production-ready relational database in minutes. No need for infrastructure provisioning, and no need for installing and maintaining database software.
- **Highly scalable:** You can scale your database's compute and storage resources with only a few mouse clicks or an API call, often with no downtime. Many Amazon RDS engine types allow you to launch one or more Read Replicas to offload read traffic from your primary database instance.

- **Available and durable:** Amazon RDS runs on the same highly reliable infrastructure used by other Amazon Web Services. When you provision a Multi-AZ DB Instance, Amazon RDS synchronously replicates the data to a standby instance in a different Availability Zone (AZ). Amazon RDS has many other features that enhance reliability for critical production databases, including automated backups, database snapshots, and automatic host replacement.
- **Fast:** Amazon RDS supports the most demanding database applications. You can choose between two SSD-backed storage options: one optimized for high-performance OLTP applications, and the other for cost-effective general-purpose use. In addition, Amazon Aurora provides performance on par with commercial databases at 1/10th the cost.
- **Secure:** Amazon RDS makes it easy to control network access to your database. Amazon RDS also lets you run your database instances in Amazon Virtual Private Cloud (Amazon VPC), which enables you to isolate your database instances and to connect to your existing IT infrastructure through an industry-standard encrypted IPsec VPN. Many Amazon RDS engine types offer encryption at rest and encryption in transit.
- **Value-for-Money:** You pay very low rates and only for the resources you actually consume. In addition, you benefit from the option of On-Demand pricing with no up-front or long-term commitments or even lower hourly rates via our Reserved Instance pricing.

75.2. Backup/Restore and Disaster Recovery

This service allows for backup and recovery. The service can back up block volumes. Users control this via a manual or scheduled API call. Users schedule and recover backups through a web interface.

75.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

75.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/rds/index.html
- **Service quotas:** https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/CHAP_Limits.html
- **Service FAQs:** https://aws.amazon.com/rds/faqs/

75.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/rds/index.html and the following links for comprehensive technical documentation regarding this service.

- Amazon RDS User Guide: Describes all Amazon RDS concepts and provides instructions on using the various features with both the console and the command line interface.
- Amazon RDS in the AWS CLI Reference: Describes all the CLI commands for Amazon RDS in detail. Provides all syntax. Also provides examples for the most common commands.

- Amazon RDS API Reference: Describes all the API operations for Amazon RDS in detail. Also provides sample requests, responses, and errors for the supported web services protocols.

76. Amazon Route 53

76.1. Service Overview

Amazon Route 53 is a highly available and scalable cloud Domain Name System (DNS) web service. It is designed to give developers and businesses an extremely reliable and cost effective way to route end users to Internet applications by translating names like www.example.com into the numeric IP addresses like 192.0.2.1 that computers use to connect to each other. Amazon Route 53 is fully compliant with IPv6 as well.

Amazon Route 53 effectively connects user requests to infrastructure running in AWS – such as Amazon EC2 instances, Elastic Load Balancing load balancers, or Amazon S3 buckets – and can also be used to route users to infrastructure outside of AWS. You can use Amazon Route 53 to configure DNS health checks, then continuously monitor your applications' ability to recover from failures and control application recovery with Route 53 Application Recovery Controller.

Amazon Route 53 Traffic Flow makes it easy for you to manage traffic globally through a variety of routing types, including Latency Based Routing, Geo DNS, Geoproximity, and Weighted Round Robin—all of which can be combined with DNS Failover in order to enable a variety of low-latency, fault-tolerant architectures. Using Amazon Route 53 Traffic Flow's simple visual editor, you can easily manage how your end-users are routed to your application's endpoints—whether in a single AWS region or distributed around the globe. Amazon Route 53 also offers Domain Name Registration – you can purchase and manage domain names such as example.com and Amazon Route 53 will automatically configure DNS settings for your domains.

76.1.1. Features

- **Route 53 Resolver**: Get recursive DNS for your Amazon VPC and on-premises networks. Create conditional forwarding rules and DNS endpoints to resolve custom names mastered in Amazon Route 53 private hosted zones or in your on-premises DNS servers.
- **Route 53 Resolver DNS Firewall:** Protect your recursive DNS queries within the Route 53 Resolver. Create domain lists and build firewall rules that filter outbound DNS traffic against these rules.
- **Route 53 Application Recovery Controller: Readiness Check:** Ensure that your resources across Availability Zones or Regions are continually audited for recovery readiness.
- **Route 53 Application Recovery Controller: Routing Control:** Use simple on/off switches, integrated with DNS records of your top-level resources, to failover traffic.
- **Route 53 Application Recovery Controller: Safety Rules:** Make sure that specific rules are followed during failover to protect automated recovery actions from impairing availability.
- **Traffic flow**: Easy-to-use and cost-effective global traffic management: route end users to the best endpoint for your application based on geoproximity, latency, health, and other considerations.

- **Latency based routing**: Route end users to the AWS region that provides the lowest possible latency.
- **Geo DNS**: Route end users to a particular endpoint that you specify based on the end user's geographic location.
- **Private DNS for Amazon VPC**: Manage custom domain names for your internal AWS resources without exposing DNS data to the public Internet.
- **DNS Failover**: Automatically route your website visitors to an alternate location to avoid site outages.

76.1.2. Benefits

- **Highly available and reliable:** Amazon Route 53 is built using AWS's highly available and reliable infrastructure. The distributed nature of our DNS servers helps ensure a consistent ability to route your end users to your application. Features such as Amazon Route 53 Traffic Flow and routing control help you improve reliability with easily-configured failover to reroute your users to an alternate location if your primary application endpoint becomes unavailable. Amazon Route 53 is designed to provide the level of dependability required by important applications. Amazon Route 53 is backed by the Amazon Route 53 Service Level Agreement.
- **Flexible:** Amazon Route 53 Traffic Flow routes traffic based on multiple criteria, such as endpoint health, geographic location, and latency. You can configure multiple traffic policies and decide which policies are active at any given time. You can create and edit traffic policies using the simple visual editor in the Route 53 console, AWS SDKs, or the Route 53 API. Traffic Flow's versioning feature maintains a history of changes to your traffic policies, so you can easily roll back to a previous version using the console or API.
- **Designed for use with other Amazon Web Services:** Amazon Route 53 is designed to work well with other AWS features and offerings. You can use Amazon Route 53 to map domain names to your Amazon EC2 instances, Amazon S3 buckets, Amazon CloudFront distributions, and other AWS resources. By using the AWS Identity and Access Management (IAM) service with Amazon Route 53, you get fine grained control over who can update your DNS data. You can use Amazon Route 53 to map your zone apex (example.com versus www.example.com) to your Elastic Load Balancing instance, Amazon CloudFront distribution, AWS Elastic Beanstalk environment, API Gateway, VPC endpoint, or Amazon S3 website bucket using a feature called Alias record.
- **Simple:** With self-service sign-up, Amazon Route 53 can start to answer your DNS queries within minutes. You can configure your DNS settings with the AWS Management Console or our easy-to-use API. You can also programmatically integrate the Amazon Route 53 API into your overall web application. For instance, you can use Amazon Route 53's API to create a new DNS record whenever you create a new EC2 instance. Amazon Route 53 Traffic Flow makes it easy to set up sophisticated routing logic for your applications by using the simple visual policy editor.
- **Fast:** Using a global anycast network of DNS servers around the world, Amazon Route 53 is designed to automatically route your users to the optimal location depending on network conditions. As a result, the service offers low query latency for your end users, as well as low update latency for your DNS record management needs. Amazon Route 53 Traffic Flow lets you further improve your customers' experience by running your application in multiple locations around the world and using traffic policies to ensure your end users are routed to the closest healthy endpoint for your application.

- **Cost-effective:** Amazon Route 53 passes on the benefits of AWS's scale to you. You pay only for the resources you use, such as the number of queries that the service answers for each of your domains, hosted zones for managing domains through the service, and optional features such as traffic policies and health checks, all at a low cost and without minimum usage commitments or any up-front fees.
- **Secure:** By integrating Amazon Route 53 with AWS Identity and Access Management (IAM), you can grant unique credentials and manage permissions for every user within your AWS account and specify who has access to which parts of the Amazon Route 53 service. When you enable Amazon Route 53 Resolver DNS firewall, you can configure it to inspect outbound DNS requests against a list of known malicious domains.
- **Scalable:** Route 53 is designed to automatically scale to handle very large query volumes without any intervention from you.
- **Simplify the hybrid cloud:** Amazon Route 53 Resolver provides recursive DNS for your Amazon VPC and on-premises networks over AWS Direct Connect or AWS Managed VPN.

76.2. Backup/Restore and Disaster Recovery

This requirement is not applicable for this service. For additional information beyond what is described herein, please refer to http://aws.amazon.com/documentation/.

76.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

76.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/route53/index.html
- **Service quotas:** https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/DNSLimitations.html
- **Service FAQs:** https://aws.amazon.com/route53/faqs/

76.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/route53/index.html and the following links for comprehensive technical documentation regarding this service.

- Developer Guide: Provides an overview of Amazon Route 53, detailed feature descriptions, and procedures for using the console.
- API Reference: Describes all the API operations for Amazon Route 53 in detail. Also provides sample requests, responses, and errors for the supported web services protocols.

77. Amazon SageMaker

77.1. Service Overview

Amazon SageMaker is a fully managed service that provides every developer and data scientist with the ability to prepare build, train, and deploy machine learning (ML) models quickly. SageMaker removes the heavy lifting from each step of the machine learning process to make it

easier to develop high quality models. SageMaker provides all of the components used for machine learning in a single toolset so models get to production faster with much less effort and at lower cost.

77.1.1. Features

- **Prepare Data in Minutes:** Using Amazon SageMaker Data Wrangler, you can quickly and easily prepare data and create model features. You can connect to data sources and use built-in data transformations to engineer model features.
- **Transparency:** Amazon SageMaker Clarify provides data to improve model quality through bias detection during data preparation and after training. SageMaker Clarify also provides model explainability reports so stakeholders can see how and why models make predictions.
- **Security and Privacy:** Amazon SageMaker allows you to operate on a fully secure ML environment from day one. You can use a comprehensive set of security features to help support a broad range of industry regulations.
- **Data Labeling:** Amazon SageMaker Ground Truth makes it easy to build highly accurate training datasets for machine learning. Get started with labeling your data in minutes through the SageMaker Ground Truth console using custom or built-in data labeling workflows including 3D point clouds, video, images, and text.
- **Data Processing at Scale:** Amazon SageMaker Processing extends the ease, scalability, and reliability of SageMaker to running data processing workloads. SageMaker Processing allows you to connect to existing storage, spin up the resources required to run your job, save the output to persistent storage, and provides logs and metrics.
- **One-click Jupyter Notebooks:** Amazon SageMaker Studio Notebooks are one-click Jupyter notebooks and the underlying compute resources are fully elastic, so you can easily dial up or down the available resources. Notebooks are shared with a single click so colleagues get the same notebook, saved in the same place.
- **Auto ML:** Amazon SageMaker Autopilot automatically builds, trains, and tunes the best machine learning models, based on your data while allowing to maintain full control and visibility. You then can directly deploy the model to production with just one click, or iterate to improve the model quality.
- **Reinforcement Learning:** Amazon SageMaker supports reinforcement learning in addition to traditional supervised and unsupervised learning. SageMaker has built-in, fully-managed reinforcement learning algorithms, including some of the newest and best performing in the academic literature.
- **Experiment Management and Tracking:** Amazon SageMaker Experiments helps you track iterations to ML models by capturing the input parameters, configurations, and results, and storing them as ‘experiments’. In SageMaker Studio you can browse active experiments, search for previous experiments, review previous experiments with their results, and compare experiment results.

77.1.2. Benefits

- **Make ML more accessible:** Enable more people to innovate with ML through a choice of tools—integrated development environments for data scientists and no-code visual interfaces for business analysts.
- **Prepare data at scale:** Access, label, and process large amounts of structured data (tabular data) and unstructured data (photos, video, and audio) for ML.

- **Accelerate ML development:** Reduce training time from hours to minutes with optimized infrastructure. Boost team productivity up to 10 times with purpose-built tools.
- **Streamline the ML lifecycle:** Automate and standardize MLOps practices across your organization to build, train, deploy, and manage models at scale.
- **High-performance, low-cost ML at scale:** Amazon SageMaker is built on Amazon's two decades of experience developing real-world machine learning applications, including product recommendations, personalization, intelligent shopping, robotics, and voice-assisted devices.

77.2. Backup/Restore and Disaster Recovery

This requirement is not applicable for this service. For additional information beyond what is described herein, please refer to http://aws.amazon.com/documentation/.

77.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

77.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/sagemaker/index.html
- **Service quotas:** https://docs.aws.amazon.com/general/latest/gr/sagemaker.html
- **Service FAQs:** https://aws.amazon.com/sagemaker/faqs/?nc=sn&loc=4

77.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/sagemaker/index.html and the following links for comprehensive technical documentation regarding this service.

- Developer Guide: Provides a conceptual overview of Amazon SageMaker and offers step-by-step instructions for building, training, and deploying models.
- API Reference: Describes all the API operations for Amazon SageMaker in detail.
- Amazon SageMaker Python SDK: Use the SageMaker Python SDK library to train and deploy models using popular deep learning frameworks and algorithms.
- AWS SDK for Python (Boto 3): Use the AWS SDK for Python (Boto 3) to format model data and build applications to build, train, and deploy machine learning models.

78. Amazon Simple Email Service (SES)

78.1. Service Overview

Amazon Simple Email Service (SES) is a cost-effective, flexible, and scalable email service that enables developers to send mail from within any application. You can configure Amazon SES quickly to support several email use cases, including transactional, marketing, or mass email communications. Amazon SES's flexible IP deployment and email authentication options help drive higher deliverability and protect sender reputation, while sending analytics measure the impact of each email. With Amazon SES, you can send email securely, globally, and at scale.

78.1.1. Features

- **Sender Configuration Options:** Amazon SES offers several methods of sending email, including the Amazon SES console, the Simple Mail Transfer Protocol (SMTP) interface, and the Amazon SES API. You can access the API using the AWS Command Line Interface (AWS CLI), or by using an AWS Software Development Kit (SDK).
- **Shared IP Addresses:** By default, Amazon SES sends email from IP addresses that are shared with other Amazon SES customers. Shared addresses are a great option for many customers who want to start sending immediately with established IPs. They are included in the base Amazon SES pricing, and their reputations are carefully monitored to ensure high deliverability.
- **Dedicated IP Addresses:** For customers that want to manage their own IP reputation, you can lease dedicated IP addresses to use with your Amazon SES account. You can also use the dedicated IP pools feature to create pools of those IP addresses. Customers can either send all traffic from these dedicated IPs or use configuration sets to align specific use cases to specific IPs.
- **Owned IP Addresses:** Amazon SES also supports Bring Your Own IP (BYOIP). This feature lets you use a range of IP addresses that you already own to send email with Amazon SES. This makes leveraging current investments and migrating from other email service providers easy.
- **Sender Identity Management and Security:** Amazon SES supports all industry-standard authentication mechanisms, including Domain Keys Identified Mail (DKIM), Sender Policy Framework (SPF), and Domain-based Message Authentication, Reporting and Conformance (DMARC). When an internet service provider (ISP) receives an email, they check to see if it is authenticated before attempting to deliver it to the recipient. Authentication demonstrates to the ISP that you own the email address you are sending from. Amazon SES also enables customers to connect an Amazon SES SMTP endpoint to a virtual private cloud (VPC) through a VPC endpoint powered by AWS PrivateLink. With this feature, customers can access the Amazon SES SMTP endpoint securely without requiring an Internet Gateway in a VPC.
- **Sending Statistics:** Amazon SES provides a few methods for monitoring your email sending activity, helping you fine-tune your email sending strategy. Amazon SES can capture information about the entire email response funnel, including the numbers of sends, deliveries, opens, clicks, bounces, complaints, and rejections. This data is shared by default in the Sending Statistics report in the Amazon SES console. Use the Global suppression list to remove bounced emails from your sending list, or configure your own account-level suppression list. Sending data can be stored in an Amazon S3 bucket or an Amazon Redshift database, sent to Amazon SNS for real-time notifications, or analyzed using Amazon Kinesis Analytics.
- **Reputation Dashboard:** The Amazon SES console includes a reputation dashboard that you can use to track issues that could impact the delivery of your emails. This dashboard tracks the overall bounce and feedback loops for your account, and can inform you when other deliverability-impacting events occur, such as spamtrap hits, references to blocked domains in your emails, and reports from reputable anti-spam organizations. Amazon SES automatically publishes the bounce and complaint metrics from this dashboard to Amazon CloudWatch. You can use CloudWatch to create alarms that notify you when your bounce or complaint rates reach certain thresholds. With this information, you can take immediate action on issues that could impact your sender reputation.

- **Deliverability Dashboard:** The Deliverability Dashboard (via the SES API v2) helps you understand and remediate issues that could impact the delivery of your emails, such as suboptimal email content, and attempting to email users who have unsubscribed or bounced in the past.
- **Email Receiving:** When you use Amazon SES to receive incoming emails, you have complete control over which emails you accept, and what to do with them after you receive them. You can accept or reject mail based on the email address, IP address, or domain of the sender. Once Amazon SES has accepted the email, you can store it in an Amazon S3 bucket, execute custom code using an AWS Lambda function, or publish notifications to Amazon SNS.
- **Mailbox Simulator:** The Amazon SES mailbox simulator makes it easy to test how your application handles certain scenarios, such as bounces or complaints, without impacting your sender reputation. Using the mailbox simulator is as easy as sending a test email to a specific address. You can use the mailbox simulator to simulate successful deliveries, hard bounces, out-of-office responses or feedback.

78.1.2. Benefits

- **Integrate quickly:** Using either the Amazon SES console, APIs, or SMTP, you can configure email sending in minutes. Amazon SES also supports email receiving, enabling you to interact with your customers at scale. Regardless of use case or sending volume, you only pay for what you use with Amazon SES.
- **Optimize your deliverability:** Use the reputation dashboard, which includes account performance insights and anti-spam feedback, to maximize your deliverability. You have flexible deployment options that range from shared, dedicated, and customer-owned IPs that helps you influence your sending reputation. Amazon SES has relationships with experts like M3AAWG to improve delivery to your customers through industry best practices.
- **Send messages efficiently:** Email sending statistics, including email deliveries, bounces, and feedback loop results, help you to measure the effectiveness of each email outreach. Additional insights like email open or click-through rates measure how engaged your customers are in your email communications.
- **Scale securely:** Amazon SES authentication options such as Sender Policy Framework (SPF) and DomainKeys Identified Mail (DKIM) confirms your right to send on behalf of your domain. Virtual private cloud (VPC) support makes email sending from any application secure. Amazon SES is globally available with HIPAA eligibility, in-region compliance (C5, IRAP) and global certifications (Fed-Ramp, ISO, GDPR).

78.2. Backup/Restore and Disaster Recovery

This service allows for backup and recovery. The service can back up emails to S3. Users control this via a manual or scheduled API call. Users schedule and recover backups through a web interface.

78.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

78.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/ses/index.html
- **Service quotas:** https://docs.aws.amazon.com/ses/latest/dg/quotas.html
- **Service FAQs:** https://aws.amazon.com/ses/faqs/

78.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/ses/index.html and the following links for comprehensive technical documentation regarding this service.

- Developer Guide: Learn about Amazon SES and how to use its features through the redesigned new console
- API Reference: Get syntax and examples for the Amazon SES API.
- Amazon SES in the AWS CLI Reference: Describes the AWS CLI commands for Amazon SES API.
- API v2 Reference: Get syntax and examples for version 2 of the Amazon SES API.
- Amazon SES v2 in the AWS CLI Reference: Describes the AWS CLI commands for version 2 of the Amazon SES API.

79. Amazon Simple Notification Service (SNS)

79.1. Service Overview

Amazon Simple Notification Service (Amazon SNS) is a fully managed messaging service for both application-to-application (A2A) and application-to-person (A2P) communication.

The A2A pub/sub functionality provides topics for high-throughput, push-based, many-to-many messaging between distributed systems, microservices, and event-driven serverless applications. Using Amazon SNS topics, your publisher systems can fanout messages to a large number of subscriber systems, including Amazon SQS queues, AWS Lambda functions, HTTPS endpoints, and Amazon Kinesis Data Firehose, for parallel processing. The A2P functionality enables you to send messages to users at scale via SMS, mobile push, and email.

79.1.1. Features

- **Standard Topics**: Standard topics can be used in many scenarios, as long as your application can process messages that arrive more than once and out of order, for example: fanning out messages to media encoding, fraud detection, tax calculation, search index, and critical alerting applications.
- **FIFO Topics:** FIFO topics are designed to enhance messaging between applications when the order of operations and events is critical, or where duplicates can't be tolerated, for example: fanning out messages to bank transaction logging, stock monitoring, flight tracking, inventory management, and price update applications.
- **Event sources and destinations:** Event-driven computing is a model in which subscriber services automatically perform work in response to events triggered by publisher services. This paradigm can be applied to automate workflows while decoupling the services that collectively and independently work to fulfill these workflows. Amazon SNS is an event-driven hub that has native integration with a wide variety of AWS event sources and event destinations.
- **Message publishing and batching:** Message publishing enables you to send data, in the form of messages, to an Amazon SNS topic which delivers the messages asynchronously to the applications that are subscribed to the topic. You can publish from

1 to 10 messages per API request. You may choose to batch messages together to reduce your Amazon SNS costs. Each message can contain up to 256KB of data. If your use case requires larger data payloads, the Amazon SNS Extended Client Library stores the payload (up to 2GB) in an Amazon S3 bucket and publishes the reference of the stored Amazon S3 object to the Amazon SNS topic.

- **Message filtering:** Message filtering empowers your subscriber applications to create filter policies, so that these applications can receive only the notifications that they are interested in, as opposed to receiving every message published to the topic. This enables you to simplify your architecture, offloading the message filtering logic from subscriber applications as well as the message routing logic from publisher applications.
- **Message fanout and delivery:** When you publish a message to a topic, Amazon SNS replicates and delivers the message to applications subscribed to the topic. Amazon SNS supports application-to-application (A2A) and application-to-person (A2P) message delivery. Amazon SNS also supports cross-region and cross-account message delivery, in addition to message delivery status logging with Amazon CloudWatch.
- **Message durability:** Amazon SNS uses a number of mechanisms that work together to provide message durability. To start, published messages are stored across multiple, geographically-separated servers and data centers. If a subscribed endpoint isn't available, Amazon SNS executes a message delivery retry policy. To preserve any messages that aren't delivered before the delivery retry policy ends, you can use a dead-letter queue powered by Amazon SQS. Moreover, you can subscribe Amazon Kinesis Data Firehose delivery streams to Amazon SNS topics, which allows messages to be sent to durable endpoints such as Amazon S3 buckets or Amazon Redshift tables.
- **Message encryption:** Amazon SNS provides encrypted topics to protect your messages from unauthorized and anonymous access. When you publish messages to encrypted topics, Amazon SNS immediately encrypts your messages. The encryption takes place on the server, using a 256-bit AES-GCM algorithm and a Customer Master Key (CMK) issued with AWS Key Management Service (KMS). The messages are stored in encrypted form, and decrypted as they are delivered to subscribing endpoints, such as Amazon SQS queues, Amazon Kinesis Data Firehose streams, AWS Lambda functions, HTTP/S endpoints, phone numbers, mobile apps, and email addresses.
- **Message privacy:** Amazon SNS supports VPC Endpoints (VPCE) via AWS PrivateLink. You can use VPC Endpoints to privately publish messages to Amazon SNS topics, from an Amazon Virtual Private Cloud (VPC), without traversing the public internet. This feature brings additional security, helps promote data privacy, and aligns with assurance programs. When you use AWS PrivateLink, you don’t need to set up an Internet Gateway (IGW), Network Address Translation (NAT) device, or Virtual Private Network (VPN) connection. You don’t need to use public IP addresses, either.
- **Message archiving and analytics:** Amazon SNS provides a direct connection to Amazon Kinesis Data Firehose, allowing message storage in services such as Amazon S3, Amazon Redshift, Amazon OpenSearch Service, and MongoDB. This feature also enables message storage in analytics services, such as Datadog, New Relic, and Splunk.

79.1.2. Benefits

- **Simplify and reduce costs with message filtering and batching:** Amazon SNS helps you simplify your application architecture and reduce costs. With message batching,

publisher systems can send up to 10 messages in a single API request. With message filtering, subscriber systems receive only the messages that they are interested in.

- **Ensure accuracy with message ordering and deduplication:** Amazon SNS FIFO topics work with Amazon SQS FIFO queues to ensure messages are delivered in a strictly-ordered manner and are only processed once. This enables you to maintain accuracy and consistency when processing transactions across a single or multiple independent services.
- **Increase security with message encryption and privacy:** Amazon SNS provides encrypted topics to protect your messages from unauthorized access. The encryption uses a 256-bit AES-GCM algorithm and a customer master key (CMK) issued with AWS Key Management Service (KMS). Amazon SNS also supports VPC endpoints via AWS PrivateLink, so you can privately publish messages to Amazon SNS topics, from an Amazon Virtual Private Cloud (VPC) subnet, without traversing the Internet.
- **Increase durability with message archiving, delivery retries, and DLQ:** Amazon SNS stores each message published across geographically-separated data centers. If a subscribed system isn't available, Amazon SNS executes a message delivery retry policy. To preserve any messages that can't be delivered before the delivery retry policy ends, Amazon SNS can move them to dead-letter queues (DLQ). Amazon SNS can also archive messages in Amazon S3 via Amazon Kinesis Data Firehose subscriptions.
- **Capture and fan out events from AWS services:** Amazon SNS is an event-driven hub that has native integration with a wide variety of AWS event sources and destinations. Amazon SNS can capture and fan out events from more than 60 AWS services, from a number of AWS categories, such as Analytics, Compute, Containers, Databases, IoT, Machine Learning, Security, Serverless, and Storage.
- **Send A2P notifications via SMS, mobile push, and email:** Amazon SNS enables you to send notifications directly to your customers. Amazon SNS supports SMS text messaging to over 200 countries, mobile push notifications to Amazon, Android, Apple, Baidu, and Microsoft devices, and also email notifications. Amazon SNS provides redundancy across multiple SMS providers, and enables you to send mobile push notifications using a single API for all mobile platforms.

79.2. Backup/Restore and Disaster Recovery

This requirement is not applicable for this service. For additional information beyond what is described herein, please refer to http://aws.amazon.com/documentation/.

79.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

79.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/sns/
- **Service quotas:** https://docs.aws.amazon.com/general/latest/gr/sns.html
- **Service FAQs:** https://aws.amazon.com/sns/faqs/

79.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/sns/ and the following links for comprehensive technical documentation regarding this service.

- Developer Guide: Provides a conceptual overview of Amazon SNS and includes detailed development instructions for using the various features.
- API Reference: Describes all the API operations for Amazon SNS in detail. Also provides sample requests, responses, and errors for the supported web services protocols.

80. Amazon Simple Queue Service (SQS)

80.1. Service Overview

Amazon Simple Queue Service (SQS) is a fully managed message queuing service that enables you to decouple and scale microservices, distributed systems, and serverless applications. SQS eliminates the complexity and overhead associated with managing and operating message-oriented middleware, and empowers developers to focus on differentiating work. Using SQS, you can send, store, and receive messages between software components at any volume, without losing messages or requiring other services to be available. Get started with SQS in minutes using the AWS console, Command Line Interface or SDK of your choice, and three simple commands.

SQS offers two types of message queues. Standard queues offer maximum throughput, best-effort ordering, and at-least-once delivery. SQS FIFO queues are designed to guarantee that messages are processed exactly once, in the exact order that they are sent.

80.1.1. Features

- **Unlimited queues and messages:** Create unlimited Amazon SQS queues with an unlimited number of messages in any Region
- **Payload Size:** Message payloads can contain up to 256KB of text in any format. Each 64KB ‘chunk’ of payload is billed as 1 request. For example, a single API call with a 256KB payload will be billed as four requests. To send messages larger than 256KB, you can use the Amazon SQS Extended Client Library for Java, which uses Amazon Simple Storage Service (S3) to store the message payload. A reference to the message payload is sent using SQS.
- **Batches:** Send, receive, or delete messages in batches of up to 10 messages or 256KB. Batches cost the same amount as single messages, meaning SQS can be even more cost effective for customers that use batching.
- **Long polling:** Reduce extraneous polling to minimize cost while receiving new messages as quickly as possible. When your queue is empty, long-poll requests wait up to 20 seconds for the next message to arrive. Long poll requests cost the same amount as regular requests.
- **Retain messages in queues for up to 14 days.**
- **Send and read messages simultaneously.**
- **Message locking**: When a message is received, it becomes “locked” while being processed. This keeps other computers from processing the message simultaneously. If the message processing fails, the lock will expire and the message will be available again.

- **Queue sharing:** Securely share Amazon SQS queues anonymously or with specific AWS accounts. Queue sharing can also be restricted by IP address and time-of-day.
- **Server-side encryption (SSE):** Protect the contents of messages in Amazon SQS queues using keys managed in the AWS Key Management Service (AWS KMS). SSE encrypts messages as soon as Amazon SQS receives them. The messages are stored in encrypted form and Amazon SQS decrypts messages only when they are sent to an authorized consumer.
- **Dead Letter Queues (DLQ):** Handle messages that a consumer has not successfully processed with dead- letter queues (DLQs). When a message's maximum receive count is exceeded, Amazon SQS moves the message to the DLQ associated with the original queue. DLQs must be of the same type as the source queue (standard or FIFO). You can inspect the messages in DLQs to understand why your consumer has not successfully received them. Once you have remediated the issues, you can move the messages from the DLQ to their respective source queues.

80.1.2. Benefits

- **Eliminate administrative overhead:** AWS manages all ongoing operations and underlying infrastructure needed to provide a highly available and scalable message queuing service. With SQS, there is no upfront cost, no need to acquire, install, and configure messaging software, and no time-consuming build-out and maintenance of supporting infrastructure. SQS queues are dynamically created and scale automatically so you can build and grow applications quickly and efficiently.
- **Reliably deliver messages:** Use Amazon SQS to transmit any volume of data, at any level of throughput, without losing messages or requiring other services to be available. SQS lets you decouple application components so that they run and fail independently, increasing the overall fault tolerance of the system. Multiple copies of every message are stored redundantly across multiple availability zones so that they are available whenever needed.
- **Keep sensitive data secure:** You can use Amazon SQS to exchange sensitive data between applications using server-side encryption (SSE) to encrypt each message body. Amazon SQS SSE integration with AWS Key Management Service (KMS) allows you to centrally manage the keys that protect SQS messages along with keys that protect your other AWS resources. AWS KMS logs every use of your encryption keys to AWS CloudTrail to help meet your regulatory and compliance needs.
- **Scale elastically and cost-effectively:** Amazon SQS leverages the AWS cloud to dynamically scale based on demand. SQS scales elastically with your application so you don’t have to worry about capacity planning and pre-provisioning. There is no limit to the number of messages per queue, and standard queues provide nearly unlimited throughput. Costs are based on usage which provides significant cost saving versus the “always-on” model of self-managed messaging middleware.

80.2. Backup/Restore and Disaster Recovery

This requirement is not applicable for this service. For additional information beyond what is described herein, please refer to http://aws.amazon.com/documentation/.

80.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

80.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/sqs/index.html
- **Service quotas:** https://docs.aws.amazon.com/AWSSimpleQueueService/latest/SQSDeveloperGuide/sqs-quotas.html
- **Service FAQs:** https://aws.amazon.com/step-functions/faqs/

80.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/sqs/index.html and the following links for comprehensive technical documentation regarding this service.

- Developer Guide: Provides a conceptual overview of Amazon SQS and includes detailed development instructions for using the various features.
- API Reference: Describes all the API operations for Amazon SQS in detail. Also provides sample requests, responses, and errors for the supported web services protocols.
- Amazon SQS in the AWS CLI Reference: Describes the AWS CLI commands that you can use to work with queues.

81. Amazon Simple Storage Service (S3)

81.1. Service Overview

Amazon Simple Storage Service (Amazon S3) is an object storage service offering industry-leading scalability, data availability, security, and performance. Customers of all sizes and industries can store and protect any amount of data for virtually any use case, such as data lakes, cloud-native applications, and mobile apps. With cost-effective storage classes and easy-to-use management features, you can optimize costs, organize data, and configure fine-tuned access controls to meet specific business, organizational, and compliance requirements.

81.1.1. Features

- **High Performance:** Amazon S3 supports parallel requests, which means you can scale your S3 performance by the factor of your compute cluster, without making any customizations to your application. Performance scales per prefix, so you can use as many prefixes as you need in parallel to achieve the required throughput. There are no limits to the number of prefixes. Amazon S3 performance supports at least 3,500 requests per second to add data and 5,500 requests per second to retrieve data. Each S3 prefix can support these request rates, making it simple to increase performance significantly.
- **Consistent:** Amazon S3 delivers strong read-after-write consistency automatically for all applications, without changes to performance or availability, without sacrificing regional isolation for applications, and at no additional cost. With strong consistency, S3 simplifies the migration of on-premises analytics workloads by removing the need to make changes to applications, and reduces costs by removing the need for extra infrastructure to provide strong consistency.
- **Query Data:** Amazon S3 has a built-in feature and complementary services that query data without needing to copy and load it into a separate analytics platform or data

warehouse. This means you can run big data analytics directly on your data stored in Amazon S3. **S3 Select** is an S3 feature designed to increase query performance by up to 400%, and reduce querying costs as much as 80%. It works by retrieving a subset of an object's data (using simple SQL expressions) instead of the entire object, which can be up to 5 terabytes in size.

- **S3 Object Lambda:** With S3 Object Lambda you can add your own code to S3 GET requests to modify and process data as it is returned to an application. For the first time, you can use custom code to modify the data returned by standard S3 GET requests to filter rows, dynamically resize images, redact confidential data, and much more. Powered by AWS Lambda functions, your code runs on infrastructure that is fully managed by AWS, eliminating the need to create and store derivative copies of your data or to run expensive proxies, all with no changes required to applications.
- **Replication:** With S3 Replication, you can replicate objects (and their respective metadata and object tags) to one or more destination buckets into the same or different AWS Regions for reduced latency, compliance, security, disaster recovery, and other use cases. S3 Cross-Region Replication (CRR) can be configured to replicate from a source S3 bucket to one or more destination buckets in different AWS Regions. Amazon S3 Same-Region Replication (SRR) replicates objects between buckets in the same AWS Region. Amazon S3 Replication Time Control (S3 RTC) helps you meet compliance requirements for data replication by providing an SLA and visibility into replication times.
- **Storage Management:** With S3 bucket names, prefixes, object tags, and S3 Inventory, you have a range of ways to categorize and report on your data, and subsequently can configure other S3 features to take action. Whether you store thousands of objects or a billion, **S3 Batch Operations** makes it simple to manage your data in Amazon S3 at any scale. With S3 Batch Operations, you can copy objects between buckets, replace object tag sets, modify access controls, and restore archived objects from S3 Glacier Flexible Retrieval and S3 Glacier Deep Archive storage classes, with a single S3 API request or a few clicks in the S3 console. You can also use S3 Batch Operations to run AWS Lambda functions across your objects to execute custom business logic, such as processing data or transcoding image files.
- **Multi-region action points:** Amazon S3 Multi-Region Access Points accelerate performance by up to 60% when accessing data sets that are replicated across multiple AWS Regions. Based on AWS Global Accelerator, S3 Multi-Region Access Points consider factors like network congestion and the location of the requesting application to dynamically route your requests over the AWS network to the lowest latency copy of your data. S3 Multi-Region Access Points provide a single global endpoint that you can use to access a replicated data set, spanning multiple buckets in S3. This allows you to build multi-region applications with the same simple architecture that you would use in a single region, and then to run those applications anywhere in the world.
- **S3 Object Lock:** You can enforce write-once-read-many (WORM) policies with S3 Object Lock. This S3 management feature blocks object version deletion during a customer-defined retention period so that you can enforce retention policies as an added layer of data protection or to meet compliance obligations. You can migrate workloads from existing WORM systems into Amazon S3, and configure S3 Object Lock at the object- and bucket-levels to prevent object version deletions prior to a pre-defined Retain Until Date or Legal Hold Date. Objects with S3 Object Lock retain WORM protection, even if they are moved to different storage classes with an S3 Lifecycle policy. To track

what objects have S3 Object Lock, you can refer to an S3 Inventory report that includes the WORM status of objects. S3 Object Lock can be configured in one of two modes.

- **Multiple storage classes:** With Amazon S3, you can store data across a range of different S3 storage classes purpose-built for specific use cases and access patterns: S3 Intelligent-Tiering, S3 Standard, S3 Standard-Infrequent Access (S3 Standard-IA), S3 One Zone-Infrequent Access (S3 One Zone-IA), S3 Glacier Instant Retrieval, S3 Glacier Flexible Retrieval, S3 Glacier Deep Archive, and S3 Outposts. Every S3 storage class supports a specific data access level at corresponding costs or geographic location.
- **S3 Storage Class Analysis:** Amazon S3 Storage Class Analysis analyzes storage access patterns to help you decide when to transition the right data to the right storage class. This Amazon S3 feature observes data access patterns to help you determine when to transition less frequently accessed storage to a lower-cost storage class. You can use the results to help improve your S3 Lifecycle policies. You can configure storage class analysis to analyze all the objects in a bucket. Or, you can configure filters to group objects together for analysis by common prefix, by object tags, or by both prefix and tags.

81.1.2. Benefits

- **Durability:** Scale storage resources to meet fluctuating needs with 99.99999999999% (11 9s) of data durability.
- **Multiple storage classes:** Store data across Amazon S3 storage classes to reduce costs without upfront investment or hardware refresh cycles.
- **Highly secure:** Protect your data with unmatched security, compliance and audit capabilities.
- **Easy to manage:** Easily manage data at any scale with robust access controls, flexible replication tools, and organization-wide visibility.
- **Build a data lake:** Run big data analytics, artificial intelligence (AI), machine learning (ML), and high performance computing (HPC) applications to unlock data insights.
- **Back up and restore critical data:** Meet Recovery Time Objectives (RTO), Recovery Point Objectives (RPO), and compliance requirements with S3's robust replication features.
- **Archive data at the lowest cost:** Move on-premises archives to the low-cost Amazon S3 Glacier and Amazon S3 Glacier Deep Archive storage classes to eliminate operational complexities.
- **Run cloud-native applications:** Build fast, powerful mobile and web-based cloud-native apps that scale automatically in a highly available configuration.

81.2. Backup/Restore and Disaster Recovery

This service allows for backup and recovery. The service can back up using cross-region replication, same-region replication, bucket versioning, and lifecycle rules. Users control this via a manual or scheduled API call. Users schedule and recover backups through a web interface.

81.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

81.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/s3/index.html
- **Service quotas:** https://docs.aws.amazon.com/general/latest/gr/s3.html
- **Service FAQs:** https://aws.amazon.com/s3/faqs/?nc=sn&loc=7

81.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/s3/index.html and the following links for comprehensive technical documentation regarding this service.

- User Guide: Provides detailed information and instructions for getting started, developing, and working with Amazon S3 using the AWS Management Console, AWS CLI, AWS SDKs, and REST API.
- API Reference: Describes all the Amazon S3 API operations in detail. Also provides sample requests, responses, and errors for the supported web services protocols.
- Developer Guide: Provides detailed information about setting up and working with Amazon S3 Glacier using the REST API and the AWS SDK for Java and AWS SDK for .NET.

82. Amazon Simple Workflow Service (SWF)

82.1. Service Overview

Amazon SWF helps developers build, run, and scale background jobs that have parallel or sequential steps. You can think of Amazon SWF as a fully-managed state tracker and task coordinator in the Cloud.

If your app's steps take more than 500 milliseconds to complete, you need to track the state of processing, and you need to recover or retry if a task fails, Amazon SWF can help you.

82.1.1. Features

- **Fully Managed Workflow:** Amazon SWF is a fully managed workflow service for building scalable, resilient applications.
- **Simple API Calls:** Amazon SWF provides simple API calls that can be executed from code written in any language and run on your EC2 instances, or any of your machines located anywhere in the world that can access the Internet.

82.1.2. Benefits

- **Logical Separation:** Amazon SWF promotes a separation between the control flow of your background job's stepwise logic and the actual units of work that contain your unique business logic. This allows you to separately manage, maintain, and scale "state machinery" of your application from the core business logic that differentiates it. As your business requirements change, you can easily change application logic without having to worry about the underlying state machinery, task dispatch, and flow control.
- **Reliable:** Amazon SWF runs within Amazon’s high-availability data centers, so the state tracking and task processing engine is available whenever applications need them. Amazon SWF redundantly stores the tasks, reliably dispatches them to application components, tracks their progress, and keeps their latest state.

- **Simple:** Amazon SWF replaces the complexity of custom-coded workflow solutions and process automation software with a fully managed cloud workflow web service. This eliminates the need for developers to manage the infrastructure plumbing of process automation so they can focus their energy on the unique functionality of their application.
- **Scalable:** Amazon SWF seamlessly scales with your application's usage. No manual administration of the workflow service is required as you add more cloud workflows to your application or increase the complexity of your workflows.
- **Flexible:** Amazon SWF lets you write your application components and coordination logic in any programming language and run them in the cloud or on-premises.

82.2. Backup/Restore and Disaster Recovery

This requirement is not applicable for this service. For additional information beyond what is described herein, please refer to http://aws.amazon.com/documentation/.

82.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

82.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/swf/index.html
- **Service quotas:** https://docs.aws.amazon.com/amazonswf/latest/developerguide/swf-dg-limits.html
- **Service FAQs:** https://aws.amazon.com/swf/faqs/

82.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/swf/index.html and the following links for comprehensive technical documentation regarding this service.

- Developer Guide: Provides a conceptual overview of Amazon SWF and includes detailed development instructions for using the various features.
- API Reference: Describes all the API operations for Amazon SWF in detail. Also provides sample requests, responses, and errors for the supported web services protocols.
- Developer Guide: Looks closely at the AWS Flow Framework for Java, which is a programming framework that enables you to build asynchronous and distributed applications with Amazon Simple Workflow Service using the features of Java.
- API Reference: Describes all the API operations for the AWS Flow Framework for Java in detail. Also provides sample requests, responses, and errors for the supported web services protocols.

83. Amazon Textract

83.1. Service Overview

Amazon Textract is a machine learning (ML) service that automatically extracts text, handwriting, and data from scanned documents. It goes beyond simple optical character recognition (OCR) to identify, understand, and extract data from forms and tables. Today, many

companies manually extract data from scanned documents such as PDFs, images, tables, and forms, or through simple OCR software that requires manual configuration (which often must be updated when the form changes). To overcome these manual and expensive processes, Textract uses ML to read and process any type of document, accurately extracting text, handwriting, tables, and other data with no manual effort. You can quickly automate document processing and act on the information extracted, whether you're automating loans processing or extracting information from invoices and receipts. Textract can extract the data in minutes instead of hours or days. Additionally, you can add human reviews with Amazon Augmented AI to provide oversight of your models and check sensitive data.

83.1.1. Features

- **Optical character recognition:** Amazon Textract uses optical character recognition (OCR) to automatically detect printed text, handwriting, and numbers in a scan or rendering of a document, such as a legal document or a scan of a book.
- **Form extraction:** You can detect key-value pairs in document images automatically and retain the context without manual intervention. A key-value pair is a set of linked data items. For instance, in a document, the field "First Name" is the key and "Jane" is the value. This makes it easy to import the extracted data into a database or provide it as a variable in an application. With traditional OCR solutions, keys and values are extracted as simple text, and their relationship is lost unless hard-coded rules are written and maintained for each form.
- **Table extraction:** Amazon Textract preserves the composition of data stored in tables during extraction. This is helpful for documents that are largely composed of structured data, such as financial reports or medical records with tables in columns and rows. You can automatically load the extracted data into a database using a predefined schema. For example, rows of item numbers and quantities in an inventory report will retain their association so an inventory management application can easily increment item totals.
- **Handwriting recognition:** Many documents, such as medical intake forms and employment applications, include both handwritten and printed text. Amazon Textract can extract both from documents written in English with high confidence scores, whether the text is free-form or embedded in tables. Documents can also contain a mix of typed text or handwritten text.
- **Invoices and receipts:** Invoices and receipts can have a wide variety of layouts, which makes it difficult and time-consuming to manually extract data at scale. Amazon Textract uses machine learning (ML) to understand the context of invoices and receipts and automatically extracts relevant data such as vendor name, invoice number, item prices, total amount, and payment terms.
- **Identity documents:** Amazon Textract uses machine learning (ML) to understand the context of identity documents such as U.S. passports and driver's licenses without the need for templates or configuration. You can automatically extract specific information such as date of expiry and date of birth, as well as intelligently identify and extract implied information such as name and address. Using Analyze ID, businesses providing ID verification services and those in finance, healthcare, and insurance can easily automate account creation, appointment scheduling, employment applications, and more by allowing customers to submit a picture or scan of their identity document.
- **Bounding boxes:** All extracted data is returned with bounding box coordinates—polygon frames that encompass each piece of identified data, such as a word, a line, a

table, or individual cells within a table. This helps you audit where a word or number came from in the source document and guides you when search results provide scans of original documents. For example, when searching medical records for patient history details, you can easily find the source document and take note for future searches.

- **Adjustable confidence thresholds:** When extracting information from documents, Amazon Textract returns a confidence score for everything it identifies so you can make informed decisions about how to use the results. For instance, if you extract information from tax records and want to ensure high accuracy, you can flag any item with a confidence score below 95% to be reviewed by a human. You can set a lower threshold for other documents where errors would have fewer negative consequences, such as when processing resumes or digitizing archived records.
- **Human review workflow:** Amazon Textract is directly integrated with Amazon Augmented AI (A2I) so you can easily implement human review of printed text and handwriting extracted from documents. Many text-extraction applications require humans to review low-confidence predictions to ensure the results are correct, but building human review systems can be time-consuming and expensive.
- **Amazon Textract pricing:** Amazon Textract is a machine learning (ML) service that uses optical character recognition (OCR) to automatically extract text, handwriting, and data from scanned documents such as PDFs. With Amazon Textract, you pay only for what you use. There are no minimum fees and no upfront commitments. Amazon Textract charges only for pages processed whether you extract text, text with tables, or form data.

83.1.2. Benefits

- **Machine learning:** Extract text and structured data such as tables and forms from documents using artificial intelligence (AI)—no configuration or templates necessary.
- **Information extraction:** Go beyond simple optical character recognition (OCR) by extracting relationships, structure, and text from documents.
- **Security:** Improve security and compliance through robust data privacy, encryption, security controls, and support compliance standards such as HIPAA, GDPR, and more.
- **Human reviews:** Easily implement human reviews with Amazon Augmented AI (A2I) to manage nuanced or sensitive workflows and audit predictions.

83.2. Backup/Restore and Disaster Recovery

This requirement is not applicable for this service. For additional information beyond what is described herein, please refer to http://aws.amazon.com/documentation/.

83.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

83.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/textract/
- **Service quotas:** https://docs.aws.amazon.com/general/latest/gr/textract.html

- **Service FAQs:** https://aws.amazon.com/textract/faqs/

83.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/textract/ and the following links for comprehensive technical documentation regarding this service.

- Developer Guide: Use Amazon Textract to detect and analyse text in your documents.

84. Amazon Timestream

84.1. Service Overview

Amazon Timestream is a fast, scalable, and serverless time series database service for IoT and operational applications that makes it easy to store and analyse trillions of events per day up to 1,000 times faster and at as little as 1/10th the cost of relational databases. Amazon Timestream saves you time and cost in managing the lifecycle of time series data by keeping recent data in memory and moving historical data to a cost optimized storage tier based upon user defined policies. Amazon Timestream's purpose-built query engine lets you access and analyse recent and historical data together, without needing to specify explicitly in the query whether the data resides in the in-memory or cost-optimized tier. Amazon Timestream has built-in time series analytics functions, helping you identify trends and patterns in your data in near real-time. Amazon Timestream is serverless and automatically scales up or down to adjust capacity and performance, so you don't need to manage the underlying infrastructure, freeing you to focus on building your applications.

84.1.1. Features

- **Serverless auto-scaling architecture:** Amazon Timestream features a fully decoupled architecture where data ingestion, storage, and query can scale independently, allowing it to offer virtually infinite scale for an application's needs. With Amazon Timestream, you don't need to manage infrastructure or provision capacity. Data ingest and query auto-scale based on your workload.
- **Data storage tiering:** Amazon Timestream simplifies your data lifecycle management with a memory store for recent data and a magnetic store for historical data. The memory store is optimized for fast point in time queries, and the magnetic store is optimized for fast analytic queries. With Amazon Timestream, you don't need to configure, monitor, and manage a complex data archival process. You can simply configure data retention policies to automatically move data from the memory store to the magnetic store, and to delete it from the magnetic store when it reaches a certain age.
- **Purpose-built adaptive query engine:** With Amazon Timestream, you don't need to use disparate tools for data access. Amazon Timestream's purpose-built adaptive query engine allows you to access data across storage tiers using a single SQL statement. It transparently accesses and combines data across storage tiers without requiring you to specify the data location.
- **Built-in time series analytics:** Amazon Timestream supports time series analytics and defines time series as a native data type. It supports advanced aggregates, window functions, and complex data types such as arrays and rows.
- **Always-encrypted data:** All data in Amazon Timestream is automatically encrypted, so you don't need to manually encrypt data at rest or in transit. Amazon Timestream also

enables you to specify an AWS KMS customer managed key (CMK) for encrypting data in the magnetic store.

- **Integrates with popular data collection, visualization, and machine learning tools:** Amazon Timestream integrates with commonly used services for data collection, visualization, and machine learning. You can send data to Amazon Timestream using AWS IoT Core, Amazon Kinesis, Amazon MSK, and open source Telegraf. You can visualize data using Amazon QuickSight, Grafana, and business intelligence tools through JDBC. You can also use Amazon SageMaker with Amazon Timestream for machine learning.
- **Performant and cost-effective time-series analytics:** Amazon Timestream's scheduled queries offer a fully managed, serverless, and scalable solution for calculating and storing aggregates, rollups, and other real-time analytics used to power frequently accessed operational dashboards, business reports, applications, and device-monitoring systems.

84.1.2. Benefits

- **High performance at low cost:** Amazon Timestream is designed to enable interactive and affordable real-time analytics, with up to 1,000 times faster query performance, and as little as 1/10th the cost of relational databases. With product features such as scheduled queries, multi-measure records, and data storage tiering, you can process, store, and analyse your time-series data at a fraction of the cost of existing time-series solutions. Amazon Timestream can help you derive faster and more affordable insights from your data so you can continue to make more data driven business decisions.
- **Serverless with auto-scaling:** Amazon Timestream is serverless – there are no servers to manage and no capacity to provision, so you can focus on building your applications. Amazon Timestream gives you the scale to process trillions of events and millions of queries per day. As your application needs change, it automatically scales to adjust capacity.
- **Data lifecycle management:** Amazon Timestream simplifies the complex process of data lifecycle management. It offers storage tiering, with a memory store for recent data and a magnetic store for historical data. Amazon Timestream automates the transfer of data from the memory store to the magnetic store based upon user configurable policies.
- **Simplified data access:** With Amazon Timestream, you no longer need to use disparate tools to access recent and historical data. Amazon Timestream's purpose-built query engine transparently accesses and combines data across storage tiers without you having to specify the data location.
- **Purpose-built for time series:** You can quickly analyse time series data using SQL, with built-in time series functions for smoothing, approximation, and interpolation. Amazon Timestream also supports advanced aggregates, window functions, and complex data types such as arrays and rows.
- **Always encrypted:** Amazon Timestream ensures that your time series data is always encrypted, whether at rest or in transit. Amazon Timestream also enables you to specify an AWS KMS customer managed key (CMK) for encrypting data in the magnetic store.

84.2. Backup/Restore and Disaster Recovery

This requirement is not applicable for this service. For additional information beyond what is described herein, please refer to http://aws.amazon.com/documentation/.

84.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

84.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/timestream/
- **Service quotas:** https://docs.aws.amazon.com/general/latest/gr/timestream.html
- **Service FAQs:** https://aws.amazon.com/timestream/faq/

84.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/timestream/ and the following links for comprehensive technical documentation regarding this service.

- Developer Guide: Provides a conceptual overview of Amazon Timestream, includes detailed instructions for using the various features, and provides a complete API reference for developers.

85. Amazon Transcribe

85.1. Service Overview

Amazon Transcribe is an automatic speech recognition service that makes it easy to add speech to text capabilities to any application. Transcribe's features enable you to ingest audio input, produce easy to read and review transcripts, improve accuracy with customization, and filter content to ensure customer privacy.

85.1.1. Features

- **Audio inputs:** Transcribe is designed to process live and recorded audio or video input to provide high quality transcriptions for search and analysis. We also offer separate APIs that uniquely understand customer calls (Amazon Transcribe Call Analytics) and medical conversations (Amazon Transcribe Medical).
- **Streaming & batch transcription:** You can process your existing audio recordings or stream the audio for real-time transcription. Using a secure connection, you can send a live audio stream to the service, and receive a stream of text in response.
- **Domain specific models:** Select a model that is tuned to telephone calls or multimedia video content. For example, Transcribe adapts to low-fidelity phone audio common in contact centers.
- **Automatic language identification:** With Amazon Transcribe, you can automatically identify the dominant language in an audio file and generate transcriptions. This is useful when your media library contains audio files in different languages. You can also use this feature for media content classification and verify that the main spoken language in your videos and podcasts is correctly labeled.

- **Easy to read transcripts:** Amazon Transcribe enables you produce accurate transcripts that are easy to read, review, and integrate into your specific applications. We work to make the output ready for downstream activities such as call transcript analysis, subtitling, and content search.
- **Punctuation & number normalization:** Amazon Transcribe automatically adds punctuation and number formatting, so that the output closely matches the quality of manual transcription at a fraction of the time and expense. Numbers are also transcribed into digits or “normal form” instead of words.
- **Timestamp generation:** Amazon Transcribe returns a timestamp for each word, so that you can easily find a word or phrase in the original recording or add subtitles to video.
- **Recognize multiple speakers:** Speaker changes are automatically recognized and attributed in the text to capture scenarios like telephone calls, meetings, and television shows accurately. To learn more about speaker identification.
- **Channel identification:** Contact centers can submit a single audio file to Amazon Transcribe, and the service will identify produce a single transcript annotated by channel labels automatically.
- **Customize your output:** Accuracy is critical and we provide you many options to customize transcripts to your specific business needs and vernacular. Transcribe also provides up to 10 alternative transcriptions for each sentence, so you can quickly choose the best option that applies to your content and domain. This is useful for human in-the-loop subtitling workflows.
- **Custom vocabulary:** With custom vocabulary, you can add new words to the base vocabulary to generate more accurate transcriptions for domain-specific words and phrases like product names, technical terminology, or names of individuals.
- **Custom language models:** When needed, you can build and train your own custom language model (CLM) for your use case and domain by submitting a corpus of text data to Amazon Transcribe. CLM is a suitable feature for enhancing speech recognition accuracy with your own data.
- **User safety & privacy features:** Ensuring customer privacy and safety is critical. When needed, Transcribe enables you to mask or remove words that are sensitive or unsuitable for your audience from transcription results.
- **Vocabulary filtering:** You can specify a list of words to remove from transcripts with vocabulary filtering. For example, you can specify a list of profane or offensive words and Amazon Transcribe removes them from transcripts automatically.
- **Automatic content redaction / PII redaction:** When instructed, Amazon Transcribe can help customers identify and redact sensitive personally identifiable information (PII) from the supported language transcripts. This allows contact centers to easily review and share the transcripts for customer experience insight and agent training.
- **Data Protection:** Secure data at rest using Amazon S3 key (SSE-S3) or specify your own AWS Key Management Service key. Amazon Transcribe uses TLS (Transport Layer Security) 1.2, a cryptographic protocol that enables authenticated connections and secure data transport over the internet via HTTP, with AWS certificates to encrypt data in transit. This includes streaming transcriptions.

- **Amazon Transcribe Call Analytics:** Extract conversation insights like call sentiment and speech loudness to improve agent productivity and customer experience with Amazon Transcribe Call Analytics.
- **Extract detailed call analytics & conversation insights:** Using the power of machine learning, you can quickly apply speech-to-text and natural language processing capabilities to uncover valuable conversation insights. You can then integrate insights such as customer and agent sentiment, detected issues, and speech characteristics like non-talk time, interruptions, and talk-speed into your inbound and outbound call analytics applications. This can help your supervisors more readily identify potential customer issues, agent coaching opportunities, and call trends.
- **Improve compliance & monitoring with automated call categorization:** Monitor your calls at scale to track compliance with company policies or regulatory requirements. Build and train your own custom categories based on your specified criteria (e.g. words/phrases or conversation characteristics). For example, you can setup category labels to see what percentage of calls are upsells or account cancellation.
- **Produce rich call transcripts:** Give your agents access to the conversation details from past interactions. The turn-by-turn transcripts provide insights such as customer sentiment, detected issues and interruptions.
- **Protect sensitive customer data:** Conversations often contain sensitive customer data such as names, addresses, credit card numbers, and social security numbers. Transcribe Call Analytics helps customers identify and redact this information from both the audio and text.

85.1.2. Benefits

- **Business insights:** Extract key business insights from customer calls, video files, clinical conversations, and more.
- **Business outcomes:** Improve business outcomes with state of the art speech recognition models that are fully managed and continuously trained.
- **Accuracy:** Enhance accuracy with custom models that understand your domain specific vocabulary.
- **Privacy:** Ensure customer privacy and safety by masking sensitive information.

85.2. Backup/Restore and Disaster Recovery

This requirement is not applicable for this service. For additional information beyond what is described herein, please refer to http://aws.amazon.com/documentation/.

85.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

85.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/transcribe/
- **Service quotas:** https://docs.aws.amazon.com/transcribe/latest/dg/limits-guidelines.html

- **Service FAQs:** https://aws.amazon.com/transcribe/faqs/

85.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/transcribe/ and the following links for comprehensive technical documentation regarding this service.

- Developer Guide: Provides a conceptual overview of Amazon Transcribe, offers detailed instructions for using the various features, and includes a complete API reference for developers.

86. Amazon Transcribe Medical

86.1. Service Overview

Amazon Transcribe Medical is an automatic speech recognition (ASR) service that makes it easy for you to add medical speech-to-text capabilities to your voice-enabled applications. Conversations between health care providers and patients provide the foundation of a patient's diagnosis and treatment plan and clinical documentation workflow. It's critically important that this information is accurate. However, accurate medical transcriptions such as dictation recorders and scribes are expensive, time consuming, and disruptive to the patient experience. Some organizations use existing medical transcription software, but find them inefficient and low in quality.

Driven by state-of-the-art machine learning, Amazon Transcribe Medical accurately transcribes medical terminologies such as medicine names, procedures, and even conditions or diseases. Amazon Transcribe Medical can serve a diverse range of use cases such as transcribing physician-patient conversations for clinical documentation, capturing phone calls in pharmacovigilance, or subtitling telehealth consultations.

Amazon Transcribe Medical is available as a set of public APIs that can address both batch workloads and real-time speech-to-text applications. The service is HIPAA-eligible and prioritizes patient data privacy and security. Amazon Transcribe Medical provides transcription expertise for primary care and specialty care areas such as cardiology, neurology, obstetrics-gynaecology, paediatrics, oncology, radiology and urology.

86.1.1. Features

- **Audio inputs:** Transcribe is designed to process live and recorded audio or video input to provide high quality transcriptions for search and analysis. We also offer separate APIs that uniquely understand customer calls (Amazon Transcribe Call Analytics) and medical conversations (Amazon Transcribe Medical).
- **Streaming & batch transcription:** You can process your existing audio recordings or stream the audio for real-time transcription. Using a secure connection, you can send a live audio stream to the service, and receive a stream of text in response.
- **Domain specific models:** Select a model that is tuned to telephone calls or multimedia video content. For example, Transcribe adapts to low-fidelity phone audio common in contact centers.
- **Automatic language identification:** With Amazon Transcribe, you can automatically identify the dominant language in an audio file and generate transcriptions. This is useful when your media library contains audio files in different languages. You can also use

this feature for media content classification and verify that the main spoken language in your videos and podcasts is correctly labeled.

- **Easy to read transcripts:** Amazon Transcribe enables you produce accurate transcripts that are easy to read, review, and integrate into your specific applications. We work to make the output ready for downstream activities such as call transcript analysis, subtitling, and content search.
- **Punctuation & number normalization:** Amazon Transcribe automatically adds punctuation and number formatting, so that the output closely matches the quality of manual transcription at a fraction of the time and expense. Numbers are also transcribed into digits or “normal form” instead of words.
- **Timestamp generation:** Amazon Transcribe returns a timestamp for each word, so that you can easily find a word or phrase in the original recording or add subtitles to video.
- **Recognize multiple speakers:** Speaker changes are automatically recognized and attributed in the text to capture scenarios like telephone calls, meetings, and television shows accurately. To learn more about speaker identification.
- **Channel identification:** Contact centers can submit a single audio file to Amazon Transcribe, and the service will identify produce a single transcript annotated by channel labels automatically.
- **Customize your output:** Accuracy is critical and we provide you many options to customize transcripts to your specific business needs and vernacular. Transcribe also provides up to 10 alternative transcriptions for each sentence, so you can quickly choose the best option that applies to your content and domain. This is useful for human in-the-loop subtitling workflows.
- **Custom vocabulary:** With custom vocabulary, you can add new words to the base vocabulary to generate more accurate transcriptions for domain-specific words and phrases like product names, technical terminology, or names of individuals.
- **Custom language models:** When needed, you can build and train your own custom language model (CLM) for your use case and domain by submitting a corpus of text data to Amazon Transcribe. CLM is a suitable feature for enhancing speech recognition accuracy with your own data.
- **User safety & privacy features:** Ensuring customer privacy and safety is critical. When needed, Transcribe enables you to mask or remove words that are sensitive or unsuitable for your audience from transcription results.
- **Vocabulary filtering:** You can specify a list of words to remove from transcripts with vocabulary filtering. For example, you can specify a list of profane or offensive words and Amazon Transcribe removes them from transcripts automatically.
- **Automatic content redaction / PII redaction:** When instructed, Amazon Transcribe can help customers identify and redact sensitive personally identifiable information (PII) from the supported language transcripts. This allows contact centers to easily review and share the transcripts for customer experience insight and agent training.
- **Data Protection:** Secure data at rest using Amazon S3 key (SSE-S3) or specify your own AWS Key Management Service key. Amazon Transcribe uses TLS (Transport Layer Security) 1.2, a cryptographic protocol that enables authenticated connections and

secure data transport over the internet via HTTP, with AWS certificates to encrypt data in transit. This includes streaming transcriptions.

- **Amazon Transcribe Call Analytics:** Extract conversation insights like call sentiment and speech loudness to improve agent productivity and customer experience with Amazon Transcribe Call Analytics.
- **Extract detailed call analytics & conversation insights:** Using the power of machine learning, you can quickly apply speech-to-text and natural language processing capabilities to uncover valuable conversation insights. You can then integrate insights such as customer and agent sentiment, detected issues, and speech characteristics like non-talk time, interruptions, and talk-speed into your inbound and outbound call analytics applications. This can help your supervisors more readily identify potential customer issues, agent coaching opportunities, and call trends.
- **Improve compliance & monitoring with automated call categorization:** Monitor your calls at scale to track compliance with company policies or regulatory requirements. Build and train your own custom categories based on your specified criteria (e.g. words/phrases or conversation characteristics). For example, you can setup category labels to see what percentage of calls are upsells or account cancellation.
- **Produce rich call transcripts:** Give your agents access to the conversation details from past interactions. The turn-by-turn transcripts provide insights such as customer sentiment, detected issues and interruptions.
- **Protect sensitive customer data:** Conversations often contain sensitive customer data such as names, addresses, credit card numbers, and social security numbers. Transcribe Call Analytics helps customers identify and redact this information from both the audio and text.
- **Amazon Transcribe Medical:** Easily transcribe your medical conversations with Transcribe Medical, a HIPAA-eligible automatic speech recognition (ASR) service.
- **Dictation mode:** Accurately transcribe single-speaker audio commonly found in medical dictation use cases.
- **Conversational mode:** Accurately transcribe multi-speaker conversational audio consisting of clinicians and/or patients alike.
- **Medical specialties:** Transcribe speech to text across a diverse range of medical specialties.
- **Batch API:** Transcribe recorded medical audio files at scale with high concurrency.
- **Streaming API:** Transcribe audio streams in near real time via either WebSocket Secure or HTTP/2 protocols.
- **Custom vocabulary:** Boost transcription accuracy by using custom vocabulary for potentially out-of-lexicon terminology.
- **Channel identification:** Concurrently transcribe multi-channel audio at no extra charge. Get one final coherent transcript.
- **Speaker diarization:** Separate speech from different speakers within any mono-channel audio.

86.1.2. Benefits

- **Save time with highly accurate transcriptions:** Amazon Transcribe Medical provides accurate medical speech-to-text capabilities that can be easily integrated into voice applications. Whether users want to dictate medical notes or transcribe drug-safety monitoring phone calls for downstream analysis, the service offers accurate speech recognition that is both scalable and cost-effective.
- **Lower medical transcription costs:** Amazon Transcribe Medical is a scalable transcription service that lives in the cloud. Pay only for what you transcribe, with no fixed costs, upfront commitments, or long-term licenses. Flexibly scale up or down the usage based on the needs of your application and use case.
- **Ease of use:** Using Amazon Transcribe Medical is straightforward. No prior machine learning knowledge or experience is required. Developers can focus on building their medical voice applications, by simply integrating with the service's easy-to-use APIs. Transcribe Medical handles the heavy lifting of developing state-of-the-art speech recognition models.
- **Data security and privacy:** Amazon Transcribe Medical is a HIPAA-eligible speech recognition service that prioritizes patient data security and privacy. The service is stateless, which means that it neither stores inbound audio nor output text. Users have full control over their data and determine whether they prefer to store transcriptions in local environments or self-managed cloud storage.

86.2. Backup/Restore and Disaster Recovery

This requirement is not applicable for this service. For additional information beyond what is described herein, please refer to http://aws.amazon.com/documentation/.

86.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

86.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/transcribe/
- **Service quotas:** https://docs.aws.amazon.com/transcribe/latest/dg/limits-guidelines.html
- **Service FAQs:** https://aws.amazon.com/transcribe/faqs/

86.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/transcribe/ and the following links for comprehensive technical documentation regarding this service.

- Developer Guide: Provides a conceptual overview of Amazon Transcribe, offers detailed instructions for using the various features, and includes a complete API reference for developers.

87. Amazon Translate

87.1. Service Overview

Amazon Translate is a neural machine translation service that delivers fast, high-quality, affordable, and customizable language translation. Neural machine translation is a form of language translation automation that uses deep learning models to deliver more accurate and more natural sounding translation than traditional statistical and rule-based translation algorithms. With Amazon Translate, you can localize content such as websites and applications for your diverse users, easily translate large volumes of text for analysis, and efficiently enable cross-lingual communication between users.

87.1.1. Features

- **Broad Language Coverage:** Amazon Translate supports translation between a wide range of languages. For up to date list of supported languages, see the Amazon Translate documentation.
- **Neural Network-Based:** Amazon Translate uses deep learning techniques to produce more accurate and fluent translation than traditional statistical and rule-based translation models. The neural machine translation system is built on a neural network that takes into account the entire context of the source sentence as well as the translation it has generated so far, to create more accurate and fluent translations. In comparison, conventional phrased-based machine translation only translates within the context of a few words before and after the translated word.
- **Customized Machine Translation:** Using Active Custom Translation (ACT), Amazon Translate allows you to take greater control over the output of your machine translation. Now you can bring your data (Parallel data) to Amazon Translate to customize the machine translated output to suit your needs. ACT produces custom-translated output without the need to build and maintain a custom translation model.
- **Named Entity Translation Customization:** Using Custom Terminology, Amazon Translate allows you to define how terms or names that are unique to certain organizations, domain, and industry get translated. The ability to customize output with Custom Terminology can decrease the number of translations that need to be edited by professional translators, resulting in cost savings and faster translations.
- **Language Identification:** Amazon Translate automatically identifies the source language when it is not specified. For example, user-generated content such as customer reviews and social media streams often do not contain a language code. Amazon Translate can automatically identify languages with high accuracy.
- **Batch and Real-Time Translations:** Amazon Translate is great for performing both batch translation when you have large quantities of pre-existing text to translate and real-time translation when you want to deliver on-demand translations of content as a feature of your applications.
- **Secure Machine Translation:** Communication between your webpage or applications and the Amazon Translate service is protected by SSL encryption. Any content processed by Amazon Translate is encrypted and stored at rest in the AWS Region where you are using the service. Administrators can also control access to Amazon Translate through an AWS Identity and Access Management (IAM) permissions policy – ensuring that sensitive information is kept secure and confidential.

- **Pay-Per-Use:** With Amazon Translate you pay only for what you use, making it easy and cost effective to scale your translation needs. You are charged based on the total number of characters sent to the API for translation.

87.1.2. Benefits

- **Highly Accurate & Continuously Improving:** Amazon Translate is a neural machine translation service. The translation engines are always improving from new and expanded datasets to produce more accurate translations for a wide range of use cases.
- **Easy to Integrate into Your Applications:** Amazon Translate removes the complexity of building real-time and batch translation capabilities into your applications with a simple API call. This makes it easy to localize an application or web site, or process multilingual data within your existing workflows.
- **Customizable:** With Custom Terminology and Active Custom Translation, Amazon Translate lets you customize your machine translated output. Use Custom Terminology to define how your brand names, model names, and other unique terms get translated. Use Active Custom Translation to generate a custom machine translated output that is tailored for your domain specific needs. You don't need to build a custom translation model, you can update your model as often as you need, and you only pay for the number of characters you translate.
- **Cost effective:** With the power of machine translation, Amazon Translate is about a 1000x cheaper than having your content manually translated by a professional translator.
- **Scalable:** Whether it's just a few words or large volumes of text, Amazon Translate easily scales to meet your translation needs with fast and reliable translations.
- **Versatile:** Amazon Translate can translate various content formats including Word documents, PowerPoint presentations, and Excel spreadsheets.

87.2. Backup/Restore and Disaster Recovery

This requirement is not applicable for this service. For additional information beyond what is described herein, please refer to http://aws.amazon.com/documentation/.

87.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

87.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/translate/
- **Service quotas:** https://docs.aws.amazon.com/general/latest/gr/translate-service.html
- **Service FAQs:** https://aws.amazon.com/translate/faqs/

87.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/translate/ and the following links for comprehensive technical documentation regarding this service.

- Developer Guide: Provides a conceptual overview of Amazon Translate, offers detailed instructions for using the various features, and includes a complete API reference for developers.

88. Amazon Virtual Private Cloud (VPC)

88.1. Service Overview

Amazon Virtual Private Cloud (Amazon VPC) gives you full control over your virtual networking environment, including resource placement, connectivity, and security. Get started by setting up your VPC in the AWS service console. Next, add resources to it such as Amazon Elastic Compute Cloud (EC2) and Amazon Relational Database Service (RDS) instances. Finally, define how your VPCs communicate with each other across accounts, Availability Zones, or AWS Regions. In the example below, network traffic is being shared between two VPCs within each Region.

88.1.1. Features

- **Flow Logs:** You can monitor your VPC flow logs delivered to Amazon Simple Storage Service (Amazon S3) or Amazon CloudWatch to gain operational visibility into your network dependencies and traffic patterns, detect anomalies and prevent data leakage, and troubleshoot network connectivity and configuration issues. The enriched metadata in flow logs helps you learn more about who initiated your TCP connections and the packet-level source and destination for traffic flowing through intermediate layers (such as a NAT gateway).
- **IP Address Manager (IPAM):** IPAM makes it easier for you to plan, track, and monitor IP addresses for your AWS workloads. IPAM automates IP address assignments to your Amazon VPC, removing the need to use homegrown or spreadsheet-based planning applications. It also enhances your network observability by showing IP usage across multiple accounts and VPCs in a unified operational view.
- **IP Addressing:** IP addresses enable resources in your VPC to communicate with each other and with resources over the internet. Amazon VPC supports both the IPv4 and IPv6 addressing protocols. In a VPC, you can create IPv4-only, dual-stack, and IPv6-only subnets and launch Amazon EC2 instances in these subnets. Amazon also gives you multiple options to assign public IP addresses to your instances. You can use the Amazon provided public IPv4 addresses, Elastic IPv4 addresses, or an IP address from the Amazon provided IPv6 CIDRs. Apart from this, you have the option to bring your own IPv4 or IPv6 addresses within the Amazon VPC that can be assigned to these instances.
- **Ingress Routing:** With this feature, you can route all incoming and outgoing traffic flowing to/from an internet gateway or virtual private gateway to a specific Amazon EC2 instance's elastic network interface. Configure your virtual private cloud to send all traffic to a gateway or an Amazon EC2 instance before it reaches your business workloads.
- **Network Access Analyzer:** Network Access Analyzer helps you verify that your network on AWS conforms to your network security and compliance requirements. Network Access Analyzer lets you specify your network security and compliance requirements, and identifies unintended network access that does not meet your specified requirements. You can use Network Access Analyzer to understand network access to your resources, helping you identify improvements to your cloud security posture and easily demonstrate compliance.

- **Network Access Control List:** A network access control list (network ACL) is an optional layer of security for your VPC that acts as a firewall for controlling traffic in and out of one or more subnets. You might set up network ACLs with rules similar to those of your security groups.
- **Reachability Analyzer:** This static configuration analysis tool enables you to analyse and debug network reachability between two resources in your VPC. After you specify the source and destination resources, Reachability Analyzer produces hop-by-hop details of the virtual path between them when they are reachable, and identifies the blocking component when they are unreachable.
- **Security Groups:** Create security groups to act as a firewall for associated Amazon EC2 instances, controlling inbound and outbound traffic at the instance level. When you launch an instance, you can associate it with one or more security groups. If you don't specify a group, the instance is automatically associated with the VPC's default group. Each instance in your VPC can belong to a different set of groups.
- **Traffic Mirroring:** This feature allows you to copy network traffic from an elastic network interface of Amazon EC2 instances and send it to out-of-band security and monitoring appliances for deep packet inspection. You can detect network and security anomalies, gain operational insights, implement compliance and security controls, and troubleshoot issues. Traffic Mirroring gives you direct access to the network packets flowing through your VPC.

88.1.2. Benefits

- **Secure:** Secure and monitor connections, screen traffic, and restrict instance access inside your virtual network.
- **Efficient:** Spend less time setting up, managing, and validating your virtual network.
- **Customizable:** Customize your virtual network by choosing your own IP address range, creating subnets, and configuring route tables.

88.2. Backup/Restore and Disaster Recovery

This requirement is not applicable for this service. For additional information beyond what is described herein, please refer to http://aws.amazon.com/documentation/.

88.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

88.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/vpc/index.html
- **Service quotas:** https://docs.aws.amazon.com/vpc/latest/userguide/amazon-vpc-limits.html
- **Service FAQs:** https://aws.amazon.com/vpc/faqs/

88.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/vpc/index.html and the following links for comprehensive technical documentation regarding this service.

- User Guide: Create and configure your virtual private cloud (VPC).
- Peering Guide: Connect your VPC with other VPCs and access resources across VPCs.
- Traffic Mirroring Guide: Capture and mirror network traffic for your Amazon EC2 instances.
- Reachability Analyzer Guide: Analyze and debug network reachability between resources in your VPCs.
- Network Access Analyzer Guide: Identify unintended network access to resources in your VPCs.

89. Amazon WorkDocs

89.1. Service Overview

Amazon WorkDocs is a fully managed, secure content creation, storage, and collaboration service. With Amazon WorkDocs, you can easily create, edit, and share content, and because it's stored centrally on AWS, access it from anywhere on any device. Amazon WorkDocs makes it easy to collaborate with others, and lets you easily share content, provide rich feedback, and collaboratively edit documents. You can use Amazon WorkDocs to retire legacy file share infrastructure by moving file shares to the cloud. Amazon WorkDocs lets you integrate with your existing systems, and offers a rich API so that you can develop your own content-rich applications. Amazon WorkDocs is built on AWS, where your content is secured on the world's largest cloud infrastructure.

89.1.1. Features

- **Secure Storage and Sharing:** You can store virtually any type of file on Amazon WorkDocs. Each individual user account on Amazon WorkDocs includes 1 TB of storage capacity by default. Administrators can set storage limits for individual users and also purchase additional storage for users on a pay-as-you-go basis.
- **Search:** Searching made easy with Amazon WorkDocs Smart Search: Amazon WorkDocs Smart Search speeds up your content searches so you can spend more time creating, editing, and sharing files with colleagues.
- **Encryption:** With Amazon WorkDocs, your content is encrypted in transit and at rest to ensure the security of your data and help you meet regulatory and compliance requirements. Amazon WorkDocs is built on AWS, where security is our number one priority.
- **WorkDocs Drive:** Amazon WorkDocs Drive is a desktop application that combines the ease of working in Windows File Explorer or Mac Finder with the scale of Amazon WorkDocs. With Amazon WorkDocs Drive, all of your files are available on-demand from your device without consuming valuable disk space on your PC or Mac. You can use Amazon WorkDocs Drive as your primary user drive, and you don't need to use network shares to store your content.

- **Collaboration:** With Amazon WorkDocs, you can add private comments, format comment text, resolve comments, and respond to comments in a threaded conversation. When providing feedback in Amazon WorkDocs, you can add overall comments or comment on specific sections of a file.
- **Editing in Microsoft Office:** You can edit documents directly in Microsoft Office with Amazon WorkDocs Companion. Amazon WorkDocs Companion is an app that lets you edit Microsoft Office, .pdf, and .txt files from your browser. You can save your changes as a new version on your Amazon WorkDocs site. With Amazon WorkDocs Companion, you no longer need to manually download, save, and upload files when accessing Amazon WorkDocs from your browser.
- **Activity Feed:** Amazon WorkDocs Activity Feed helps you easily keep track of engagements with your content by other users on your Amazon WorkDocs site. You can search actions by file, folder, or user name in real-time. You can also use multiple options to filter your results.
- **Active Directory Integration:** Amazon WorkDocs lets you use your Active Directory to manage your users. If you use Active Directory, you can create user groups, enable multi-factor authentication (MFA), and configure single sign-on (SSO) for your Amazon WorkDocs site. Your users can also log in with their existing credentials when you use Active Directory with Amazon WorkDocs.
- **Data Residency:** You can specify in which AWS Region to store your content to help meet data residency requirements. Your users can access your Amazon WorkDocs site from anywhere in the world regardless of which AWS region you choose. Refer to AWS Regions to see where Amazon WorkDocs is currently available.
- **For Developers:** The Amazon WorkDocs SDK helps you build content collaboration and management capabilities into your solutions and applications by providing full administrator and user level access to Amazon WorkDocs site resources. Using the Amazon WorkDocs SDK, you can also integrate the activity feed with your analytics solutions to create real-time monitoring of your Amazon WorkDocs users and files.

89.1.2. Benefits

- **Migrate your on-premise file servers and reduce costs significantly**: Amazon WorkDocs is a fully managed service that lets you retire expensive network file shares and painlessly move content to the cloud. With Amazon WorkDocs pay-as-you-go pricing, you only pay for the active user accounts on your site. With Amazon WorkDocs Drive, you can launch content directly from Windows File Explorer, Mac Finder, or Amazon WorkSpaces, all without consuming valuable local disk space.
- **Securely share with internal teams and external users in real-time:** Amazon WorkDocs allows you to easily share with teams and invite external users for cross-organizational collaboration. You can use real-time Activity Feed to track site-wide collaboration actions by file, folder, or user name. With Amazon WorkDocs you can reduce long email threads with commenting, highlighting, and requesting feedback capabilities. Granular searching helps you find feedback from colleagues across documents as well.
- **Secure your content in the cloud:** Amazon WorkDocs lets you store your content on the world's largest global cloud infrastructure, built to satisfy the requirements of our most security-sensitive customers. Your content is encrypted in transit and at rest. You

can review user and admin activity tracking to know who is accessing what. With Amazon WorkDocs, you can also easily maintain compliance: Amazon WorkDocs is HIPAA eligible, GDPR and PCI DSS compliant, evaluated with SOC reports 1-3, and aligned with ISO compliance requirements.

- **Bring content into your applications and processes:** Amazon WorkDocs makes it easier for you to add content rich features to your web and mobile applications by using Amazon WorkDocs as a content repository. With Amazon WorkDocs, you can use the user and admin level capabilities of the AWS SDK to integrate with your business tools like anti-virus and malware detection applications as well.
- **Route your documents using approval workflow:** Using Amazon WorkDocs approval workflow, you can now route documents and other files stored in WorkDocs for approvals. Approval Workflow automatically routes the files, assigns review tasks to approvers and sends reminders and notifications. You can also track the status of your approval requests as well as those awaiting your approval.
- **Extend your desktop to the cloud:** Amazon WorkDocs Drive is a native desktop application that combines the ease of working in Windows File Explorer or Mac Finder with the scale of Amazon WorkDocs. With Amazon WorkDocs Drive, all of your files are available on-demand from your device without having to consume valuable disk space on your PC or macOS. You can use Amazon WorkDocs Drive as your primary user drive, and you don't need to use network shares to store your content.

89.2. Backup/Restore and Disaster Recovery

This requirement is not applicable for this service. For additional information beyond what is described herein, please refer to http://aws.amazon.com/documentation/.

89.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

89.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/workdocs/index.html
- **Service quotas:** https://docs.aws.amazon.com/general/latest/gr/workdocs.html
- **Service FAQs:** https://aws.amazon.com/workdocs/faq/

89.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/workdocs/index.html and the following links for comprehensive technical documentation regarding this service

- Administration Guide: Helps you use Amazon WorkDocs to perform several administrative tasks, such as creating a new directory in the cloud, connecting to your on-premises directory, or setting user privileges and defaults.
- User Guide: Helps you use the Amazon WorkDocs collaboration and utility applications. File collaboration applications are available for desktop web browsers, as well as several tablets.

- Developer Guide: Helps you use Amazon WorkDocs as a developer through the AWS SDK.

90. Amazon WorkMail

90.1. Service Overview

Amazon WorkMail is a secure, managed business email and calendar service with support for existing desktop and mobile email client applications. Amazon WorkMail gives users the ability to seamlessly access their email, contacts, and calendars using the client application of their choice, including Microsoft Outlook, native iOS and Android email applications, any client application supporting the IMAP protocol, or directly through a web browser. You can integrate Amazon WorkMail with your existing corporate directory, use email journaling to meet compliance requirements, and control both the keys that encrypt your data and the location in which your data is stored. You can also set up interoperability with Microsoft Exchange Server, and programmatically manage users, groups, and resources using the Amazon WorkMail SDK.

90.1.1. Features

- **Microsoft Outlook Compatible:** Amazon WorkMail provides native support for Microsoft Outlook on Windows and Mac OS X, so users can continue to use the email client they are already using without needing to install any additional software, such as plug-ins for Microsoft Outlook. Amazon WorkMail supports all of the advanced Microsoft Outlook capabilities your users depend on, such as free/busy scheduling, delegation, and out of office replies.
- **Enterprise-Grade Security:** Amazon WorkMail automatically encrypts all of your data at rest with encryption keys you control, using the AWS Key Management Service. Amazon WorkMail also allows you to retain full control over data locality by choosing the AWS region where all of your data is stored. All data in transit is encrypted using industry-standard SSL encryption.
- **Active Directory Integration:** You can integrate your existing Microsoft Active Directory with Amazon WorkMail using AWS Directory Service AD Connector or run AWS Directory Service for Microsoft Active Directory Enterprise Edition ("Microsoft AD") so you don’t have to manage users in two places and users can continue to use their existing Microsoft Active Directory credentials. Alternatively, you can have Amazon WorkMail create and manage a Simple AD directory for you and have users in that directory created when you add them to your Amazon WorkMail organization.
- **Interoperability with Microsoft Exchange Server:** You can set up interoperability with Microsoft Exchange Server 2010 and 2013. Interoperability allows you to use both Amazon WorkMail and Microsoft Exchange Server at the same time so that you can use the same corporate domain for mailboxes across both environments.
- **Administrative SDK:** Amazon WorkMail provides an administrative SDK so you can natively integrate WorkMail with your existing services. The SDK enables programmatic user, email group, and meeting room or equipment resource management.
- **IMAP Protocol Support:** Amazon WorkMail allows you to access your email with any email client application that supports the IMAP protocol. This allows you to choose from a broad range of clients on the platform of your choice, including those that do not offer support for Microsoft Exchange compatible protocols like Microsoft Exchange Web Services and Microsoft ActiveSync.

- **Feature-Rich Web Client:** Amazon WorkMail provides a feature-rich web client so that users can access their email and calendars, view shared calendars, quickly schedule meetings with co-workers, or search the company address book using their browser. Users can also manage their out-of-office replies and book resources, such as conference rooms.
- **Mobile Device Management:** Amazon WorkMail enforces security policies that you specify for your users' mobile devices via the Microsoft Exchange ActiveSync protocol.
- **Large Mailboxes:** Amazon WorkMail provides a 50 GB mailbox size by default, eliminating the need for users to have to worry about dealing with large attachments or mailbox size restrictions.
- **AWS CloudTrail Integration:** Amazon WorkMail integrates with AWS CloudTrail. CloudTrail is a service that enables governance, compliance, operational auditing, and risk auditing of your AWS account.

90.1.2. Benefits

- **Managed service:** Amazon WorkMail makes it easy to manage your corporate email infrastructure and eliminates the need for up-front investments to license and provision on-premises email servers. There is no complex software to install or maintain and no hardware to manage. Amazon WorkMail automatically handles all of the patches, back-ups, and upgrades.
- **Enterprise-grade security:** Amazon WorkMail automatically encrypts all of your data at rest with encryption keys you control, using the AWS Key Management Service (KMS). Amazon WorkMail also allows you to retain full control over data locality by choosing the AWS region where all of your data is stored.
- **Outlook compatible:** Amazon WorkMail provides native support for Microsoft Outlook on both Windows and Mac OS X, so users can continue to use the email client they are already using without needing to install any additional software, such as plug-ins for Microsoft Outlook.
- **Anywhere access:** Users can synchronize their mailboxes with iOS, Android, Amazon Fire, and Windows Phone devices. If you are migrating from an on-premises Microsoft Exchange server, your users' mobile devices can automatically connect to Amazon WorkMail with no end-user reconfiguration required, and no change in user experience. A feature-rich web client is also available for users to access their email, calendar, and contacts.
- **Active directory integration:** Amazon WorkMail securely integrates with your existing Microsoft Active Directory so that users can access their mailbox using their existing credentials. This also makes it easy to manage users and groups with familiar systems management tools, such as Active Directory Users and Computers.
- **Low cost:** Amazon WorkMail features simple, low, monthly per-user pricing and costs $4 per user per month which includes 50GB of storage per user.

90.2. Backup/Restore and Disaster Recovery

This requirement is not applicable for this service. For additional information beyond what is described herein, please refer to http://aws.amazon.com/documentation/.

90.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

90.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/workmail/index.html
- **Service quotas:** https://docs.aws.amazon.com/workmail/latest/adminguide/workmail_limits.html
- **Service FAQs:** https://aws.amazon.com/workmail/faqs/

90.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/workmail/index.html and the following links for comprehensive technical documentation regarding this service

- User Guide: Helps you configure your desktop email client or mobile device for Amazon WorkMail, and use the Amazon WorkMail web application.
- Administrator Guide: Helps you perform administrative tasks for Amazon WorkMail, such as setting up Amazon WorkMail for your organization, adding a domain, and managing users, groups, and mobile devices.

91. Amazon WorkSpaces

91.1. Service Overview

Amazon WorkSpaces offers you an easy way to provide a secure, managed, cloud-based virtual desktop experience to your end-users. Unlike traditional on-premises Virtual Desktop Infrastructure (VDI) solutions, you don't have to worry about procuring, deploying, and managing a complex environment – Amazon WorkSpaces takes care of the heavy lifting and provides a fully managed service. With Amazon WorkSpaces, you can deliver a high quality portable desktop, and applications, to your users on the device of their choice.

91.1.1. Features

- **Streaming protocols:** Amazon WorkSpaces utilizes streaming protocols to provide users a secure and high quality experience. These protocols analyse the hosted desktop, network, and user's device to select compression and decompression algorithms (codecs) that encode a rendering of the user's desktop and transmit it as a pixel stream to the user's device.
- **Amazon WorkSpaces Bundles:** To get started, select from a choice of Amazon WorkSpaces bundles that offer different hardware and software options, and launch the number of WorkSpaces you require. When WorkSpaces are provisioned, users receive an email providing instructions on where to download the WorkSpaces client applications they need, and how to connect to their WorkSpace.
- **Bring Your Own Licenses:** You can bring your existing Windows 10 Desktop licenses to Amazon WorkSpaces and run them on hardware that is physically dedicated to you. When you bring your existing Windows licenses to WorkSpaces, you can save up to 16% ($4 per month per WorkSpace) over WorkSpaces with a new Windows license

included. To be eligible, your organization must meet the licensing requirements set by Microsoft, and you must commit to running at least 100 Amazon WorkSpaces in a given AWS region each month. If you plan to use GPU-enabled (Graphics, GraphicsPro, Graphics.g4dn, and GraphicsPro.g4dn) bundles, verify that you will run a minimum of 4 AlwaysOn or 20 AutoStop GPU-enabled WorkSpaces in a Region per month on dedicated hardware. To learn more about this licensing option. To learn more about this licensing option, please see the Amazon WorkSpaces FAQ.

- **Easy provisioning:** Provisioning desktops with Amazon WorkSpaces is easy. Whether you choose to launch one or many Amazon WorkSpaces, all you need to do is to choose the bundles that best meet the needs of your users, and the number of Amazon WorkSpaces that you would like to launch.
- **Secure and encrypted:** Amazon WorkSpaces provides a high quality desktop experience that helps you meet compliance and security requirements, including HIPAA-eligibility and PCI compliance. With WorkSpaces, your organization's data is not sent to or stored on end-user devices. The PC-over-IP (PCoIP) remote display protocol used by WorkSpaces provides the familiar desktop experience to the user while the data remains in the AWS cloud or in your on-premises environment.
- **Active Directory and RADIUS integration:** Amazon WorkSpaces allows you to use your on-premises Microsoft Active Directory to manage your WorkSpaces and your end user credentials. By integrating with your on-premises Active Directory, your users can log in with their existing credentials, you can apply Group Policies to your WorkSpaces, you can deploy software to your WorkSpaces using your existing tools, and you can use your existing RADIUS server to enable multi-factor authentication (MFA).
- **Persistent storage:** Amazon WorkSpaces provides each user with access to varying amounts of persistent storage (SSD Volumes) in the AWS cloud based on the bundle selected. Data that users store on the 'user volume' attached to the WorkSpace is automatically backed up to Amazon S3 on a regular basis. Amazon S3 is designed for 99.999999999% durability of objects, providing you with peace of mind about your users' data

91.1.2. Benefits

- **Onboard contingent workers:** Easily assign and remove desktops for contractors while keeping your sensitive data secure in the cloud.
- **Facilitate remote work:** Enable work-from-home and remote workers to access fully functional Windows and Linux desktops from any location.
- **Run powerful desktops:** Provide high-performance desktops for developers and engineers to store and access proprietary models, designs, and code.
- **Let contact centre agents work from anywhere:** Enable contact centre agents to work from anywhere with a secure, easy-to-use agent experience.

91.2. Backup/Restore and Disaster Recovery

This requirement is not applicable for this service. For additional information beyond what is described herein, please refer to http://aws.amazon.com/documentation/.

91.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

91.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/workspaces/
- **Service quotas:** https://docs.aws.amazon.com/workspaces/latest/adminguide/workspaces-limits.html
- **Service FAQs:** https://aws.amazon.com/workmail/faqs/

91.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/workspaces/ and the following links for comprehensive technical documentation regarding this service

- Administration Guide: Helps you get started using Amazon WorkSpaces. Learn how to quickly and easily provision and maintain one or more WorkSpaces.
- User Guide: Helps you get started using an Amazon WorkSpaces client application or web browser to access your WorkSpace.

92. AWS Amplify

92.1. Service Overview

AWS Amplify is a set of purpose-built tools and features that lets frontend web and mobile developers quickly and easily build full-stack applications on AWS, with the flexibility to leverage the breadth of AWS services as your use cases evolve. With Amplify, you can configure a web or mobile app backend, connect your app in minutes, visually build a web frontend UI, and easily manage app content outside the AWS console. Ship faster and scale effortlessly—with no cloud expertise needed.

92.1.1. Features

- **Authentication:** Create seamless on-boarding flows with a fully-managed user directory and pre-built sign-up, sign-in, forgot password, and multi-factor auth workflows. Amplify also supports login with a social provider such as Facebook, Google Sign-In, or Login With Amazon and provides fine grained access control to mobile and web applications. Powered by Amazon Cognito.
- **DataStore:** Use a multi-platform (iOS/Android/React Native/Web) on-device persistent storage engine that automatically synchronizes data between mobile/web apps and the cloud, powered by GraphQL. DataStore provides a programming model for leveraging shared and distributed data without writing additional code for offline and online scenarios, which makes working with distributed, cross-user data just as simple as working with local-only data. Powered by AWS AppSync.
- **Analytics:** Understand the behaviour of your web, iOS or Android users. Use auto tracking to track user sessions and web page metrics or create custom user attributes and in-app metrics. Get access to real time data stream and analyse the data for

customer insights and build data driven marketing strategies to drive customer adoption, engagement, and retention. Powered by Amazon Pinpoint and Amazon Kinesis.

- **API:** Make secure HTTP requests to GraphQL and REST endpoints to access, manipulate, and combine data from one or more data sources such Amazon DynamoDB, Amazon Aurora Serverless, and your custom data sources with AWS Lambda. Amplify enables you to easily build scalable applications that require real-time updates, local data access for offline scenarios, and data synchronization with customizable conflict resolution when devices are back online. Powered by AWS AppSync and Amazon API Gateway.
- **Functions:** Add a Lambda function to your project which you can use alongside a REST API or as a data source in your GraphQL API using the @function directive in the Amplify CLI. You can update the Lambda execution role policies for your function to access other resources generated and maintained by the CLI, using the CLI. Amplify CLI enables you to create, test and deploy Lambda functions across various runtimes and once a runtime is selected, you can select a function template for the runtime to help bootstrap your Lambda function.
- **Geo:** Add location-aware features like maps and location search to your JavaScript-based web app in minutes. Amplify Geo includes pre-integrated map UI components (based on the popular MapLibre open-source library) and it updates the Amplify Command Line Interface (CLI) tool with support for provisioning all required cloud location services. You can customize embedded maps to match your app's theme, or choose from many community-developed MapLibre plugins for more flexibility and advanced visualization options. Powered by Amazon Location Service.
- **Interactions:** Build interactive and engaging conversational bots with the same deep learning technologies that power Amazon Alexa with just a single line of code. Create great user experiences through chat bots when it comes to tasks such as automated customer chat support, product information/recommendations or streamlining common work activities etc.
- **Predictions:** Enhance your app by adding AI/ML capabilities. You can easily achieve use cases like text translation, speech generation from text, entities recognition in image, interpretation of text, and transcribing text. Amplify enables simplified orchestration of advanced use cases like uploading images for automatic training and using GraphQL directives for chaining multiple AI/ML actions. Powered by Amazon Machine Learning services, such as Amazon Sagemaker.
- **PubSub:** Pass messages between your app instances and your app's backend creating real-time interactive experiences. Amplify provides connectivity with cloud-based message-oriented middleware. Powered by AWS IoT services and Generic MQTT Over WebSocket Providers.
- **Push notifications:** Improve customer engagement by using marketing and analytics capabilities. Leverage customer insights to segment and target your customers more effectively. You can tailor your content and communicate through multiple channels including email, texts as well as push notifications. Powered by Amazon Pinpoint.
- **Storage:** Store and manage user generated content such as photos, videos securely on device or in the cloud. The AWS Amplify Storage module provides a simple mechanism for managing user content for your app in public, protected or private storage buckets.

Leverage cloud scale storage so that you can easily take your application from prototype to production. Powered by Amazon S3.

- **Amplify Hosting:** AWS Amplify offers a fully managed service for deploying and hosting fullstack web applications, with built-in CI/CD workflows that accelerate your application release cycle. A fullstack serverless app consists of a backend built with cloud resources such as GraphQL or REST APIs, file and data storage, and a frontend built with single page application frameworks such as React, Angular, Vue, or Gatsby. Simply connect your application's code repository in the Amplify console, and changes to your frontend and backend are deployed in a single workflow on every code commit.
- **Access AWS resources:** Amplify is built on top of Infrastructure-as-Code that deploys resources within your account. Add business logic to your backend using Amplify's Function and Container support. Grant functions access to an SNS topic to send an SMS or allow your container to access an existing database.
- **Import AWS resources:** Integrate your existing resources, such as your Amazon Cognito user pool and federated identities (identity pool), or Storage resources like DynamoDB + S3, into an Amplify project via Amplify Studio. This will enable your API (GraphQL), Storage (S3), and other resources to leverage your existing authentication mechanism.
- **Command hooks:** Use Command Hooks to run custom scripts before, during, and after Amplify CLI commands (“amplify push”, “amplify api gql-compile”, and more). Customers can trigger validation checks, run credential scans, or clean-up build artefacts during deployment. This allows you to extend Amplify’s best-practice defaults to meet your organization’s security guidelines and operational requirements.
- **Export Infrastructure-as-Code:** You can use Amplify with your existing DevOps guidelines and tools to enforce deployment policies or to integrate into your in-house deployment systems. Amplify’s export feature lets you export your Amplify project to your preferred tooling, using CDK. “amplify export” exports the Amplify CLI build artefacts (including CloudFormation templates, API resolver code, and client-side code generation).
- **Amplify Libraries:** AWS Amplify offers use case-centric open source libraries in the Amplify Framework to build cloud powered mobile and web apps. Amplify libraries are powered by AWS services and can be used with new backends created with the Amplify CLI and Amplify Studio, or your existing AWS backend.
- **Amplify UI components:** Amplify UI Components is an open-source UI toolkit that encapsulates cloud-connected workflows inside of cross-framework UI components. AWS Amplify provides drop-in UI components for authentication, storage and interactions, with a style guide for your apps that automatically integrate with your configured cloud services.
- **Amplify Studio:** Amplify Studio offers an easy way to develop app backends and manage app content. Amplify Studio provides a visual interface for modelling data, adding authentication, authorization and managing users and groups. As you create backend resources, Amplify Studio generates automation templates that enable seamless integration with the Amplify CLI, giving you the ability to extend your app backend with additional capabilities and create multiple environments for testing and team collaboration. With Amplify Studio, you can provide access to team members

without an AWS account so developers and non-developers can access the data they need to develop and manage apps more efficiently.

- **Amplify CLI toolchain:** The Amplify Command Line Interface (CLI) is a toolchain to configure and maintain your app backend from your local desktop. Configure cloud functionality using the CLI's interactive workflow and intuitive use cases such as auth, storage, API. Test features locally and deploy multiple environments. All configured resources are available to customers as infrastructure-as-code templates enabling better team collaboration and easy integration with Amplify's CI/CD workflow.

92.1.2. Benefits

- **Full-stack:** Create full-stack apps, frontend UI and backends visually, with authentication, storage, data, and more.
- **Connectivity:** Connect web and mobile apps to new and existing AWS resources in a few lines of code.
- **Easy to use:** Build, deploy, and host static websites, single-page web apps, and server-side rendered apps in just a few clicks.
- **Wide reaching:** Access 175+ AWS services to support new use cases, DevOps practices, and user growth.

92.2. Backup/Restore and Disaster Recovery

The AWS Amplify Console connects to your private or public code repositories. You should consider how your code repositories meet your organisational disaster recovery policy. For deployment of an application, the AWS Amplify Console leverages the Amazon CloudFront Global Edge Network to distribute your web app globally.

92.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

92.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.amplify.aws/
- **Service quotas:** https://docs.aws.amazon.com/general/latest/gr/amplify.html
- **Service FAQs:** https://aws.amazon.com/amplify/faqs/

92.5. Technical Requirements

Please refer to https://docs.amplify.aws/ for comprehensive technical documentation regarding this service.

93. AWS App Mesh

93.1. Service Overview

AWS App Mesh is a service mesh that provides application-level networking to make it easy for your services to communicate with each other across multiple types of compute infrastructure. App Mesh gives end-to-end visibility and high-availability for your applications.

Modern applications are typically composed of multiple services. Each service may be built using multiple types of compute infrastructure such as Amazon EC2, Amazon ECS, Amazon EKS, and AWS Fargate. As the number of services grow within an application, it becomes difficult to pinpoint the exact location of errors, re-route traffic after failures, and safely deploy code changes. Previously, this has required you to build monitoring and control logic directly into your code and redeploy your service every time there are changes.

AWS App Mesh makes it easy to run services by providing consistent visibility and network traffic controls, and helping you deliver secure services. App Mesh removes the need to update application code to change how monitoring data is collected or traffic is routed between services. App Mesh configures each service to export monitoring data and implements consistent communications control logic across your application.

You can use App Mesh with AWS Fargate, Amazon EC2, Amazon ECS, Amazon EKS, and Kubernetes running on AWS, to better run your application at scale. App Mesh also integrates with AWS Outposts for your applications running on-premises. App Mesh uses the open source Envoy proxy, making it compatible with a wide range of AWS partner and open source tools.

93.1.1. Features

- **Open source proxy:** App Mesh uses the open source Envoy proxy to manage all traffic into and out of a service's containers. App Mesh configures this proxy to automatically handle all of the service's application communications. Envoy has a vibrant ecosystem of community-built integrations that work with App Mesh.
- **Compatible AWS services:** Amazon CloudWatch* – monitoring and logging service for complete visibility of resources and applications. AWS X-Ray* – tracing service for an end-to-end view of application performance.
- **Compatible AWS partner and open source tools:** Datadog, Alcide, HashiCorp, Sysdig, Signalfx, Spotinst, Tetrate, Neuvector, Weaveworks, Twistlock, Wavefront by VMware, Aqua.
- **Traffic Routing:** App Mesh lets you configure services to connect directly to each other instead of requiring code within the application or using a load balancer. When each service starts, its proxies connect to App Mesh and receives configuration data about the locations of other services in the mesh. You can use controls in App Mesh to dynamically update traffic routing between services with no changes to your application code.
- **Client-side Traffic Policies:** The proxies automatically load balance traffic from all clients in the mesh, and add and remove load balancing endpoints based on health checks and service registration. These capabilities make it easier to deploy new versions of your services and help tune applications to be resilient to failures.
- **Service-to-Service Authentication:** Mutual TLS (mTLS) enables transport layer authentication, which provides service-to-service identity verification for the application components running in and outside service meshes. It allows customers to extend the security perimeter to the applications running in AWS App Mesh by provisioning certificates from AWS Certificate Manager Private Certificate Authority or a customer-managed Certificate Authority (CA) to workloads in the service mesh, and to enforce automatic authentication for client applications connecting to services.
- **Container orchestration native user experience:** App Mesh works with services managed by Amazon ECS, Amazon EKS, AWS Fargate, Kubernetes running on EC2.

For containerized workloads running on ECS, EKS, Fargate, or Kubernetes, you include the provided App Mesh proxy as part of the task or pod definition for each microservice and configure the services' application container to communicate directly with the proxy. When the service starts, the proxy automatically checks in with and is configured by App Mesh.

- **Fully managed:** AWS App Mesh is a managed and highly available service. App Mesh allows you to manage services communications without needing to install or manage application-level infrastructure for communications management.

93.1.2. Benefits

- **Get end-to-end visibility:** App Mesh captures metrics, logs, and traces from all of your applications. You can combine and export this data to Amazon CloudWatch, AWS X-Ray, and compatible AWS partner and community tools for monitoring and tracing. This lets you quickly identify and isolate issues with any service to optimize your entire application.
- **Streamline your operations:** App Mesh provides controls to configure and standardize how traffic flows between your services. You can easily implement custom traffic routing rules so that your service is highly available during deployments, after failures, and as your application scales. This removes the need to configure communication protocols for each service, write custom code, or implement libraries to operate your application.
- **Enhance network security:** App Mesh helps encrypt all requests between services even when they are in your private networks. Further, you can add authentication controls to ensure that only services that you allow interconnect.

93.2. Backup/Restore and Disaster Recovery

This requirement is not applicable for this service. For additional information beyond what is described herein, please refer to http://aws.amazon.com/documentation/.

93.3. Pricing Overview

There is no additional charge for using AWS App Mesh. You pay only for the AWS resources (EC2 instances or requested Fargate CPU and memory) consumed by the lightweight proxy that is deployed alongside your containers.

93.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/app-mesh/
- **Service quotas:** https://docs.aws.amazon.com/general/latest/gr/appmesh.html
- **Service FAQs:** https://aws.amazon.com/app-mesh/faqs/

93.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/app-mesh/ and the following links for comprehensive technical documentation regarding this service.

- User Guide: Describes key concepts and provides instructions for using the features of AWS App Mesh.

94. AWS App Runner

94.1. Service Overview

AWS App Runner is a fully managed service that makes it easy for developers to quickly deploy containerized web applications and APIs, at scale and with no prior infrastructure experience required. Start with your source code or a container image. App Runner builds and deploys the web application automatically, load balances traffic with encryption, scales to meet your traffic needs, and makes it easy for your services to communicate with other AWS services and applications that run in a private Amazon VPC. With App Runner, rather than thinking about servers or scaling, you have more time to focus on your applications.

94.1.1. Features

- **Automatic Deployments:** When you connect App Runner to your code repository or container image registry, App Runner can automatically build and deploy your application when you update your source code or container image.
- **Load Balancing:** App Runner automatically load balances traffic to provide high levels of reliability and availability for your applications.
- **Auto Scaling:** Enabled by default, App Runner automatically scales the number of containers up or down to meet the needs of your application.
- **Logs and Metrics:** App Runner makes it easy to monitor and optimize your containerized applications by providing detailed build, deployment, and runtime logs. You also get a full set of compute metrics with built-in Amazon CloudWatch integration.
- **Certificate Management:** App Runner includes fully managed TLS with no setup needed. App Runner automatically renews the certificates before their expiration date.
- **Cost Management:** Easily pause and resume your App Runner applications using the console, CLI, or API. You're only billed when the service is running.
- **Amazon Virtual Private Cloud (Amazon VPC) Access:** App Runner services can communicate with AWS services running in a private Amazon VPC, enabling you to easily add support for your applications on App Runner and ensure that network access is contained within your VPC.

94.1.2. Benefits

- **Easy to use:** With App Runner, you can build and run secure web scale applications in just a few clicks with no prior containers or infrastructure experience required. You don't need knowledge about server configuration, networking, load balancing, or deployment pipelines.
- **Scales with traffic:** App Runner makes it easy and cost effective to run your applications at web scale with high availability. App Runner seamlessly scales up resources in response to your traffic, and automatically scales down to your configured number of provisioned container instances to eliminate cold starts and ensure consistently low latency.
- **Saves time:** App Runner resources and infrastructure components are fully managed by AWS and benefit from our security and operational best practices. This enables you to meet your infrastructure and compliance requirements while staying completely focused on your application.
- **Ensure a compliant environment:** With Amazon VPC support on App Runner, you can easily connect to database, cache, and message queue services on AWS to support

your applications on App Runner. No public subnets are required—helping you protect resources in your VPC.

94.2. Backup/Restore and Disaster Recovery

This requirement is not applicable for this service. For additional information beyond what is described herein, please refer to http://aws.amazon.com/documentation/.

94.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

94.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/apprunner/?id=docs_gateway
- **Service quotas:** https://docs.aws.amazon.com/general/latest/gr/aws_service_limits.html
- **Service FAQs:** https://aws.amazon.com/apprunner/faqs/?refid=ps_a134p000006gb2oaau&trkcampaign=acq_paid_search_brand

94.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/apprunner/?id=docs_gateway and the following links for comprehensive technical documentation regarding this service.

- Developer Guide: Provides a conceptual overview of AWS App Runner, detailed feature descriptions, and instructions on how to use the service and deploy web applications.
- API Reference: Describes all the API actions and data types for AWS App Runner and provides call examples.
- App Runner section of AWS CLI Reference: Documents the AWS App Runner commands available in the AWS Command Line Interface (AWS CLI).
- Release Notes: Provides details about AWS App Runner releases—new features, updates, and fixes related to the service, runtimes, and console.

95. AWS Application Discovery Service

95.1. Service Overview

AWS Application Discovery Service helps enterprise customers plan migration projects by gathering information about their on-premises data centres.

Planning data centre migrations can involve thousands of workloads that are often deeply interdependent. Server utilization data and dependency mapping are important early first steps in the migration process. AWS Application Discovery Service collects and presents configuration, usage, and behaviour data from your servers to help you better understand your workloads.

The collected data is retained in encrypted format in an AWS Application Discovery Service data store. You can export this data as a CSV file and use it to estimate the Total Cost of Ownership (TCO) of running on AWS and to plan your migration to AWS. In addition, this data is

also available in AWS Migration Hub, where you can migrate the discovered servers and track their progress as they get migrated to AWS.

95.1.1. Features

- **Discover on-premises infrastructure:** AWS Application Discovery Service collects server hostnames, IP addresses, MAC addresses, and resource allocation and utilization details of key resources including CPU, network, memory, and disk. This information can then be used to size AWS resources when you migrate.
- **Identify server dependencies:** AWS Application Discovery Service agents record inbound and outbound network activity for each server. This data can then be used to understand the dependencies across servers.
- **Measure server performance:** AWS Application Discovery Service captures performance information about applications and processes by measuring host CPU, memory, and disk use, as well as disk and network performance (e.g., latency and throughput). This information lets you establish a performance baseline to use as a comparison after you migrate to AWS.
- **Data Exploration in Amazon Athena:** Explore the data collected from your on-premises servers with Amazon Athena by running pre-defined queries to analyse the time-series system performance for each server, the type of processes that are running on them and the network dependencies between different servers.

95.1.2. Benefits

- **Reliable Discovery for Migration Planning:** AWS Application Discovery Service collects server specification information, performance data, and details of running processes and network connections. This data can be used to perform a detailed cost estimate in advance of migrating to AWS, or to group servers into applications for planning purposes.
- **Integrated with Migration Hub:** AWS Application Discovery Service is integrated with AWS Migration Hub, which simplifies your migration tracking. After performing discovery and grouping your servers as applications, you can use Migration Hub to track the status of migrations across your application portfolio.
- **Protect Data with Encryption:** AWS Application Discovery Service provides protection for the collected data by encrypting it both in transit to AWS and at rest within the Application Discovery Service data store.
- **Engage with Migration Experts:** AWS Professional Services and APN Migration Partners have helped many enterprise customers successfully complete their migration to the cloud. These professionals are trained to analyse Application Discovery Service output and can help you gain further insights about your on-premises environment and recommend appropriate migration strategies.

95.2. Backup/Restore and Disaster Recovery

This requirement is not applicable for this service. For additional information beyond what is described herein, please refer to http://aws.amazon.com/documentation/.

95.3. Pricing Overview

You only pay for the AWS resources (e.g., Amazon S3, Amazon Athena, or Amazon Kinesis Firehose) that are provisioned to store your on-premises data. You only pay for what you use, as you use it; there are no minimum fees and no upfront commitments.

95.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/application-discovery/
- **Service quotas:** https://docs.aws.amazon.com/application-discovery/latest/userguide/ads_service_limits.html
- **Service FAQs:** https://aws.amazon.com/application-discovery/faqs/

95.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/application-discovery/ and the following links for comprehensive technical documentation regarding this service.

- User Guide: Walks through how to set up the AWS Application Discovery Service and integrate it with other services.

96. AWS Application Migration Service

96.1. Service Overview

AWS Application Migration Service (AWS MGN) allows you to quickly realize the benefits of migrating applications to the cloud without changes and with minimal downtime.

AWS Application Migration Service minimizes time-intensive, error-prone manual processes by automatically converting your source servers from physical, virtual, or cloud infrastructure to run natively on AWS. It further simplifies your migration by enabling you to use the same automated process for a wide range of applications.

And by launching non-disruptive tests before migrating, you can be confident that your most critical applications such as SAP, Oracle, and SQL Server will work seamlessly on AWS.

96.1.1. Features

- Operate the service from the AWS Management Console.
- Control permissions and access using AWS Identity and Access Management (IAM).
- Operate the service without a connection to the public internet.
- Store your migration metadata in the same AWS Region as your migrated instances.
- Utilize an agentless replication option (for vCenter), if needed.
- Use a new API that is better suited for migration-specific workflows, as well as a CLI and SDKs.
- Use Amazon CloudWatch and AWS CloudTrail to monitor AWS Application Migration Service.
- Better control how your test and cutover instances are launched using Amazon EC2 launch templates (rather than Blueprints).
- Use tags to organize your source servers and control access permissions.

96.1.2. Benefits

- **Access to advanced technology:** Simplify operations and get better insights with AWS Application Migration Service's integration with AWS Identity and Access Management (IAM), Amazon CloudWatch, AWS CloudTrail, and other AWS services.
- **Minimal downtime during migration:** With AWS Application Migration Service, you can maintain normal business operations throughout the replication process.
  It continuously replicates source servers, which means little to no performance impact. Continuous replication also makes it easy to conduct non-disruptive tests and shortens cutover windows.
- **Reduced costs:** AWS Application Migration Service reduces overall migration costs as there is no need to invest in multiple migration solutions, specialized cloud development, or application-specific skills. This is because it can be used to lift and shift any application from any source infrastructure that runs supported operating systems (OS).

96.2. Backup/Restore and Disaster Recovery

This requirement is not applicable for this service. For additional information beyond what is described herein, please refer to http://aws.amazon.com/documentation/.

96.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

96.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/mgn/?id=docs_gateway
- **Service quotas:** https://docs.aws.amazon.com/general/latest/gr/mgn.html
- **Service FAQs:** https://aws.amazon.com/application-migration-service/faqs/

96.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/mgn/?id=docs_gateway and the following links for comprehensive technical documentation regarding this service.

- User Guide: Learn how to set up and use AWS Application Migration Service.
- API Reference: Describes all the API operations for AWS Application Migration Service in detail.

97. AWS AppSync

97.1. Service Overview

AWS AppSync is a fully managed service that makes it easy to develop GraphQL APIs by handling the heavy lifting of securely connecting to data sources like AWS DynamoDB, Lambda, and more. Adding caches to improve performance, subscriptions to support real-time updates, and client-side data stores that keep off-line clients in sync are just as easy. Once deployed, AWS AppSync automatically scales your GraphQL API execution engine up and down to meet API request volumes.

97.1.1. Features

- **Simplified Data Access and Querying:** AWS AppSync uses GraphQL, a data language that enables client apps to fetch, change and subscribe to data from servers. In a GraphQL query, the client specifies how the data is to be structured when it is returned by the server. This makes it possible for the client to query only for the data it needs, in the format that it needs it in. GraphQL also includes a feature called “introspection” which lets new developers on a project discover the data available without requiring knowledge of the backend.
- **Immediate updates across clients and devices:** AWS AppSync lets you specify which portions of your data should be available in a real-time manner using GraphQL Subscriptions. GraphQL Subscriptions are simple statements in the application code that tell the service what data should be updated in real-time.
- **Interact with and update your data, even when offline:** The Amplify DataStore provides a queryable on-device DataStore for web, mobile and IoT developers with a local-first and familiar programming model to interact with data seamlessly whether you’re online or offline. When combined with AWS AppSync the DataStore can leverage advanced versioning, conflict detection and resolution in the cloud allowing to automatically merge data from different clients as well as providing data consistency and integrity.
- **Preconfigured access to AWS data sources**: AWS AppSync gives client applications the ability to specify data requirements with GraphQL so that only the needed data is fetched, allowing for both server and client filtering. Since AWS AppSync supports AWS Lambda, Amazon DynamoDB and Amazon Elasticsearch, the GraphQL operations can be simple lookups, complex queries & mappings, full text searches, fuzzy/keyword searches or geo lookups.
- **Cache your data that doesn't change frequently for improved performance**: AWS AppSync’s server-side data caching capabilities reduce the need to directly access data sources by making data available in high speed in-memory managed caches, delivering data at low latency. Being fully managed, it eliminates the operational overhead of managing cache clusters. By providing the flexibility to selectively cache data fields and operations defined in the GraphQL schema with customizable expiration, data caching further enables developers to configure optimal performance for their business needs.
- **Enterprise security and fine-grained access control:** AWS AppSync allows several levels of data access and authorization depending on the needs of an application. Simple access can be protected by a key and more restrictive permission can be done with AWS Identity and Access Management using Roles. Additionally, AWS AppSync integrates with Amazon Cognito User Pools for email and password functionality, social providers (Facebook, Google+, and Login with Amazon), and enterprise federation with SAML. Customers can use the Group functionality for logical organization of users and roles as well as OAuth features for application access.
- **Use your own domain name to access a GraphQL endpoints:** AWS AppSync enables customers to use custom domain names with their AWS AppSync API to access their GraphQl endpoint and real-time endpoint. To create a custom domain name in AppSync, you simply provide a domain name you own and indicate a valid AWS Certificate Manager (ACM) certificate that covers your domain. Once the custom domain name is created, you can associate the domain name with any available AppSync API in your account. After you have updated your DNS record to map to to the AppSync-provided domain name, you can configure your applications to use the new GraphQL

and real-time endpoints. You can change the API association on your custom domain at any time without having to update your applications. When AppSync receives a request on the custom domain endpoint, it routes it to the associated API for handling.

97.1.2. Benefits

- **Simple & secure data access:** Power your applications with the right data, from one or more data sources with a single network request using GraphQL. AWS AppSync makes it easy to secure your app data using multiple concurrent authentication modes and allows you to define security, caching and fine-grained access control at the data definition level directly from your GraphQL schema.
- **Built-in real-time & offline capabilities:** With managed GraphQL subscriptions, AWS AppSync can push real-time data updates over Websockets to millions of clients. For mobile and web applications, AppSync also provides local data access when devices go offline, and data synchronization with customizable conflict resolution, when they are back online.
- **No servers to manage:** AWS AppSync offers fully managed GraphQL API setup, administration, and maintenance, with high-availability serverless infrastructure built in. Create a GraphQL API in seconds via the AWS CLI, console, Amplify CLI or CloudFormation. And, easily monitor and acquire insights via CloudWatch and X-Ray for flawless operation.

97.2. Backup/Restore and Disaster Recovery

This requirement is not applicable for this service. For additional information beyond what is described herein, please refer to http://aws.amazon.com/documentation/.

97.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

97.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/appsync/
- **Service quotas:** https://docs.aws.amazon.com/general/latest/gr/appsync.html
- **Service FAQs:** https://aws.amazon.com/appsync/faqs/

97.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/appsync/ and the following links for comprehensive technical documentation regarding this service.

- Developer Guide: Create and interact with data sources using GraphQL from your application. You can build a new application or integrate existing data sources with AWS AppSync.
- API Reference: Describes all the API operations for AWS AppSync in detail. Also provides sample requests, responses, and errors for the supported web service protocols.

98. AWS Artifact

98.1. Service Overview

AWS Artifact is your go-to, central resource for compliance-related information that matters to you. It provides on-demand access to AWS' security and compliance reports and select online agreements. Reports available in AWS Artifact include our Service Organization Control (SOC) reports, Payment Card Industry (PCI) reports, and certifications from accreditation bodies across geographies and compliance verticals that validate the implementation and operating effectiveness of AWS security controls. Agreements available in AWS Artifact include the Business Associate Addendum (BAA) and the Nondisclosure Agreement (NDA).

98.1.1. Features

- **AWS Artifact Reports:** AWS Artifact Reports provides several compliance reports from third-party auditors who have tested and verified our compliance with a variety of global, regional, and industry specific security standards and regulations. When new reports are released, they are made available in AWS Artifact.
- **AWS Artifact Agreements:** AWS Artifact Agreements enables you to review, accept, and manage agreements with AWS for an individual account, and for all accounts that are part of your organization in AWS Organizations. You can also use AWS Artifact to terminate agreements you have previously accepted if they are no longer required.

98.1.2. Benefits

- **Comprehensive resource:** Access all of AWS' auditor issued reports, certifications, accreditations and other third-party attestations.
- **Agreement governance:** Review, accept, and manage your agreements with AWS. Apply your AWS agreements to all current and future accounts within your organization.
- **Deep insights:** Perform due-diligence of AWS with enhanced transparency into our security control environment. Continuously monitor the security and compliance of AWS with immediate access to new reports.

98.2. Backup/Restore and Disaster Recovery

This requirement is not applicable for this service. For additional information beyond what is described herein, please refer to http://aws.amazon.com/documentation/.

98.3. Pricing Overview

This is a no cost, self-service portal for on-demand access to AWS' compliance reports.

98.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/artifact/index.html
- **Service FAQs:** https://aws.amazon.com/artifact/faq/

98.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/artifact/index.html for comprehensive technical documentation regarding this service.

- User Guide: Describes key concepts of AWS Artifact and provides instructions for using the features of AWS Artifact.

99. AWS Audit Manager

99.1. Service Overview

AWS Audit Manager helps you continuously audit your AWS usage to simplify how you assess risk and compliance with regulations and industry standards. Audit Manager automates evidence collection to reduce the “all hands on deck” manual effort that often happens for audits and enable you to scale your audit capability in the cloud as your business grows. With Audit Manager, it is easy to assess if your policies, procedures, and activities – also known as controls – are operating effectively. When it is time for an audit, AWS Audit Manager helps you manage stakeholder reviews of your controls and enables you to build audit-ready reports with much less manual effort.

AWS Audit Manager’s prebuilt frameworks help translate evidence from cloud services into auditor-friendly reports by mapping your AWS resources to the requirements in industry standards or regulations, such as CIS AWS Foundations Benchmark, the General Data Protection Regulation (GDPR), and the Payment Card Industry Data Security Standard (PCI DSS). You can also fully customize a framework and its controls for your unique business requirements. Based on the framework you select, Audit Manager launches an assessment that continuously collects and organizes relevant evidence from your AWS accounts and resources, such as resource configuration snapshots, user activity, and compliance check results.

You can get started quickly in the AWS Management Console. Just select a prebuilt framework to launch an assessment and begin automatically collecting and organizing evidence.

99.1.1. Features

- **Prebuilt frameworks:** AWS Audit Manager offers prebuilt frameworks that cover a range of compliance standards, and they are developed with AWS best practices in mind. These frameworks help map your AWS resources to the requirements for industry standards and regulations.
- **Custom frameworks and controls:** AWS Audit Manager enables you to build your own framework using either custom controls or AWS-managed controls which help you meet your audit requirements. Customizing an Audit Manager framework helps you evaluate controls in your existing framework for compliance with your particular business requirements.
- **Automated evidence collection:** Once an assessment has been defined and launched, AWS Audit Manager automatically collects data for the AWS account and services you have defined to be in scope for an audit. The evidence contains both the data captured from that resource as well as metadata that indicates which control the data supports to help you demonstrate security, change management, business continuity, and software licensing compliance
- **Multi-account evidence collection:** AWS Audit Manager supports multiple accounts via integration with AWS Organizations. Audit Manager assessments can run over multiple accounts and will collect and consolidate evidence into a delegated administrator account in AWS Organizations.
- **Delegation workflow:** You can delegate control sets to team members who are specialized in certain topic areas, such as network infrastructure, identity management,

software licensing, or personnel policies. The delegation feature enables the support team members to review the control set and related evidence, add comments, upload additional evidence, and update the status of each control.

- **Audit-ready reports:** AWS Audit Manager automates evidence collect and organizes the evidence as defined by the control set in the framework you selected. You and your team can review evidence, comment on evidence, upload other supporting evidence, and update the status of each control. You then select the relevant evidence to include in your assessment report and generate a final assessment report to share with your auditors.

99.1.2. Benefits

- **Easily map your AWS usage to controls:** AWS Audit Manager provides prebuilt frameworks that map your AWS resources to control requirements, which are grouped in accordance to the requirements of an industry standard or regulation, such as CIS AWS Foundations Benchmarks, GDPR, or PCI DSS. You can fully customize these prebuilt frameworks and controls to tailor them to your unique needs.
- **Save time with automated collection of evidence:** AWS Audit Manager saves you time by automatically collecting and organizing evidence as defined by each control requirement. Instead of manually collecting evidence, you can focus on reviewing the relevant evidence to ensure your controls are working as intended. For example, you can configure an Audit Manager assessment to automatically collect resource configuration snapshots on a daily, weekly, or monthly basis, subject to underlying AWS service configurations.
- **Streamline collaboration across teams:** You can assign controls in your assessment to a subject matter expert to review. For example, you might delegate a security control to a network security engineer to confirm the evidence properly demonstrates compliance. Audit Manager also allows team members to comment on evidence, upload manual evidence, and update the status of each control.
- **Always be prepared to produce audit-ready reports:** An audit-ready report includes a report summary file that contains links to the relevant evidence folders, which are named and organized according to the controls that are specified in your assessment. The evidence Audit Manager continuously collects from the AWS services you use becomes a record containing the information needed to demonstrate compliance with control requirements. You can review and select the relevant evidence to include in your final audit report.
- **Ensure assessment report and evidence integrity:** AWS Audit Manager securely stores evidence in its own managed storage repository with read-only permissions to your end-users. When you generate audit-ready reports, Audit Manager produces a report file checksum so you can validate that the report evidence remains unaltered. Both the summary report and evidence can be downloaded to share with your auditors.

99.2. Backup/Restore and Disaster Recovery

This requirement is not applicable for this service. For additional information beyond what is described herein, please refer to http://aws.amazon.com/documentation/.

99.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

99.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/audit-manager/
- **Service quotas:** https://docs.aws.amazon.com/audit-manager/latest/userguide/service-quotas.html
- **Service FAQs:** https://aws.amazon.com/audit-manager/faqs/

99.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/audit-manager/ and the following links for comprehensive technical documentation regarding this service.

- User Guide: Describes the key concepts of AWS Audit Manager and provides detailed instructions for using the various features.
- AWS CloudFormation User Guide for AWS Audit Manager: Documents the reference information for all AWS Audit Manager resource and property types that are supported by AWS CloudFormation.

100. AWS Auto Scaling

100.1. Service Overview

AWS Auto Scaling monitors your applications and automatically adjusts capacity to maintain steady, predictable performance at the lowest possible cost. Using AWS Auto Scaling, it's easy to setup application scaling for multiple resources across multiple services in minutes. The service provides a simple, powerful user interface that lets you build scaling plans for resources including Amazon EC2 instances and Spot Fleets, Amazon ECS tasks, Amazon DynamoDB tables and indexes, and Amazon Aurora Replicas. AWS Auto Scaling makes scaling simple with recommendations that allow you to optimize performance, costs, or balance between them. If you're already using Amazon EC2 Auto Scaling to dynamically scale your Amazon EC2 instances, you can now combine it with AWS Auto Scaling to scale additional resources for other AWS services. With AWS Auto Scaling, your applications always have the right resources at the right time.

100.1.1. Features

Copy and paste features from service features page

- **Unified scaling:** Using AWS Auto Scaling, you can configure automatic scaling for all of the scalable resources powering your application from a single unified interface, including:
  - **Amazon EC2**: Launch or terminate Amazon EC2 instances in an Amazon EC2 Auto Scaling group.
  - **Amazon EC2 Spot Fleets**: Launch or terminate instances from an Amazon EC2 Spot Fleet, or automatically replace instances that get interrupted for price or capacity reasons.

- **Amazon ECS**: Adjust ECS service desired count up or down to respond to load variations.
- **Amazon DynamoDB**: Enable a DynamoDB table or a global secondary index to increase its provisioned read and write capacity to handle sudden increases in traffic without throttling.
- **Amazon Aurora**: Dynamically adjust the number of Aurora Read Replicas provisioned for an Aurora DB cluster to handle sudden increases in active connections or workload.

- **Automatic resource discovery:** AWS Auto Scaling scans your environment and automatically discovers the scalable cloud resources underlying your application, so you don't have to manually identify these resources one by one through individual service interfaces.
- **Built-in scaling strategies:** Using AWS Auto Scaling, you can select one of three predefined optimization strategies designed to optimize performance, optimize costs, or balance the two. If you prefer, you can set your own target resource utilization. Using your selected scaling strategy, AWS Auto Scaling will create the scaling policies for each of your resources for you.
- **Predictive Scaling:** Predictive Scaling predicts future traffic, including regularly-occurring spikes, and provisions the right number of EC2 instances in advance of predicted changes. Predictive Scaling's machine learning algorithms detect changes in daily and weekly patterns, automatically adjusting their forecasts. This removes the need for manual adjustment of Auto Scaling parameters over time, making Auto Scaling simpler to configure and consume. Auto Scaling enhanced with Predictive Scaling delivers faster, simpler, and more accurate capacity provisioning resulting in lower cost and more responsive applications.
- **Fully-managed:** AWS Auto Scaling automatically creates target tracking scaling policies for all of the resources in your scaling plan, using your selected scaling strategy to set the target values for each metric. AWS Auto Scaling also creates and manages the Amazon CloudWatch alarms that trigger scaling adjustments for each of your resources.
- **Smart scaling policies:** AWS Auto Scaling continually calculates the appropriate scaling adjustments and immediately adds and removes capacity as needed to keep your metrics on target. AWS target tracking scaling policies are self-optimizing, and learn your actual load patterns to minimize fluctuations in resource capacity. This results in smoother, smarter scaling and you pay only for the resources you actually need.

100.1.2. Benefits

- **Setup scaling quickly:** AWS Auto Scaling lets you set target utilization levels for multiple resources in a single, intuitive interface. You can quickly see the average utilization of all of your scalable resources without having to navigate to other consoles. For example, if your application uses Amazon EC2 and Amazon DynamoDB, you can use AWS Auto Scaling to manage resource provisioning for all of the EC2 Auto Scaling groups and database tables in your application.
- **Make Smart Scaling Decisions:** AWS Auto Scaling lets you build scaling plans that automate how groups of different resources respond to changes in demand. You can optimize availability, costs, or a balance of both. AWS Auto Scaling automatically creates all of the scaling policies and sets targets for you based on your preference. AWS Auto Scaling monitors your application and automatically adds or removes capacity from your resource groups in real-time as demands change.

- **Automatically Maintain Performance:** Using AWS Auto Scaling, you maintain optimal application performance and availability, even when workloads are periodic, unpredictable, or continuously changing. AWS Auto Scaling continually monitors your applications to make sure that they are operating at your desired performance levels. When demand spikes, AWS Auto Scaling automatically increases the capacity of constrained resources so you maintain a high quality of service.
- **Pay Only for What You Need:** AWS Auto Scaling can help you optimize your utilization and cost efficiencies when consuming AWS services so you only pay for the resources you actually need. When demand drops, AWS Auto Scaling will automatically remove any excess resource capacity so you avoid overspending. AWS Auto Scaling is free to use, and allows you to optimize the costs of your AWS environment.

100.2. Backup/Restore and Disaster Recovery

This requirement is not applicable for this service. For additional information beyond what is described herein, please refer to http://aws.amazon.com/documentation/.

100.3. Pricing Overview

There is no additional charge for AWS Auto Scaling. You pay only for the AWS resources needed to run your applications and Amazon CloudWatch monitoring fees.

100.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/autoscaling/index.html
- **Service quotas:** https://docs.aws.amazon.com/autoscaling/plans/userguide/scaling-plan-quotas.html
- **Service FAQs:** https://aws.amazon.com/autoscaling/faqs/

100.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/autoscaling/index.html and the following links for comprehensive technical documentation regarding this service.

- User Guide: Walks you through using the AWS Auto Scaling console to identify your scalable resources and to create a scaling plan for the first time.
- AWS Auto Scaling Plans section of the AWS CLI Reference: Provides syntax and examples for the AWS CLI commands for use with scaling plans.
- API Reference: Describes all the API operations for scaling plans in detail. Also provides sample requests, responses, and errors for the supported web services protocols.

101. AWS Backup

101.1. Service Overview

AWS Backup enables you to centralize and automate data protection across AWS services and hybrid workloads. AWS Backup offers a cost-effective, fully managed, policy-based service that further simplifies data protection at scale. AWS Backup also helps you support your regulatory compliance or business policies for data protection. Together with AWS Organizations, AWS Backup enables you to centrally deploy data protection policies to configure, manage, and govern your backup activity across your organization's AWS accounts and resources, including

Amazon Elastic Compute Cloud (Amazon EC2) instances, Amazon Elastic Block Store (Amazon EBS) volumes, Amazon Relational Database Service (Amazon RDS) databases (including Amazon Aurora clusters), Amazon DynamoDB tables, Amazon Neptune databases, Amazon DocumentDB (with MongoDB compatibility) databases, Amazon Elastic File System (Amazon EFS) file systems, Amazon FSx for Lustre file systems, Amazon FSx for Windows File Server file systems, and AWS Storage Gateway volumes, and VMware workloads on premises and in VMware CloudTM on AWS. AWS Backup also offers a preview for backup and restore of Amazon Simple Storage Service (Amazon S3) buckets.

101.1.1. Features

- **Centralized backup management:** AWS Backup provides a centralized backup console, a set of backup APIs, and a command line interface to manage backups across the AWS services that your applications run on, including Amazon Elastic Block Store (EBS), Amazon FSx, Amazon Elastic Compute Cloud (EC2), Amazon Relational Database Service (RDS), Amazon DynamoDB, Amazon Elastic File System (EFS), AWS Storage Gateway, Amazon Neptune, Amazon DocumentDB (with MongoDB compatibility). as well hybrid applications like VMware workloads running on premises and in VMware Cloud™ on AWS. AWS Backup also provides a preview of AWS Backup for Amazon Simple Storage Service (S3). With AWS Backup, you can centrally manage backup policies that meet your backup requirements and apply them to your AWS resources across AWS services and hybrid cloud workloads, enabling you to back up your application data in a consistent and compliant manner. The AWS Backup centralized backup console offers a consolidated view of your backups and backup activity logs, making it easier to audit your backups and ensure compliance.
- **Policy-based backup solution:** With AWS Backup, you can create backup policies called backup plans that enable you to define your backup requirements and then apply them to the AWS resources you want backed up. You can create separate backup plans that meet specific business and regulatory compliance requirements, helping to ensure that each of your AWS resources is backed up and protected. Backup plans make it easy to implement your backup strategy across your organization and across your applications.
- **Tag-based backup policies:** AWS Backup allows you to apply backup plans to your AWS resources by simply tagging them, making it easier to implement your backup strategy across all your applications and ensure that all your AWS resources are backed up and protected. AWS tags are a great way to organize and classify your AWS resources. Integration with AWS tags enables you to quickly apply a backup plan to a group of AWS resources so that they are backed up in a consistent and compliant manner.
- **Automated backup scheduling:** AWS Backup allows you to create backup schedules that you can customize to meet your business and regulatory backup requirements. You can also choose from predefined backup schedules based on common best practices. AWS Backup will automatically back up your AWS resources according to the policies and schedules you define. A backup schedule includes the backup start time, backup frequency, and backup window.
- **Automated retention management:** With AWS Backup, you can set backup retention policies that will automatically retain and expire backups according to your business and regulatory backup compliance requirements. Automated backup retention management makes it easy to minimize backup storage costs by retaining backups for only as long as they are needed.

- **Backup activity monitoring:** AWS Backup provides a dashboard that makes it simple to monitor backup and restore activity across AWS services. With just a few clicks in the AWS Backup console, you can view the status of recent backup jobs and restore jobs across AWS services to ensure that your AWS resources are properly protected. AWS Backup integrates with AWS CloudTrail, which provides you with a consolidated view of backup activity logs that make it quick and easy to audit resources are backed up and how. AWS Backup also integrates with Amazon Simple Notification Service (SNS), which can automatically alert you on backup activity, such as when a backup succeeds or a restore has been initiated.
- **AWS Backup Audit Manager:** AWS Backup Audit Manager allows you to audit and report on the compliance of your data protection policies to help you meet your business and regulatory needs. AWS Backup Audit Manager provides built-in compliance controls and allows you to customize these controls to define your data protection policies (such as backup frequency or retention period). It is designed to automatically detect violations of your defined data protection policies and will prompt you to take corrective actions. With AWS Backup Audit Manager, you can continuously evaluate backup activity and generate audit reports that can help you demonstrate compliance with regulatory requirements.
- **AWS Backup Vault Lock:** AWS Backup Vault Lock allows you to protect your backups from deletion or changes to their lifecycle by inadvertent or malicious changes. You can use the AWS CLI, AWS Backup API, or AWS Backup SDK to apply the AWS Backup Vault Lock protection to an existing vault or a new one. AWS Backup Vault Lock works seamlessly with backup policies such as retention periods, cold storage transitioning, cross-account, and cross-Region copy, providing you an additional layer of protection and helping you meet your compliance requirements. While AWS Backup Vault Lock helps you implement safeguards that ensure you are storing your backups using a Write-Once-Read-Many (WORM) model, the feature has not yet been assessed for compliance with the Securities and Exchange Commission (SEC) rule 17a-4(f) and the Commodity Futures Trading Commission (CFTC) in regulation 17 C.F.R. 1.31(b)-(c).
- **Lifecycle management policies:** AWS Backup enables you to meet compliance requirements while minimizing backup storage costs by storing backups in a low-cost cold storage tier. You can configure lifecycle policies that will automatically transition backups from warm storage to cold storage according to a schedule that you define.
- **Incremental backups:** AWS Backup efficiently stores your periodic backups incrementally. The first backup of an AWS resource backs up a full copy of your data. For each successive incremental backup, only the changes to your AWS resources are backed up. Incremental backups enable you to benefit from the data protection of frequent backups while minimizing storage costs. Currently, Amazon DynamoDB, Amazon Aurora, Amazon DocumentDB (with MongoDB compatibility), and Amazon Neptune do not support incremental backups.

101.1.2. Benefits

- **Centrally manage backups:** Configure backup policies from a central backup console, simplifying backup management and making it easy to ensure that your application data across AWS services is backed up and protected. Use AWS Backup's central console, APIs, or command line interface to back up, restore, and set backup retention policies across AWS services.
- **Automate backup processes:** Save time and money with AWS Backup's fully managed, policy-based solution. AWS Backup provides automated backup schedules,

retention management, and lifecycle management, removing the need for custom scripts and manual processes. With AWS Backup, you can apply backup policies to your AWS resources by simply tagging them, making it easy to implement your backup strategy across all your AWS resources and ensuring that all your application data is appropriately backed up.

- **Improve backup compliance:** Enforce your backup policies, encrypt your backups, protect your backups from manual deletion and prevent changes to your backup lifecycle settings, and audit and report on backup activity from a centralized console to help meet your backup compliance requirements. Backup policies make it simple to align your backup strategy with your internal or regulatory requirements. AWS Backup secures your backups by encrypting your data in transit and at rest. Consolidated backup activity logs across AWS services make it easier to perform compliance audits. AWS Backup is PCI and ISO compliant as well as HIPAA eligible.

101.2. Backup/Restore and Disaster Recovery

AWS Backup is the backup service storing data resiliently on S3. S3 is designed to provide 99.999999999% durability of objects over a given year. However, data can also be copied between AWS regions using S3 cross region replication. more information can be found here, https://docs.aws.amazon.com/AmazonS3/latest/dev/crr.html.

101.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

101.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/aws-backup/index.html
- **Service quotas:** https://docs.aws.amazon.com/aws-backup/latest/devguide/aws-backup-limits.html
- **Service FAQs:** https://aws.amazon.com/backup/faqs/?nc=sn&loc=6&refid=ps_a134p000006gb41aae&trkcampaign=acq_paid_search_brand

101.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/aws-backup/index.html and the following links for comprehensive technical documentation regarding this service.

Developer Guide: Provides a conceptual overview of AWS Backup, detailed instructions for using the various features, and a complete API reference for developers.

102. AWS Batch

102.1. Service Overview

AWS Batch enables developers, scientists, and engineers to easily and efficiently run hundreds of thousands of batch computing jobs on AWS. AWS Batch dynamically provisions the optimal quantity and type of compute resources (e.g., CPU or memory optimized instances) based on the volume and specific resource requirements of the batch jobs submitted. With AWS Batch, there is no need to install and manage batch computing software or server clusters that you use

to run your jobs, allowing you to focus on analyzing results and solving problems. AWS Batch plans, schedules, and executes your batch computing workloads across the full range of AWS compute services and features, such as AWS Fargate, Amazon EC2 and Spot Instances.

102.1.1. Features

- **Dynamic compute resource provisioning and scaling:** When using Fargate or Fargate Spot with Batch, you only need to set up a few concepts in Batch (a CE, job queue, and job definition), and you have a complete queue, scheduler, and compute architecture without managing a single piece of compute infrastructure.
- **AWS Batch with Fargate:** AWS Batch with Fargate resources allows you to have a completely serverless architecture for your batch jobs. With Fargate, every job receives the exact amount of CPU and memory that it requests (within allowed Fargate SKU's), so there is no wasted resource time or need to wait for EC2 instance launches.
- **Support for tightly-coupled HPC workloads:** AWS Batch supports multi-node parallel jobs, which enables you to run single jobs that span multiple EC2 instances. This feature lets you use AWS Batch to easily and efficiently run workloads such as large-scale, tightly-coupled High Performance Computing (HPC) applications or distributed GPU model training.
- **Granular job definitions and simple job dependency modelling:** AWS Batch allows you to specify resource requirements, such as vCPU and memory, AWS Identity and Access Management (IAM) roles, volume mount points, container properties, and environment variables, to define how jobs are to be run. AWS Batch executes your jobs as containerized applications running on Amazon ECS. Batch also enables you to define dependencies between different jobs
- **Priority-based job scheduling:** AWS Batch enables you to set up multiple queues with different priority levels. Batch jobs are stored in the queues until compute resources are available to execute the job. The AWS Batch scheduler evaluates when, where, and how to run jobs that have been submitted to a queue based on the resource requirements of each job. The scheduler evaluates the priority of each queue and runs jobs in priority order on optimal compute resources (e.g., memory vs CPU optimized), as long as those jobs have no outstanding dependencies.
- **Support for GPU scheduling:** GPU scheduling allows you to specify the number and type of accelerators your jobs require as job definition input variables in AWS Batch. AWS Batch will scale up instances appropriate for your jobs based on the required number of GPUs and isolate the accelerators according to each job's needs, so only the appropriate containers can access them.
- **Support for popular workflow engines:** AWS Batch can be integrated with commercial and open-source workflow engines and languages such as Pegasus WMS, Luigi, Nextflow, Metaflow, Apache Airflow, and AWS Step Functions, enabling you to use familiar workflow languages to model your batch computing pipelines.
- **Integration with EC2 Launch Templates:** AWS Batch now supports EC2 Launch Templates, allowing you to build customized templates for your compute resources, and enabling Batch to scale instances with those requirements. You can specify your EC2 Launch Template to add storage volumes, specify network interfaces, or configure permissions, among other capabilities.
- **Flexible allocation strategies:** AWS Batch allows customers to choose mujltiple methods to allocate compute resources. These strategies allow customers to factor in

throughput as well as price when deciding how AWS Batch should scale instances on their behalf.

102.1.2. Benefits

- **Fully managed:** AWS Batch eliminates the need to operate third-party commercial or open source batch processing solutions. There is no batch software or servers to install or manage. AWS Batch manages all the infrastructure for you, avoiding the complexities of provisioning, managing, monitoring, and scaling your batch computing jobs.
- **Integrated with AWS:** AWS Batch is natively integrated with the AWS platform, allowing you to leverage the scaling, networking, and access management capabilities of AWS. This makes it easy to run jobs that safely and securely retrieve and write data to and from AWS data stores such as Amazon S3 or Amazon DynamoDB. You can also run AWS Batch on AWS Fargate, for fully serverless architecture, eliminating the need to manage compute infrastructure.
- **Cost optimized resource provisioning:** AWS Batch provisions compute resources and optimizes the job distribution based on the volume and resource requirements of the submitted batch jobs. AWS Batch dynamically scales compute resources to any quantity required to run your batch jobs, freeing you from the constraints of fixed-capacity clusters. AWS Batch will utilize Spot Instances or submit to Fargate Spot on your behalf, reducing the cost of running your batch jobs further.

102.2. Backup/Restore and Disaster Recovery

This requirement is not applicable for this service. For additional information beyond what is described herein, please refer to http://aws.amazon.com/documentation/.

102.3. Pricing Overview

There is no additional charge for AWS Batch. You pay for AWS resources (e.g. EC2 instances, AWS Lambda functions or AWS Fargate) you create to store and run your application. You can use your Reserved Instances, Savings Plan, EC2 Spot Instances, and Fargate with AWS Batch by specifying your compute-type requirements when setting up your AWS Batch compute environments. Discounts will be applied at billing time.

102.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/batch/index.html
- **Service quotas:** https://docs.aws.amazon.com/batch/latest/userguide/service_limits.html
- **Service FAQs:** https://aws.amazon.com/batch/faqs/?nc=sn&loc=5

102.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/batch/index.html and the following links for comprehensive technical documentation regarding this service.

- User Guide: Describes key concepts of AWS Batch and provides instructions for using the features of AWS Batch.
- API Reference: Describes all the API operations for AWS Batch in detail.

- AWS Batch section of AWS CLI Reference: Documents the AWS Batch commands available in the AWS Command Line Interface (AWS CLI).

103. AWS Budgets

103.1. Service Overview

AWS Budgets allows you to set custom budgets to track your cost and usage from the simplest to the most complex use cases. With AWS Budgets, you can choose to be alerted by email or SNS notification when actual or forecasted cost and usage exceed your budget threshold, or when your actual RI and Savings Plans' utilization or coverage drops below your desired threshold. With AWS Budget Actions, you can also configure specific actions to respond to cost and usage status in your accounts, so that if your cost or usage exceeds or is forecasted to exceed your threshold, actions can be executed automatically or with your approval to reduce unintentional over-spending.

AWS Budgets integrates with multiple other AWS services, such as AWS Cost Explorer, so you can easily view and analyse your cost and usage drivers, AWS Chatbot, so you can receive Budget alerts in your designated Slack channel or Amazon Chime room, and AWS Service Catalog, so you can track cost on your approved AWS portfolios and products.

103.1.1. Features

- **Monitor:** Monitor your AWS cost and usage, or RI and Savings Plans’ coverage and utilization. Set up your preferred budget period of day, month, quarter or year, and create specific budget limits that meet your business needs.
- **Report:** Stay informed of how actual or forecasted cost and usage, or actual RI and Savings Plans’ coverage and utilization, progresses towards your budget threshold, with scheduled reports without needing to log into the AWS console.
- **Respond:** Set up custom actions, such as Identity and Access Management (IAM) policies, Service Control Policies (SCPs), or target running instances (EC2 or RDS) that can be executed automatically or via a workflow approval process, when a budget target has been exceeded.

103.1.2. Benefits

- **Custom budgets that meet your needs:** Track your cost, usage, or coverage and utilization for your Reserved Instances and Savings Plans, across multiple dimensions, such as service, or Cost Categories. Aggregate your costs with an unblended or amortized view and include or exclude certain charges, such as tax and refunds. Configure your Budget Actions with IAM policies, Service Control Policies (SCPs), and targeted running instances.
- **Stay informed with alerts and reports:** Set up event-driven alert notifications for when actual or forecasted cost or usage exceeds your budget limit, or when your RI and Savings Plans' coverage or utilization drops below your threshold. You can also choose to be informed on a daily, weekly, or monthly basis with pre-scheduled Budgets Reports.
- **Granular budget time periods:** Create annual, quarterly, monthly, or even daily budgets depending on your business needs. This allows you to take timely actions to prevent cost or usage overage, or inefficient utilization or resource coverage of your Reserved Instances and Savings Plans.

103.2. Backup/Restore and Disaster Recovery

This requirement is not applicable for this service. For additional information beyond what is described herein, please refer to http://aws.amazon.com/documentation/.

103.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

103.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/account-billing/index.html
- **Service quotas:** https://docs.aws.amazon.com/awsaccountbilling/latest/aboutv2/billing-limits.html
- **Service FAQs:** https://aws.amazon.com/aws-cost-management/aws-budgets/faqs/

103.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/account-billing/index.html and the following links for comprehensive technical documentation regarding this service.

- AWS Billing User Guide: Describes key concepts of how to use the Billing console, integrated with the AWS Management Console.
- AWS Cost Management User Guide: Describes key concepts of the Cost Management console, and provides detailed instructions for using the various features.
- Cost and Usage Reports User Guide: Describes how to use the AWS Cost and Usage Reports feature, integrated with the AWS Billing and Cost Management console.
- API Reference for AWS Billing and Cost Management: Describes all the API operations for AWS Billing and Cost Management in detail. Also provides sample requests, responses, and errors for the supported web services protocols.

104. AWS Certificate Manager

104.1. Service Overview

AWS Certificate Manager is a service that lets you easily provision, manage, and deploy public and private Secure Sockets Layer/Transport Layer Security (SSL/TLS) certificates for use with AWS services and your internal connected resources. SSL/TLS certificates are used to secure network communications and establish the identity of websites over the Internet as well as resources on private networks. AWS Certificate Manager removes the time-consuming manual process of purchasing, uploading, and renewing SSL/TLS certificates.

With AWS Certificate Manager, you can quickly request a certificate, deploy it on ACM-integrated AWS resources, such as Elastic Load Balancers, Amazon CloudFront distributions, and APIs on API Gateway, and let AWS Certificate Manager handle certificate renewals. It also enables you to create private certificates for your internal resources and manage the certificate lifecycle centrally. Public and private certificates provisioned through AWS Certificate Manager for use with ACM-integrated services are free. You pay only for the AWS resources you create to run your application. With AWS Certificate Manager Private Certificate Authority, you pay monthly for the operation of the private CA and for the private certificates you issue.

104.1.1. Features

- **Centrally manage certificates on the AWS Cloud:** You will find it easy to centrally manage AWS Certificate Manager SSL/TLS certificates provided by AWS Certificate Manager in an AWS Region from the AWS Management Console, AWS CLI, or AWS Certificate Manager APIs. You can also audit the use of each certificate by reviewing your Amazon CloudTrail logs.
- **Private certificate authority:** AWS Certificate Manager (ACM) Private Certificate Authority (CA) is a managed private CA service that helps you easily and securely manage the lifecycle of your private certificates. ACM Private CA provides you a highly-available private CA service without the upfront investment and ongoing maintenance costs of operating your own private CA or private CA hierarchy. Learn more about ACM Private CA.
- **Secure key management:** AWS Certificate Manager is designed to protect and manage the private keys used with SSL/TLS certificates. Strong encryption and key management best practices are used when protecting and storing private keys.
- **Integrated with other AWS cloud services:** AWS Certificate Manager is integrated with other AWS services, so you can provision an SSL/TLS certificate and deploy it with your Elastic Load Balancer, Amazon CloudFront distribution or API in Amazon API Gateway. AWS Certificate Manager also works with AWS Elastic Beanstalk and AWS CloudFormation for public email-validated certificates to help you manage public certificates and use them with your applications in the AWS Cloud. To deploy a certificate with an AWS resource, you simply select the certificate you want from a drop-down list in the AWS Management Console. Alternatively, you can call an AWS API or CLI to associate the certificate with your resource. AWS Certificate Manager then deploys the certificate to the selected resource for you.
- **Import third-party certificates:** AWS Certificate Manager makes it easy to import SSL/TLS certificates issued by third-party Certificate Authorities (CAs) and deploy them with your Elastic Load Balancers, Amazon CloudFront distributions and APIs on Amazon API Gateway. You can monitor the expiration date of an imported certificate and import a replacement when the existing certificate is nearing expiration. Alternatively, you can request a free certificate from AWS Certificate Manager and let AWS manage future renewals for you. Importing certificates doesn't cost anything.

104.1.2. Benefits

- **Free public certificates for ACM-integrated services:** With AWS Certificate Manager, there is no additional charge for provisioning public or private SSL/TLS certificates you use with ACM-integrated services, such as Elastic Load Balancing and API Gateway. You pay for the AWS resources you create to run your application. For private certificates, ACM Private CA provides you the ability to pay monthly for the service and certificates you create. You pay less per certificate as you create more private certificates.
- **Managed certificate renewal:** AWS Certificate Manager manages the renewal process for the certificates managed in ACM and used with ACM-integrated services, such as Elastic Load Balancing and API Gateway. ACM can automate renewal and deployment of these certificates. With ACM Private CA APIs, ACM enables you to automate creation and renewal of private certificates for on-premises resources, EC2 instances, and IoT devices.

- **Get certificates easily:** AWS Certificate Manager removes many of the time-consuming and error-prone steps to acquire an SSL/TLS certificate for your website or application. There is no need to generate a key pair or certificate signing request (CSR), submit a CSR to a Certificate Authority, or upload and install the certificate once received. With a few clicks in the AWS Management Console, you can request a trusted SSL/TLS certificate from AWS. Once the certificate is created, AWS Certificate Manager takes care of deploying certificates to help you enable SSL/TLS for your website or application.

104.2. Backup/Restore and Disaster Recovery

This requirement is not applicable for this service. For additional information beyond what is described herein, please refer to http://aws.amazon.com/documentation/.

104.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

104.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/acm/index.html
- **Service quotas:** https://docs.aws.amazon.com/acm/latest/userguide/acm-limits.html
- **Service FAQs:** https://aws.amazon.com/certificate-manager/faqs/?nc=sn&loc=5

104.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/acm/index.html and the following links for comprehensive technical documentation regarding this service.

- User Guide: Provides conceptual overviews and explains how to provision ACM certificates on AWS based websites.
- API Reference: Describes the API operations available for ACM along with sample requests, responses, and errors for the supported web services protocols.
- ACM in the AWS CLI Reference: Describes the AWS Certificate Manager commands that are available in the AWS Command Line Interface.
- User Guide: Provides conceptual overviews and explains how to create a private certificate authority.
- API Reference: Describes the API operations available for ACM Private CA along with sample requests, responses, and errors for the supported web services protocols.
- ACM Private CA in the AWS CLI Reference: Describes the ACM Private CA commands that are available in the AWS Command Line Interface.

105. AWS Chatbot

105.1. Service Overview

AWS Chatbot is an interactive agent that makes it easy to monitor, operate, and troubleshoot your AWS workloads in your chat channels. With AWS Chatbot, you can receive alerts, run commands to retrieve diagnostic information, configure AWS resources, and initiate workflows.

With just a few clicks, you can receive AWS notifications and run AWS Command Line Interface (CLI) commands from your chat channels in a secure and efficient manner. AWS Chatbot manages the integration and security permissions between the AWS services and your Slack channels or Amazon Chime chatrooms. AWS Chatbot makes it easier for your team to stay updated, collaborate, and respond quickly to incidents, security findings, and other alerts for applications running in your AWS environment. Your team can run commands to safely configure AWS resources, resolve incidents, and run tasks from Slack channels without switching context to other AWS Management Tools.

105.1.1. Features

- **Receive notifications:** Simply choose the Slack channels or Amazon Chime chatrooms you want to receive notifications in, and then choose the Amazon Simple Notification Service (SNS) topics that should trigger notifications.
- **Retrieve diagnostic information and configure AWS resources:** AWS Chatbot supports AWS Command Line Interface (CLI) commands for most AWS services from Slack on desktop and mobile devices. Your teams can analyze and respond to events faster by retrieving diagnostic information and configuring AWS resources from chatrooms. You can issue CLI commands from Slack to restart Amazon Elastic Compute Cloud (EC2) instances, increase Lambda function concurrency limits, and change Amazon Elastic Container Service (ECS) auto-scaling limits. You can also initiate workflows by invoking Lambda functions, running AWS Systems Manager Automation runbooks, or creating AWS Support cases. AWS Chatbot commands use the already-familiar AWS CLI syntax. AWS Chatbot also helps you to complete commands by providing syntax cues and prompting for parameters.
- **Predefined AWS IAM policy templates:** AWS Chatbot provides chat channel–specific permission controls through AWS Identity Access Management (IAM). Predefined templates make it easy to select and quickly set up the permissions you want associated with a given channel or chatroom.
- **Supports Slack and Amazon Chime:** You can add AWS Chatbot to your Slack channel or Amazon Chime chatroom with just a few clicks.

105.1.2. Benefits

- **Notification in real time:** AWS Chatbot delivers preselected, event-triggered alerts to your Slack channels or Amazon Chime chatrooms, keeping your team informed and aware of operational incidents or other events that they care most about.
- **Faster response:** AWS Chatbot lets you issue commands from your Slack channels, facilitates collaboration, and helps your team to quickly respond to events, operate workloads, and resolve issues without switching context to other AWS Management Tools.
- **Quick setup:** It takes you less than five minutes to set up AWS Chatbot from the AWS Chatbot console in your Slack channels or Amazon Chime chatrooms.
- **Easily define permissions:** With AWS Chatbot, you can quickly set permissions for individual Slack channels and Amazon Chime chatrooms. Account-level settings, predefined permission templates, and guardrail policies make it easy to tailor to your organization’s security and compliance needs.

105.2. Backup/Restore and Disaster Recovery

This requirement is not applicable for this service. For additional information beyond what is described herein, please refer to http://aws.amazon.com/documentation/.

105.3. Pricing Overview

There is no additional charge for AWS Chatbot. You pay only for the underlying services (such as Amazon Simple Notification Service, Amazon CloudWatch, Amazon GuardDuty, and AWS Security Hub) that you use, in the same manner as if you were using them without AWS Chatbot. Additionally, there are no minimum fees and no upfront commitments.

105.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/chatbot/
- **Service quotas:** https://docs.aws.amazon.com/general/latest/gr/chatbot.html
- **Service FAQs:** https://aws.amazon.com/chatbot/faqs/

105.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/chatbot/ and the following link for comprehensive technical documentation regarding this service:

- Administrator Guide: Describes how to set up, configure, and use AWS Chatbot.

106. AWS Cloud Map

106.1. Service Overview

AWS Cloud Map is a cloud resource discovery service. With Cloud Map, you can define custom names for your application resources, and it maintains the updated location of these dynamically changing resources. This increases your application availability because your web service always discovers the most up-to-date locations of its resources.

Modern applications are typically composed of multiple services that are accessible over an API and perform a specific function. Each service interacts with a variety of other resources, such as databases, queues, object stores, and customer-defined microservices, and it needs to be able to find the location of all the infrastructure resources on which it depends in order to function. In most cases, you manage all these resource names and their locations manually within the application code. However, manual resource management becomes time consuming and error-prone as the number of dependent infrastructure resources increases or the number of microservices dynamically scale up and down based on traffic. You can also use third-party service discovery products, but this requires installing and managing additional software and infrastructure.

Cloud Map allows you to register any application resources, such as databases, queues, microservices, and other cloud resources, with custom names. Cloud Map then constantly checks the health of resources to make sure the location is up-to-date. The application can then query the registry for the location of the resources needed based on the application version and deployment environment.

106.1.1. Features

- **Discover resources via API calls or DNS queries:** Cloud Map allows your applications to discover any web-based service via AWS SDK, API calls, or DNS queries. Over DNS, Cloud Map provides resource locations of IP addresses or IP:port combinations using either IPv4 or IPv6. Using the discovery API, Cloud Map can return URLs or ARNs as well as IP addresses and IP:port combinations.
- **Simplified service naming:** AWS Cloud Map lets you define simple custom names for services in your application. This can include Amazon Elastic Container Service (ECS) tasks, Amazon EC2 instances, Amazon S3 buckets, Amazon DynamoDB tables, Amazon Simple Queue Service (SQS) queues, and any other cloud resource.
- **Assign custom attributes:** Cloud Map lets you define custom attributes for each resource, such as location and deployment stage. This provides you the ability to customize your deployment across different regions or environments.
- **Access control:** Cloud Map is integrated with AWS Identity and Access Management (IAM) to ensure that only authenticated services can discover resources within the registry and retrieve the location and credential for those resources.
- **Automatic health check:** Amazon Route 53 health checks ensure that only healthy endpoints are returned on discovery queries. This ensures that Cloud Map always has an up-to-date registry of healthy resources.
- **Deep integration with AWS container services:** Services and tasks managed by Amazon Elastic Container Service (ECS) or Amazon Elastic Service for Kubernetes (EKS) can be automatically registered and updated in Cloud Map. As ECS launches tasks for your service, it automatically registers them as resources with Cloud Map, and they are discoverable within five seconds.
- **Rapid change propagation:** When you are using API-based discovery, the updates on your resource locations and attributes are available within 5 seconds.
- **Fully managed:** AWS Cloud Map eliminates the need to set up, update, and manage your own service discovery tools and software.

106.1.2. Benefits

- **Increase application availability:** Cloud Map constantly monitors the health of every IP-based component of your application and dynamically updates the location of each microservice as it is added or removed. This ensures that your applications only discover the most up-to-date location of its resources, increasing the availability of the application.
- **Increase developer productivity:** Cloud Map provides a single registry for all your application services which you can define with custom names. This ensures that your development teams don't have to constantly store, track, and update resource name and location information or make changes directly within the application code.

106.2. Backup/Restore and Disaster Recovery

This requirement is not applicable for this service. For additional information beyond what is described herein, please refer to http://aws.amazon.com/documentation/.

106.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

106.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/cloud-map/index.html
- **Service quotas:** https://docs.aws.amazon.com/cloud-map/latest/dg/cloud-map-limits.html
- **Service FAQs:** https://aws.amazon.com/cloud-map/faqs/

106.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/cloud-map/index.html and the following links for comprehensive technical documentation regarding this service.

- Developer Guide: Provides an overview of AWS Cloud Map, detailed feature descriptions, and procedures for using the console.
- API Reference: Describes all the API operations for AWS Cloud Map in detail. Also provides sample requests, responses, and errors for the supported web services protocols.
- AWS Cloud Map (service discovery) in the AWS CLI Reference: Describes the commands in the AWS CLI that you can use for AWS Cloud Map (listed under "servicediscovery"). Provides syntax, options, and usage examples for each command.

107. AWS Cloud9

107.1. Service Overview

AWS Cloud9 is a cloud-based integrated development environment (IDE) that lets you write, run, and debug your code with just a browser. It includes a code editor, debugger, and terminal. Cloud9 comes pre-packaged with essential tools for popular programming languages, including JavaScript, Python, PHP, and more, so you don’t need to install files or configure your development machine to start new projects. Since your Cloud9 IDE is cloud-based, you can work on your projects from your office, home, or anywhere using an internet-connected machine. Cloud9 also provides a seamless experience for developing serverless applications enabling you to easily define resources, debug, and switch between local and remote execution of serverless applications. With Cloud9, you can quickly share your development environment with your team, enabling you to pair program and track each other's inputs in real time.

107.1.1. Features

- **Fully-featured Editor:** AWS Cloud9 includes a browser-based editor that makes it easy to write, run, and debug your projects. As you type, code completion and code hinting suggestions appear in the editor, helping you code faster and avoid errors. Code completion is based not only on the identifiers within your files but also on standard libraries. The editor also enables you to fully customize your view. You can adjust your panels in any direction with a simple drag-and-drop action.
- **Broad Selection of Run Configurations:** AWS Cloud9 supports over 40 programming languages and application types including JavaScript, Python, PHP, Ruby, Go, and C++. With Cloud9, you can either choose from the default run configurations or define custom configurations by specifying environment variables, filenames, command line options, etc.

- **Integrated Debugger:** AWS Cloud9 comes with an integrated debugger, which provides commonly used capabilities like setting breakpoints, stepping through code, and inspecting variables of any PHP, Python, JS/Node.js, C/C++ app.
- **Integrated Tools for Serverless Development:** AWS Cloud9 allows you to easily build serverless applications by providing an integrated experience to get started, write, and debug serverless application code. The Cloud9 development environment is pre-packaged with SDKs, tools, and libraries needed for serverless application development. Cloud9 also supports the Serverless Application Model (SAM) so you can use SAM templates in Cloud9 to provide a simplified way of defining resources for your serverless applications. Additionally, Cloud9 allows you to edit and debug AWS Lambda functions locally, which eliminates the need to upload your code to the Lambda console for debugging.
- **Connectivity to Any Linux Server Platform:** You have the flexibility to run AWS Cloud9 development environments on a managed Amazon EC2 Linux instance or any Linux server that you are using today. You can just choose the SSH connectivity option during Cloud9 setup when connecting to your own Linux server that could be running anywhere including AWS, on-premises, or any other cloud provider.
- **Built-in Terminal:** AWS Cloud9 provides a terminal that has full sudo privileges to your managed Amazon EC2 instance. It enables you to run commands, such as pushing code changes to git, compiling your code, or displaying command output from your servers. A pre-authenticated AWS Command Line Interface is installed in your terminal, allowing you to easily control and interact with AWS services directly from the command line.
- **Collaborative Editing and Chat:** AWS Cloud9 lets you share your development environment with your team. This makes it easy for multiple developers in your team to actively see each other type and pair-program together on the same file. Cloud9 allows you to use the built-in chat capability to communicate with your team without having to leave the IDE.
- **Continuous Delivery Toolchain:** AWS Cloud9 integrates with AWS CodeStar, allowing you to quickly setup an end-to-end continuous delivery toolchain for your application and start releasing code faster on AWS. CodeStar provides a unified experience that enables you to easily build, test, and deploy applications to AWS with the help of AWS CodeCommit, AWS CodeBuild, AWS CodePipeline, and AWS CodeDeploy. In a few clicks, you will be able to connect your Cloud9 development environment to a continuous delivery toolchain.
- **File Revision History:** AWS Cloud9 keeps the revision history of the files in your development environment. This allows you to quickly access code changes that were made in the past and revert to an earlier iteration.
- **Themes:** AWS Cloud9 allows you to choose from a variety of colour schemes that control syntax highlighting and the UI. You can also fully customize the Cloud9 UI by editing your stylesheet.
- **Keyboard Shortcuts:** In addition to the default key bindings, AWS Cloud9 offers the choice to use VIM, Emacs, and Sublime key bindings, as well as to define your own custom bindings. This allows you to use the same commands and shortcuts you are already familiar with.

- **Built-in Image Editor:** AWS Cloud9 supports the ability to edit images, enabling you to resize, crop, rotate or flip the image straight from the browser.

107.1.2. Benefits

- **Code with Just a Browser:** AWS Cloud9 gives you the flexibility to run your development environment on a managed Amazon EC2 instance or any existing Linux server that supports SSH. This means that you can write, run, and debug applications with just a browser, without needing to install or maintain a local IDE. The Cloud9 code editor and integrated debugger include helpful, time-saving features such as code hinting, code completion, and step-through debugging. The Cloud9 terminal provides a browser-based shell experience enabling you to install additional software, do a git push, or enter commands.
- **Code Together in Real Time:** AWS Cloud9 makes collaborating on code easy. You can share your development environment with your team in just a few clicks and pair program together. While collaborating, your team members can see each other type in real time, and instantly chat with one another from within the IDE.
- **Build Serverless Applications with Ease:** AWS Cloud9 makes it easy to write, run, and debug serverless applications. It preconfigures the development environment with all the SDKs, libraries, and plug-ins needed for serverless development. Cloud9 also provides an environment for locally testing and debugging AWS Lambda functions. This allows you to iterate on your code directly, saving you time and improving the quality of your code.
- **Direct Terminal Access to AWS:** AWS Cloud9 comes with a terminal that includes sudo privileges to the managed Amazon EC2 instance that is hosting your development environment and a preauthenticated AWS Command Line Interface. This makes it easy for you to quickly run commands and directly access AWS services.
- **Start New Projects Quickly:** AWS Cloud9 makes it easy for you to start new projects. Cloud9's development environment comes prepackaged with tooling for over 40 programming languages, including Node.js, JavaScript, Python, PHP, Ruby, Go, and C++. This enables you to start writing code for popular application stacks within minutes by eliminating the need to install or configure files, SDKs, and plug-ins for your development machine. Because Cloud9 is cloud-based, you can easily maintain multiple development environments to isolate your project's resources.

107.2. Backup/Restore and Disaster Recovery

This service allows for backup and recovery. The service can back up machine images and volumes. Users control this via a manual or scheduled API call. Users schedule and recover backups through a web interface.

107.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

107.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/cloud9/

- **Service quotas:** https://docs.aws.amazon.com/general/latest/gr/cloud9.html
- **Service FAQs:** https://aws.amazon.com/cloud9/faqs/

107.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/cloud9/ and the following links for comprehensive technical documentation regarding this service.

- User Guide: Describes how to get started and use AWS Cloud9.

108. AWS CloudFormation

108.1. Service Overview

AWS CloudFormation gives you an easy way to model a collection of related AWS and third-party resources, provision them quickly and consistently, and manage them throughout their lifecycles, by treating infrastructure as code. A CloudFormation template describes your desired resources and their dependencies so you can launch and configure them together as a stack. You can use a template to create, update, and delete an entire stack as a single unit, as often as you need to, instead of managing resources individually. You can manage and provision stacks across multiple AWS accounts and AWS Regions.

108.1.1. Features

- **Extensibility:** Using the AWS CloudFormation Registry, you can model and provision third-party resources and modules published by AWS Partner Network (APN) Partners and the developer community. Examples of third-party resources are monitoring, team productivity, incident management, and version control tools, along with resources from APN Partners such as MongoDB, Datadog, Atlassian Opsgenie, JFrog, Trend Micro, Splunk, Aqua Security, FireEye, Sysdig, Snyk, Check Point, Spot by NetApp, Gremlin, Stackery, and Iridium. You can also browse, discover, and choose from a collection of pre-built modules by JFrog and Stackery, along with those maintained by AWS Quick Starts. You can build your own resource providers using the AWS CloudFormation CLI, an open-source tool that streamlines the development process, including local testing and code generation capabilities.
- **Cross account & cross-region management:** CloudFormation StackSets let you provision a common set of AWS resources across multiple accounts and regions, with a single CloudFormation template. StackSets takes care of automatically and safely provisioning, updating, or deleting stacks, no matter where they are.
- **Authoring with JSON/YAML:** CloudFormation allows you to model your entire cloud environment in text files. You can use open-source declarative languages, such as JSON or YAML, to describe what AWS resources you want to create and configure. If you prefer to design visually, you can use AWS CloudFormation Designer to help you get started with AWS CloudFormation templates.
- **Authoring with familiar programming languages:** With the AWS Cloud Development Kit (AWS CDK), you can define your cloud environment using TypeScript, Python, Java, and .NET. AWS CDK is an open-source software development framework that helps you model cloud application resources using familiar programming languages, and then provision your infrastructure using CloudFormation directly from your IDE. CDK provides high-level components that preconfigure cloud resources with proven defaults, so you can build cloud applications without needing to be an expert. Learn more about AWS CDK.

- **Build serverless applications with SAM:** Build serverless applications faster with the AWS Serverless Application Model (SAM), an open-source framework that provides shorthand syntax to express functions, APIs, databases, and event source mappings. With just a few lines per resource, you can define the application you want and model it using YAML. During deployment, SAM transforms and expands the SAM syntax into CloudFormation syntax.
- **Safety controls:** CloudFormation automates provisioning and updating your infrastructure in a safe and controlled manner. There are no manual steps or controls that can lead to errors. You can use Rollback Triggers to specify the CloudWatch alarms that CloudFormation should monitor during the stack creation and update process. If any of the alarms are triggered, CloudFormation rolls back the entire stack operation to a previously deployed state. Using ChangeSets, you can preview the proposed changes that CloudFormation intends to make to your infrastructure and application resources prior to execution, so that your deployments go exactly as planned. CloudFormation determines the right operations to perform, provisions resources in the most efficient way possible, and rolls back automatically if errors are encountered. This returns the state of your infrastructure and application resources to the last known good state. Using Drift Detection, you can keep track of changes to resources outside CloudFormation, making sure you always have the most up-to-date picture of your infrastructure.
- **Preview changes to your environment:** AWS CloudFormation Change Sets allow you to preview how proposed changes to a stack might affect your running resources, for example to check whether your changes will delete or replace any critical resources. CloudFormation makes the changes to your stack only after you decide to execute the Change Set.
- **Dependency management:** AWS CloudFormation automatically manages dependencies between your resources during stack management actions. You don't need to worry about specifying the order in which resources are created, updated, or deleted; CloudFormation determines the correct sequence of actions to take for each resource when performing stack operations.

108.1.2. Benefits

- **Scalable worldwide:** Scale your infrastructure worldwide and manage resources across all AWS accounts and regions through a single operation.
- **Manage your infrastructure:** Extend and manage your infrastructure to include cloud resources published in the CloudFormation Registry, the developer community, and your library.
- **Automated resources:** Automate resource management across your organization with AWS service integrations offering turnkey application distribution and governance controls.

108.2. Backup/Restore and Disaster Recovery

This requirement is not applicable for this service. For additional information beyond what is described herein, please refer to http://aws.amazon.com/documentation/.

108.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

108.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/cloudformation/index.html
- **Service quotas:** https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/cloudformation-limits.html
- **Service FAQs:** https://aws.amazon.com/cloudformation/faqs/

108.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/cloudformation/index.html and the following links for comprehensive technical documentation regarding this service.

- User Guide: Provides a conceptual overview of AWS CloudFormation and includes instructions on using the various features with the command line interface.
- API Reference: Describes all the API operations for AWS CloudFormation in detail. Also provides sample requests, responses, and errors for the supported web services protocols.
- AWS CloudFormation in the AWS CLI Reference: Describes the AWS CloudFormation commands that are available in the AWS CLI.
- User Guide for Extension Development: Provides a conceptual overview and detailed walkthroughs on using the CloudFormation CLI to model and provision both AWS and third-party extensions through CloudFormation.
- AWS CloudFormation Guard User Guide: Provides a conceptual overview of the AWS CloudFormation Guard open-source policy-as-code evaluation tool. Includes walkthroughs for writing and testing policy rules and validating data against those rules.

109. AWS CloudHSM

109.1. Service Overview

AWS CloudHSM is a cloud-based hardware security module (HSM) that enables you to easily generate and use your own encryption keys on the AWS Cloud. With CloudHSM, you can manage your own encryption keys using FIPS 140-2 Level 3 validated HSMs. CloudHSM offers you the flexibility to integrate with your applications using industry-standard APIs, such as PKCS#11, Java Cryptography Extensions (JCE), and Microsoft CryptoNG (CNG) libraries.

CloudHSM is standards-compliant and enables you to export all of your keys to most other commercially-available HSMs, subject to your configurations. It is a fully-managed service that automates time-consuming administrative tasks for you, such as hardware provisioning, software patching, high-availability, and backups. CloudHSM also enables you to scale quickly by adding and removing HSM capacity on-demand, with no up-front costs.

109.1.1. Features

- **Tamper-resistant and FIPS 140-2 Level 3 compliant HSMs:** AWS CloudHSM offers single-tenant access to tamper-resistant HSMs that comply with the U.S. Government's FIPS 140-2 Level 3 standard for cryptographic modules.

- **Scalable HSM capacity:** AWS CloudHSM makes it easy to scale your HSM capacity. You can add and remove HSMs on-demand using the AWS Management Console and AWS API.
- **Open solution:** AWS CloudHSM is an open solution that eliminates vendor lock-in. With CloudHSM, you can transfer your keys to other commercial HSM solutions to make it easy for you to migrate keys on or off of the AWS Cloud.
- **Industry-standard APIs:** AWS CloudHSM offers integration with custom applications via industry-standard APIs and supports multiple programming languages, including PKCS#11, Java Cryptography Extensions (JCE), and Microsoft CryptoNG (CNG) libraries.
- **Secure authentication**: AWS CloudHSM supports quorum authentication for critical administrative and key management functions. CloudHSM also supports multi-factor authentication (MFA) using tokens you provide.
- **AWS-managed infrastructure:** AWS CloudHSM is a managed service that automates time-consuming administrative tasks, such as hardware provisioning, software patching, high availability, and backups.

109.1.2. Benefits

- **Generate and use encryption keys on FIPS 140-2 level 3 validated HSMs:** AWS CloudHSM enables you to generate and use your encryption keys on a FIPS 140-2 Level 3 validated hardware. CloudHSM protects your keys with exclusive, single-tenant access to tamper-resistant HSM instances in your own Amazon Virtual Private Cloud (VPC).
- **Deploy secure, compliant workloads:** Utilizing HSMs as the root of trust helps you demonstrate compliance with security, privacy and anti-tamper regulations such as HIPAA, FedRAMP and PCI. AWS CloudHSM enables you to build secure, compliant workloads with high reliability and low latency, using HSM instances in the AWS cloud.
- **Use an open HSM built on industry standards:** You can use AWS CloudHSM to integrate with custom applications using industry-standard APIs, such as PKCS#11, Java Cryptography Extensions (JCE), and Microsoft CryptoNG (CNG) libraries. You can also transfer your keys to other commercial HSM solutions to make it easy for you to migrate keys on or off of AWS.
- **Keep control of your encryption keys:** AWS CloudHSM provides you access to your HSMs over a secure channel to create users and set HSM policies. The encryption keys that you generate and use with CloudHSM are accessible only by the HSM users that you specify. AWS has no visibility or access to your encryption keys.
- **Easy to manage and scale:** AWS CloudHSM automates time-consuming HSM administrative tasks for you, such as hardware provisioning, software patching, high availability, and backups. You can scale your HSM capacity quickly by adding and removing HSMs from your cluster on-demand. AWS CloudHSM automatically load balances requests and securely duplicates keys stored in any HSM to all of the other HSMs in the cluster.
- **Control AWS KMS keys:** You can configure AWS Key Management Service (KMS) to use your AWS CloudHSM cluster as a custom key store rather than the default KMS key store. With a KMS custom key store you benefit from the integration between KMS and AWS services that encrypt data while retaining control of the HSMs that protect your KMS master keys. KMS custom key store gives you the best of both worlds, combining

single-tenant HSMs under your control with the ease of use and integration of AWS KMS.

109.2. Backup/Restore and Disaster Recovery

This service allows for backup and recovery. The service can back up users, keys and cluster policies. The supplier controls the whole backup schedule. Users recover backups through a web interface.

109.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

109.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/cloudhsm/index.html
- **Service quotas:** https://docs.aws.amazon.com/cloudhsm/latest/userguide/limits.html
- **Service FAQs:** https://aws.amazon.com/cloudhsm/faqs/

109.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/cloudhsm/index.html and the following links for comprehensive technical documentation regarding this service.

- User Guide: Explains important concepts of AWS CloudHSM and documents advanced product features and command line tools.
- API Reference: Describes the API operations for AWS CloudHSM.
- AWS CloudHSM in the AWS CLI Reference: Describes the AWS CloudHSM commands that are available in the AWS Command Line Interface.

110. AWS CloudTrail

110.1. Service Overview

AWS CloudTrail enables auditing, security monitoring, and operational troubleshooting. CloudTrail records user activity and API usage across AWS services as Events. CloudTrail Events help you answer the questions of "who did what, where, and when?"

CloudTrail records two types of events: Management events capturing control plane actions on resources such as creating or deleting Amazon Simple Storage Service (Amazon S3) buckets, and data events capturing data plane actions within a resource, such as reading or writing an Amazon S3 object.

CloudTrail uses these events in three features:

- **Trails** enables delivery and storage of events in Amazon S3, with optional delivery to Amazon CloudWatch Logs and EventBridge.
- **Insights** analyzes control plane events for anomalous behavior in API call volumes.
- **Event history** provides a 90-day history of control plane actions for free. As part of its core audit capabilities, CloudTrail provides customer managed keys for encryption and log file validation to guarantee immutability.

110.1.1. Features

- **Always on:** AWS CloudTrail is enabled on all AWS accounts and records management events across AWS services without the need for any manual setups. You can view, search, and download the most recent 90-day history of your account's management events for free using CloudTrail in the AWS console or the AWS CLI Lookup API.
- **Deliver ongoing events for storage or monitoring:** You can deliver your ongoing management and data events to Amazon S3 and optionally to Amazon CloudWatch Logs by creating trails. This lets you get the complete event details, export, and store events as you like. Learn more on Creating a trail for your AWS account in the User Guide. CloudTrail's integration with Amazon EventBridge provides a convenient way to create rules-based alerts and set automated workflows in response to events.
- **Multi-region:** You can configure AWS CloudTrail to capture and store events from multiple regions in a single location. This ensures that all settings apply consistently across all existing and newly-launched regions.
- **Multi-account:** You can configure AWS CloudTrail to capture and store events from multiple accounts in a single location. This ensures that all settings apply consistently across all existing and newly-created accounts.
- **Log file integrity validation:** You can validate the integrity of AWS CloudTrail log files stored in your Amazon S3 bucket and detect whether the log files were unchanged, modified, or deleted since CloudTrail delivered them to your Amazon S3 bucket. You can use log file integrity validation in your IT security and auditing processes.
- **Log file encryption:** By default, AWS CloudTrail encrypts all log files delivered to your specified Amazon S3 bucket using Amazon S3 server-side encryption (SSE). Optionally, add a layer of security to your CloudTrail log files by encrypting the log files with your AWS Key Management Service (KMS) key. Amazon S3automatically decrypts your log files if you have decrypt permissions.
- **CloudTrail Insights:** Identify unusual activity in your AWS accounts, such as spikes in resource provisioning, bursts of AWS Identity and Access Management (IAM) actions, or gaps in periodic maintenance activity. You can enable CloudTrail Insights events in your trails.
- **CloudTrail Lake:** AWS CloudTrail Lake is a managed audit and security lake, allowing customers to aggregate, immutably store, and query their activity logs for auditing, security investigation, and operational troubleshooting.

110.1.2. Benefits

- **Prove Compliance:** Protect your organization from penalties using CloudTrail logs to prove compliance with regulations such as SOC, PCI, and HIPAA.
- **Improve security:** Improve your security posture by recording user activity and events, and set up automated workflow rules with Amazon EventBridge.
- **Track user activity:** Capture and consolidate user activity and API usage across AWS Regions and accounts on a single, centrally controlled platform.

110.2. Backup/Restore and Disaster Recovery

This service allows for backup and recovery. The service can back up logs to S3. Users control this via a manual or scheduled API call. Users schedule and recover backups through a web interface.

110.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

110.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/cloudtrail/index.html
- **Service quotas:** https://docs.aws.amazon.com/awscloudtrail/latest/userguide/WhatIsCloudTrail-Limits.html
- **Service FAQs:** https://aws.amazon.com/cloudtrail/faqs/

110.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/cloudtrail/index.html and the following links for comprehensive technical documentation regarding this service.

- User Guide: Provides detailed descriptions of product concepts and includes instructions for using CloudTrail features with both the console and the command line interface.
- API Reference: Describes all API operations for AWS CloudTrail in detail. Also provides sample requests, responses, and errors for supported web service protocols.

111. AWS CodeArtifact

111.1. Service Overview

AWS CodeArtifact is a fully managed artefact repository service that makes it easy for organizations of any size to securely store, publish, and share software packages used in their software development process. CodeArtifact can be configured to automatically fetch software packages and dependencies from public artefact repositories so developers have access to the latest versions. CodeArtifact works with commonly used package managers and build tools like Maven, Gradle, npm, yarn, twine, pip, and NuGet making it easy to integrate into existing development workflows.

Development teams often rely on both open-source software packages and those packages built within their organization. IT leaders need to be able to control access to and validate the safety of these software packages. Teams need a way to find up-to-date packages that have been approved for use by their IT leaders. To address these challenges, IT leaders turn to central artefact repository services to store and share packages. However, existing solutions often require teams to purchase licenses for software solutions that are complex to setup, scale, and operate.

AWS CodeArtifact is a pay-as-you go artefact repository service that scales based on the needs of the organization. With CodeArtifact there is no software to update or servers to manage. In just a few clicks, IT leaders can set-up central repositories that make it easy for development teams to find and use the software packages they need. IT leaders can also approve packages and control distribution across the organization, ensuring development teams consume software packages that are safe for use.

111.1.1. Features

- **Consume packages from public artifact repositories:** You can configure CodeArtifact to fetch software packages from public repositories such as the npm Registry, Maven Central, PyPI, and NuGet.org with just a few clicks. CodeArtifact automatically downloads and stores application dependencies from these repositories, so they're always available to your developers and CI/CD systems.
- **Publish and share packages:** You can use your existing package managers such as npm, pip, yarn, twine, Maven, and NuGet to publish packages developed within your organization. Development teams can save time by retrieving packages published to and shared in a central organizational repository, rather than creating their own.
- **Approve packages for use and get visibility into package usage:** You can approve packages for use by building automated workflows using CodeArtifact APIs and AWS EventBridge. Integration with AWS CloudTrail gives leaders visibility into which packages are in use and where, making it easy to identify packages that need to be updated or removed.
- **High availability and durability:** AWS CodeArtifact operates in multiple Availability Zones and stores artifact data and metadata in Amazon S3 and Amazon DynamoDB. Your encrypted data is redundantly stored across multiple facilities and multiple devices in each facility, making it highly available and highly durable.
- **Use a fully managed service:** CodeArtifact lets you focus on delivering for your customers, not configuring and maintaining your development infrastructure. CodeArtifact is a highly available service that scales to meet the needs of any software development team. There is no software to update or servers to manage.
- **Enable access control and monitoring:** AWS CodeArtifact integrates with IAM and AWS CloudTrail, offering control over who can access software packages and visibility into who has access to your software packages. CodeArtifact also integrates with AWS Key Management Service (KMS) for package encryption.
- **Access packages within a VPC:** You can increase the security of your repositories by configuring AWS CodeArtifact to use AWS PrivateLink endpoints. This allows systems running in your VPC to access packages stored in CodeArtifact without the data being transferred over the public internet.

111.1.2. Benefits

- **Securely store and share artifacts:** CodeArtifact integrates with AWS Key Management Service (KMS) to provide encrypted storage. CodeArtifact supports AWS IAM, so IT leaders can grant the appropriate level of access to different teams across their AWS accounts.
- **Reduce operational overhead:** CodeArtifact is a fully managed service, eliminating the need to set up and operate the infrastructure required to manage artifact repositories. CodeArtifact is highly available and scales to meet the needs of organizations of all sizes.
- **Pay as you go:** With CodeArtifact, there are no upfront fees or licensing costs for features that you don’t use. You pay only for the software packages stored, the number of requests made, and the data transferred out of an AWS Region.

111.2. Backup/Restore and Disaster Recovery

This requirement is not applicable for this service. For additional information beyond what is described herein, please refer to http://aws.amazon.com/documentation/.

111.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

111.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/codeartifact/index.html
- **Service quotas:** https://docs.aws.amazon.com/codeartifact/latest/ug/service-limits.html
- **Service FAQs:** https://aws.amazon.com/codeartifact/faq/

111.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/codeartifact/index.html and the following links for comprehensive technical documentation regarding this service.

- User Guide: Describes how to use AWS CodeArtifact to create repositories and use them to share packages and their assets.
- API Reference: Describes the API operations for AWS CodeArtifact. Also provides details of related request and response syntax and errors.
- AWS CLI Reference for AWS CodeArtifact: Describes the AWS CLI commands that you can use to manage domains, repositories, and packages.

112. AWS CodeBuild

112.1. Service Overview

AWS CodeBuild is a fully managed continuous integration service that compiles source code, runs tests, and produces software packages that are ready to deploy. With CodeBuild, you don't need to provision, manage, and scale your own build servers. CodeBuild scales continuously and processes multiple builds concurrently, so your builds are not left waiting in a queue. You can get started quickly by using prepackaged build environments, or you can create custom build environments that use your own build tools. With CodeBuild, you are charged by the minute for the compute resources you use.

112.1.1. Features

- **Build and test your code:** AWS CodeBuild runs your builds in preconfigured build environments that contain the operating system, programming language runtime, and build tools (e.g., Apache Maven, Gradle, npm) required to complete the task. You just specify your source code's location and select settings for your build, such as the build environment to use and the build commands to run during a build. AWS CodeBuild builds your code and stores the artifacts into an Amazon S3 bucket, or you can use a build command to upload them to an artifact repository. You can create, manage, and initiate build projects using AWS CodePipeline, the AWS Management Console, AWS CLIs, or SDKs.

- **Preconfigured build environments:**AWS CodeBuild provides build environments for Java, Python, Node.js, Ruby, Go, Android, .NET Core for Linux, and Docker.
- **Customize build environments:** You can bring your own build environments to use with AWS CodeBuild, such as for the Microsoft .NET Framework. You can package the runtime and tools for your build into a Docker image and upload it to a public Docker Hub repository or Amazon EC2 Container Registry (Amazon ECR). When you create a new build project, you can specify the location of your Docker image, and CodeBuild will pull the image and use it as the build project configuration.
- **Specify build commands:** You can define the specific commands that you want AWS CodeBuild to perform, such as installing build tool packages, running unit tests, and packaging your code. The build specification is a YAML file that lets you choose the commands to run at each phase of the build and other settings. CodeBuild helps you get started quickly with sample build specification files for common scenarios, such as builds using Apache Maven, Gradle, or npm.
- **Select compute type:** You can select the compute type that is best suited to your development needs. You can choose from three levels of compute capacity that vary by the amount of CPU and memory. This lets you choose higher CPU and memory compute if you want your builds to complete faster, or if your builds require a minimum level of CPU and memory to complete. CodeBuild supports Linux and Windows operating systems.
- **Choose source integrations:** You can initiate builds with AWS CodeBuild in several ways. For example, you can initiate builds in CodeBuild after connecting to AWS CodeCommit, GitHub, GitHub Enterprise, Bitbucket, or Amazon S3. You can also connect CodeBuild and your source repository with AWS CodePipeline, which automatically initiates a build every time you commit a change.
- **Continuous integration and delivery workflows:** AWS CodeBuild's on-demand compute and pay-as-you-go model enables you to build and integrate code more frequently, helping you find and fix bugs early in the development process when they are easy to fix. You can integrate CodeBuild into your existing continuous integration and continuous delivery (CI/CD) workflow using its source integrations, build commands, or Jenkins integration. CodeBuild also belongs to a family of AWS Code Services that helps you practice CI/CD. You can plug CodeBuild into AWS CodePipeline, which automates building and testing code in CodeBuild each time you commit a change to your source repository. You can create this CI workflow by using the AWS CodePipeline wizard to connect your source repository and then select CodeBuild as the build provider. You can easily extend your continuous integration workflow into continuous delivery with CodePipeline by integrating third-party load or user interface testing tools (e.g. BlazeMeter, Ghost Inspector) that initiate after CodeBuild completes the build. You can then deploy to your instances or on-premises servers using services integrated with AWS CodePipeline, such as AWS CodeDeploy and AWS Elastic Beanstalk.
- **Security and permissions:** Your build artifacts are encrypted with customer-specific keys that are managed by the AWS Key Management Service (KMS). AWS CodeBuild is integrated with AWS Identity and Access Management so you can set granular controls over which users and AWS resources have access to your builds.
- **Monitoring:** You can use the AWS CodeBuild Console, AWS CLI, SDKs, and APIs, or Amazon CloudWatch to view detailed information about your builds. AWS CodeBuild shows you information such as the build's start time, end time, status and commit ID. CodeBuild also streams build metrics and logs to CloudWatch. You can use CloudWatch

to create a custom dashboard, set a CloudWatch Alarm, troubleshoot build issues, or inspect build logs.

- **Receive Notifications:** You can create notifications for events impacting your build projects. Notifications will come in the form of Amazon SNS notifications. Each notification includes a status message as well as a link to the resources whose event generated that notification.

112.1.2. Benefits

- **Fully managed build service:** AWS CodeBuild eliminates the need to set up, patch, update, and manage your own build servers and software. There is no software to install or manage.
- **Continuous scaling:** AWS CodeBuild scales up and down automatically to meet your build volume. It immediately processes each build you submit and can run separate builds concurrently, which means your builds are not left waiting in a queue.
- **Pay as you go:** With AWS CodeBuild, you are charged based on the number of minutes it takes to complete your build. This means you no longer have to worry about paying for idle build server capacity.
- **Extensible:** You can bring your own build tools and programming runtimes to use with AWS CodeBuild by creating customized build environments in addition to the prepackaged build tools and runtimes supported by CodeBuild.
- **Enables continuous integration and delivery:** AWS CodeBuild belongs to a family of AWS Code Services, which you can use to create complete, automated software release workflows for continuous integration and delivery (CI/CD). You can also integrate CodeBuild into your existing CI/CD workflow. For example, you can use CodeBuild as a worker node for your existing Jenkins server setup for distributed builds.
- **Secure:** With AWS CodeBuild, your build artifacts are encrypted with customer-specific keys that are managed by the AWS Key Management Service (KMS). CodeBuild is integrated with AWS Identity and Access Management (IAM), so you can assign user-specific permissions to your build projects.

112.2. Backup/Restore and Disaster Recovery

This requirement is not applicable for this service. For additional information beyond what is described herein, please refer to http://aws.amazon.com/documentation/.

112.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

112.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/codebuild/index.html
- **Service quotas:** https://docs.aws.amazon.com/codebuild/latest/userguide/limits.html
- **Service FAQs:** https://aws.amazon.com/codebuild/faqs/?nc=sn&loc=5

112.4.1. Technical Requirements

Please refer to https://docs.aws.amazon.com/codebuild/index.html and the following links for comprehensive technical documentation regarding this service.

- User Guide: Describes how you can use AWS CodeBuild, an AWS service that builds your software applications in the AWS cloud.
- API Reference: Describes the API operations for AWS CodeBuild. Also provides details of related request and response syntax and errors.
- AWS CLI reference for AWS CodeBuild: Describes the AWS CLI commands that you can use to automate building your source code.

113. AWS CodeCommit

113.1. Service Overview

AWS CodeCommit is a secure, highly scalable, managed source control service that hosts private Git repositories. It makes it easy for teams to securely collaborate on code with contributions encrypted in transit and at rest. CodeCommit eliminates the need for you to manage your own source control system or worry about scaling its infrastructure. You can use CodeCommit to store anything from code to binaries. It supports the standard functionality of Git, so it works seamlessly with your existing Git-based tools.

113.1.1. Features

- **Collaboration:** AWS CodeCommit is designed for collaborative software development. You can easily commit, branch, and merge your code allowing you to easily maintain control of your team's projects. CodeCommit also supports pull requests, which provide a mechanism to request code reviews and discuss code with collaborators. You can create a repository from the AWS Management Console, AWS CLI, or AWS SDKs and start working with the repository using Git.
- **Encryption:** You can transfer your files to and from AWS CodeCommit using HTTPS or SSH, as you prefer. Your repositories are also automatically encrypted at rest through AWS Key Management Service (AWS KMS) using customer-specific keys.
- **Access Control:** AWS CodeCommit uses AWS Identity and Access Management to control and monitor who can access your data as well as how, when, and where they can access it. CodeCommit also helps you monitor your repositories via AWS CloudTrail and AWS CloudWatch.
- **High Availability and Durability:** AWS CodeCommit stores your repositories in Amazon S3 and Amazon DynamoDB. Your encrypted data is redundantly stored across multiple facilities. This architecture increases the availability and durability of your repository data.
- **The repositories you need, when you need them:** AWS CodeCommit allows you to create up to 1,000 repositories by default, and additional repositories up to 25,000 by request. You can store and version any kind of file, including application assets such as images and libraries alongside your code. It's easy to create repositories when you need them, and delete them when you're done.
- **Easy Access and Integration:** You can use the AWS Management Console, AWS CLI, and AWS SDKs to manage your repositories. You can also use Git commands or Git graphical tools to interact with your repository source files. AWS CodeCommit supports

all Git commands and works with your existing Git tools. You can integrate with your development environment plugins or continuous integration/continuous delivery systems.

- **Notifications and Custom Scripts:** You can now receive notifications for events impacting your repositories. Notifications will come in the form of Amazon SNS notifications. Each notification will include a status message as well as a link to the resources whose event generated that notification. Additionally, using AWS CodeCommit repository triggers, you can send notifications and create HTTP webhooks with Amazon SNS or invoke AWS Lambda functions in response to the repository events you choose.

113.1.2. Benefits

- **Fully managed:** AWS CodeCommit eliminates the need to host, maintain, back up, and scale your own source control servers. The service automatically scales to meet the growing needs of your project.
- **Secure:** AWS CodeCommit automatically encrypts your files in transit and at rest. CodeCommit is integrated with AWS Identity and Access Management (IAM) allowing you to customize user-specific access to your repositories.
- **High availability:** AWS CodeCommit has a highly scalable, redundant, and durable architecture. The service is designed to keep your repositories highly available and accessible.
- **Collaborate on code:** AWS CodeCommit helps you collaborate on code with teammates via pull requests, branching, and merging. You can implement workflows that include code reviews and feedback by default, and control who can make changes to specific branches.
- **Faster development lifecycle:** AWS CodeCommit keeps your repositories close to your build, staging, and production environments in the AWS cloud. You can transfer incremental changes instead of the entire application. This allows you to increase the speed and frequency of your development lifecycle.
- **Use your existing tools:** AWS CodeCommit supports all Git commands and works with your existing Git tools. You can keep using your preferred development environment plugins, continuous integration/continuous delivery systems, and graphical clients with CodeCommit.

113.2. Backup/Restore and Disaster Recovery

This requirement is not applicable for this service. For additional information beyond what is described herein, please refer to http://aws.amazon.com/documentation/.

113.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

113.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/codecommit/index.html
- **Service quotas:** https://docs.aws.amazon.com/codecommit/latest/userguide/limits.html
- **Service FAQs:** https://aws.amazon.com/codecommit/faqs/

113.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/codecommit/index.html and the following links for comprehensive technical documentation regarding this service.

- User Guide: Get started privately storing and managing assets (such as documents, source code, and binary files) using AWS CodeCommit.
- API Reference: Describes all the API operations for AWS CodeCommit in detail. Also provides details of request and response syntax and errors for the supported web services protocols.
- AWS CLI Reference for AWS CodeCommit: Describes the AWS CLI commands that you can use to automate management of your assets.
- AWS CodeCommit Resource Type Reference: Learn how to automatically create resources for AWS CodeCommit by using AWS CloudFormation.
- Developer Tools Console User Guide: Get started with the Developer Tools console, which includes robust features that can help you expand and manage your AWS CodeCommit resources.
- Amazon CodeGuru Reviewer: Amazon CodeGuru Reviewer can provide analysis and recommendations for repositories in AWS CodeCommit.

114. AWS CodeDeploy

114.1. Service Overview

AWS CodeDeploy is a fully managed deployment service that automates software deployments to a variety of compute services such as Amazon EC2, AWS Fargate, AWS Lambda, and your on-premises servers. AWS CodeDeploy makes it easier for you to rapidly release new features, helps you avoid downtime during application deployment, and handles the complexity of updating your applications. You can use AWS CodeDeploy to automate software deployments, eliminating the need for error-prone manual operations. The service scales to match your deployment needs.

114.1.1. Features

- **Repeatable deployments:** You can easily repeat an application deployment across different groups of instances with AWS CodeDeploy. CodeDeploy uses a file and command-based install model, which enables it to deploy any application and reuse existing setup code. The same setup code can be used to consistently deploy and test updates across your deployment, test, and production release stages for Amazon EC2 instances. Eliminating manual steps from deployments increases both the speed and reliability of your software delivery process.
- **Automatic scaling:** AWS CodeDeploy allows you to integrate software deployment and scaling activities in order to keep your application up-to-date in a dynamic production environment. For Amazon EC2 instances, CodeDeploy integrates with Auto Scaling. Auto Scaling allows you to scale EC2 capacity according to conditions you define such as spikes in traffic. CodeDeploy is notified whenever a new instance launches into an Auto Scaling group and will automatically perform an application deployment on the new instance before it is added to an Elastic Load Balancing load balancer.
- **On-premises deployments:** You can use AWS CodeDeploy to automate software deployments across your development, test, and production environments running on any instance including instances in your own data centers (your instances will need to be

able to connect to AWS public endpoints). This enables you to use a single service to consistently deploy applications across hybrid architectures.

- **Rolling and Blue/Green updates:** Applications do not require downtime when they're being upgraded to a new revision with AWS CodeDeploy. AWS CodeDeploy can perform blue/green deployments to Amazon EC2 instances, an Amazon ECS service (both EC2 and AWS Fargate launch type), or an AWS Lambda function. With a blue/green deployment, the new version of your application is launched alongside the old version. Once the new revision is tested and declared ready, CodeDeploy can shift the traffic from your prior version to your new version according to your specifications. CodeDeploy can also perform a rolling update across a group of Amazon EC2 instances where only a fraction of the instances are taken offline at any one time. CodeDeploy progressively works its way across the instances allowing applications to remain available and continue serving traffic. For AWS Lambda functions, incoming traffic is gradually routed from the old version to the new one.
- **Deployment health tracking:** Deployment Health Tracking works in conjunction with rolling updates to keep applications highly available during deployments. Unexpected downtime can occur if bad updates are deployed. AWS CodeDeploy monitors your deployment and will stop a deployment if there are too many failed updates.
- **Stop and rollback:** You can stop an application deployment that is in process at any time using the AWS Management Console, the AWS CLI, or any of the AWS SDKs. You can simply re-deploy that revision if you want to continue the stopped deployment at a later time. You can also immediately rollback by redeploying the previous revision.
- **Monitoring and control:** You can launch, control, and monitor deployments of your software directly from the AWS Management Console or by using the AWS CLI, SDKs, or APIs. In the case of a failure, you can pinpoint the script experiencing failure. You can also set push notifications that allow you to monitor the status of your deployments via SMS or email messages through Amazon Simple Notification Service.
- **Deployment groups:** One application can be deployed to multiple deployment groups. Deployment groups are used to match configurations to specific environments, such as a staging or production environments. You can test a revision in staging and then deploy that same code with the same deployment instructions to production once you are satisfied.
- **Deployment history:** AWS CodeDeploy tracks and stores the recent history of your deployments. You can view which application versions are currently deployed to each of your target deployment groups. You can inspect the change history and success rates of past deployments to specific deployment groups. You can also investigate a timeline of past deployments for a detailed view of your deployment successes and errors.

114.1.2. Benefits

Copy and paste benefits from service landing page

- **Automated deployments:** AWS CodeDeploy fully automates your software deployments, allowing you to deploy reliably and rapidly. You can consistently deploy your application across your development, test, and production environments whether deploying to Amazon EC2, AWS Fargate, AWS Lambda, or your on-premises servers. The service scales with your infrastructure.
- **Minimize downtime:** AWS CodeDeploy helps maximize your application availability during the software deployment process. It introduces changes incrementally and tracks

application health according to configurable rules. Software deployments can easily be stopped and rolled back if there are errors.

- **Centralized control:** AWS CodeDeploy allows you to easily launch and track the status of your application deployments through the AWS Management Console or the AWS CLI. CodeDeploy gives you a detailed report allowing you to view when and to where each application revision was deployed. You can also create push notifications to receive live updates about your deployments.
- **Easy to adopt:** AWS CodeDeploy is platform and language agnostic, works with any application, and provides the same experience whether you're deploying to Amazon EC2, AWS Fargate, or AWS Lambda. You can easily reuse your existing setup code. CodeDeploy can also integrate with your existing software release process or continuous delivery toolchain (e.g., AWS CodePipeline, GitHub, Jenkins).

114.2. Backup/Restore and Disaster Recovery

This requirement is not applicable for this service. For additional information beyond what is described herein, please refer to http://aws.amazon.com/documentation/.

114.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

114.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/codedeploy/index.html
- **Service quotas:** https://docs.aws.amazon.com/codedeploy/latest/userguide/limits.html
- **Service FAQs:** https://aws.amazon.com/codedeploy/faqs/?nc=sn&loc=6

114.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/codedeploy/index.html and the following links for comprehensive technical documentation regarding this service.

- User Guide: Describes how to get started and deploy applications using AWS CodeDeploy.
- API Reference: Describes all the API operations for AWS CodeDeploy. Also provides sample requests, responses, and errors for the supported web services protocols.
- AWS CLI Reference for AWS CodeDeploy: Describes the AWS CLI commands that you can use to automate deployments.

115. AWS CodePipeline

115.1. Service Overview

AWS CodePipeline is a fully managed continuous delivery service that helps you automate your release pipelines for fast and reliable application and infrastructure updates. CodePipeline automates the build, test, and deploy phases of your release process every time there is a code change, based on the release model you define. This enables you to rapidly and reliably deliver features and updates. You can easily integrate AWS CodePipeline with third-party services

such as GitHub or with your own custom plugin. With AWS CodePipeline, you only pay for what you use. There are no upfront fees or long-term commitments.

115.1.1. Features

- **Workflow modelling:** A pipeline defines your release process workflow, and describes how a new code change progresses through your release process. A pipeline comprises a series of stages (e.g., build, test, and deploy), which act as logical divisions in your workflow. Each stage is made up of a sequence of actions, which are tasks such as building code or deploying to test environments. AWS CodePipeline provides you with a graphical user interface to create, configure, and manage your pipeline and its various stages and actions, allowing you to easily visualize and model your release process workflow.
- **Parallel Execution**: You can use CodePipeline to model your build, test, and deployment actions to run in parallel in order to increase your workflow speeds.
- **AWS integrations:** AWS CodePipeline can pull source code for your pipeline directly from AWS CodeCommit, GitHub, Amazon ECR, or Amazon S3. It can run builds and unit tests in AWS CodeBuild. CodePipeline can deploy your changes using AWS CodeDeploy, AWS Elastic Beanstalk, Amazon Elastic Container Service (Amazon ECS), or AWS Fargate. You can model AWS CloudFormation actions that let you provision, update, or delete AWS resources as part of your release process. This also enables you to continuously deliver serverless applications built using AWS Lambda, Amazon API Gateway, and Amazon DynamoDB with the AWS Serverless Application Model (AWS SAM). You can also trigger custom functions defined by code at any stage of your pipeline using CodePipeline's integration with AWS Lambda. For example, you can trigger a Lambda function that tests whether your web application deployed successfully. CodePipeline lets you configure a pipeline that ties these services together along with third-party developer tools and custom systems.
- **Pre-built plugins:** AWS CodePipeline allows you to integrate third-party developer tools, like GitHub or Jenkins, into any stage of your release process with one click. You can use third-party tools for source control, build, test, or deployment.
- **Custom plugins:** AWS CodePipeline allows you to integrate your own custom systems. You can register a custom action that allows you to hook your servers into your pipeline by integrating the CodePipeline open source agent with your servers. You can also use the CodePipeline Jenkins plugin to easily register your existing build servers as a custom action.
- **Declarative templates:** AWS CodePipeline allows you to define your pipeline structure through a declarative JSON document that specifies your release workflow and its stages and actions. These documents enable you to update existing pipelines as well as provide starting templates for creating new pipelines.
- **Access control:** AWS CodePipeline uses AWS IAM to manage who can make changes to your release workflow, as well as who can control it. You can grant users access through IAM users, IAM roles, and SAML-integrated directories.
- **Receive Notifications:** You can create notifications for events impacting your pipelines. Notifications will come in the form of Amazon SNS notifications. Each notification includes a status message as well as a link to the resources whose event generated that notification.

115.1.2. Benefits

- **Rapid delivery:** AWS CodePipeline automates your software release process, allowing you to rapidly release new features to your users. With CodePipeline, you can quickly iterate on feedback and get new features to your users faster. Automating your build, test, and release process allows you to quickly and easily test each code change and catch bugs while they are small and simple to fix. You can ensure the quality of your application or infrastructure code by running each change through your staging and release process.
- **Configurable workflow:** AWS CodePipeline allows you to model the different stages of your software release process using the console interface, the AWS CLI, AWS CloudFormation, or the AWS SDKs. You can easily specify the tests to run and customize the steps to deploy your application and its dependencies.
- **Get started fast:** With AWS CodePipeline, you can immediately begin to model your software release process. There are no servers to provision or set up. CodePipeline is a fully managed continuous delivery service that connects to your existing tools and systems.
- **Easy to integrate:** AWS CodePipeline can easily be extended to adapt to your specific needs. You can use our pre-built plugins or your own custom plugins in any step of your release process. For example, you can pull your source code from GitHub, use your on-premises Jenkins build server, run load tests using a third-party service, or pass on deployment information to your custom operations dashboard.

115.2. Backup/Restore and Disaster Recovery

This requirement is not applicable for this service. For additional information beyond what is described herein, please refer to http://aws.amazon.com/documentation/.

115.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

115.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/codepipeline/index.html
- **Service quotas:** https://docs.aws.amazon.com/codepipeline/latest/userguide/limits.html
- **Service FAQs:** https://aws.amazon.com/codepipeline/faqs/?nc=sn&loc=5

115.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/codepipeline/index.html and the following links for comprehensive technical documentation regarding this service.

- User Guide: Walks through how to set up AWS CodePipeline and integrate it with other services.
- API Reference: Describes all the API operations for AWS CodePipeline in detail. Also provides sample requests, responses, and errors for the supported web services protocols.

- AWS CLI Reference for AWS CodePipeline: Describes the AWS CLI commands that you can use to manage your assets.

116. AWS CodeStar

116.1. Service Overview

AWS CodeStar enables you to quickly develop, build, and deploy applications on AWS. AWS CodeStar provides a unified user interface, enabling you to easily manage your software development activities in one place. With AWS CodeStar, you can set up your entire continuous delivery toolchain in minutes, allowing you to start releasing code faster. AWS CodeStar makes it easy for your whole team to work together securely, allowing you to easily manage access and add owners, contributors, and viewers to your projects. Each AWS CodeStar project comes with a project management dashboard, including an integrated issue tracking capability powered by Atlassian JIRA Software. With the AWS CodeStar project dashboard, you can easily track progress across your entire software development process, from your backlog of work items to teams' recent code deployments.

116.1.1. Features

- **Project templates:** AWS CodeStar provides a number of project templates to help you quickly start developing applications for deployment on Amazon EC2, AWS Lambda, and AWS Elastic Beanstalk with support for many popular programming languages including Java, JavaScript, Python, Ruby, and PHP. With AWS CodeStar, you can use a code editor of your choice such as Visual Studio, Eclipse, or the AWS Command Line Interface.
- **Team access management:** AWS CodeStar uses AWS Identity and Access Management (IAM) to manage developer identities and provides built-in, role-based security policies that allow you to easily secure access to your team. AWS CodeStar allows you to share your projects using three levels of access: owners, contributors, and viewers.
- **Hosted Git repository:** AWS CodeStar stores your application code securely on AWS CodeCommit, a fully-managed source control service that eliminates the need to manage your own infrastructure to host Git repositories. You can also choose to have your project source code stored in a GitHub repository in your own GitHub account.
- **Fully managed build service:** AWS CodeStar compiles and packages your source code with AWS CodeBuild, a fully-managed build service that makes it possible for you to build, test, and integrate code more frequently.
- **Automated continuous delivery pipeline:** AWS CodeStar accelerates software release with the help of AWS CodePipeline, a continuous integration and continuous delivery (CI/CD) service. Each project comes pre-configured with an automated pipeline that continuously builds, tests, and deploys your code with each commit.
- **Automated deployments:** AWS CodeStar integrates with AWS CodeDeploy and AWS CloudFormation so that you can easily update your application code and deploy to Amazon EC2 and AWS Lambda.
- **IDE integrations:** After you create a project in AWS CodeStar, you can begin developing your code directly in AWS Cloud9, which makes it easy to get started developing on AWS. Cloud9 is a cloud-based IDE that lets you write, run, and debug

your code with just a browser. Cloud9 comes with a terminal that has a pre-authenticated AWS Command Line Interface giving you immediate access to a broad spectrum of AWS service. In addition to Cloud9, CodeStar allows you to pick from a number of other popular IDEs such as Microsoft Visual Studio and Eclipse.

- **Central project dashboard:** AWS CodeStar projects include a unified dashboard, so you can easily track and manage your end-to-end development toolchain. With the project dashboard, you can centrally manage activity for your CI/CD pipeline, such as code commits, builds, tests, and deployments and take remedial action where needed. AWS CodeStar also includes a project wiki, making it easy for you to provide team information such as project links, code samples, and team notes. AWS CodeStar also integrates with Amazon CloudWatch, an application monitoring service, and Atlassian JIRA Software, a third-party issue tracking and project management tool. These integrations allow you to centrally monitor application activity and manage JIRA issues in the AWS CodeStar dashboard.

116.1.2. Benefits

- **Start developing on AWS in minutes:** AWS CodeStar makes it easy for you to set up your entire development and continuous delivery toolchain for coding, building, testing, and deploying your application code. To start a project, you can choose from a variety of AWS CodeStar templates for Amazon EC2, AWS Lambda, and AWS Elastic Beanstalk. You have the option to choose AWS CodeCommit or GitHub to use as your project's source control. You also have the option to edit your source code using one of several options including AWS Cloud9, Microsoft Visual Studio, or Eclipse. After you make your selections the underlying AWS services are provisioned in minutes, allowing you to quickly start coding and deploying your applications.
- **Manage software delivery in one place:** AWS CodeStar provides an easy way to coordinate your day-to-day development activities through a unified user interface, reducing the need to switch between various service consoles. AWS CodeStar's project dashboard lets you monitor application activity, and track progress across all stages of your software development process, including code commits, builds, tests, and deployments, from a central place. AWS CodeStar integrates Atlassian JIRA Software, a third-party issue tracking and project management tool, allowing you to easily manage JIRA issues directly in the AWS CodeStar dashboard.
- **Work across your team securely:** AWS CodeStar enables you to collaborate on projects across your team in a secure manner. You can easily manage access for project owners, contributors, and viewers without needing to manually configure your own policy for each service. AWS CodeStar simplifies the process of setting up project access for teams by providing built-in role-based policies that follow AWS Identity and Access Management best practices.
- **Choose from a variety of project templates:** With AWS CodeStar project templates, you can easily develop a variety of applications such as websites, web applications, web services, and Alexa skills. AWS CodeStar project templates include the code for getting started on supported programming languages including Java, JavaScript, PHP, Ruby, C#, and Python.

116.2. Backup/Restore and Disaster Recovery

This requirement is not applicable for this service. For additional information beyond what is described herein, please refer to http://aws.amazon.com/documentation/.

116.3. Pricing Overview

There is no additional charge for AWS CodeStar. You only pay for AWS resources (e.g. Amazon EC2 instances, AWS Lambda executions, Amazon Elastic Block Store volumes, or Amazon S3 buckets) that you provision in your CodeStar projects. You only pay for what you use, as you use it; there are no minimum fees and no upfront commitments.

116.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/codestar/
- **Service quotas:** https://docs.aws.amazon.com/general/latest/gr/codestar.html
- **Service FAQs:** https://aws.amazon.com/codestar/faqs/

116.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/codestar/ and the following links for comprehensive technical documentation regarding this service.

- User Guide: Provides conceptual overviews of AWS CodeStar and explains how to use it to develop software applications on AWS.

117. AWS Compute Optimizer

117.1. Service Overview

AWS Compute Optimizer recommends optimal AWS resources for your workloads to reduce costs and improve performance by using machine learning to analyse historical utilization metrics. Over-provisioning resources can lead to unnecessary infrastructure cost, and under-provisioning resources can lead to poor application performance. Compute Optimizer helps you choose optimal configurations for three types of AWS resources: e.g. Amazon Elastic Compute Cloud (EC2) instance types, Amazon Elastic Block Store (EBS) volumes, and AWS Lambda functions, based on your utilization data.

117.1.1. Features

- **Configuration analysis:** Compute Optimizer analyzes the configuration and resource utilization of your workload to identify dozens of defining characteristics, for example, if a workload is CPU-intensive, if it exhibits a daily pattern, or if a workload accesses local storage frequently. The service processes these characteristics and identifies the hardware resource required by the workload.
- **Recommendations:** Compute Optimizer infers how the workload would have performed on various hardware platforms (e.g. Amazon EC2 instance types) or using different configurations (e.g. Amazon EBS volume IOPS settings, and AWS Lambda function memory sizes) to offer recommendations.

117.1.2. Benefits

- **Lower costs by up to 25%:** You can take advantage of the recommendations in Compute Optimizer to reduce costs by up to 25%. Compute Optimizer analyzes the configuration and resource utilization of a workload to identify AWS resources, such as Amazon EC2 instances, Amazon EBS volumes, and AWS Lambda functions, that might be under-provisioned or over-provisioned. Compute Optimizer then makes

recommendations for right-sizing your AWS resources and switching to a different instance size or instance type to save costs.

- **Optimize performance with actionable recommendations:** Compute Optimizer recommends up to 3 options from 140+ EC2 instance types, as well as a wide range of EBS volume and Lambda function configuration options, to right size your workloads. Compute Optimizer also projects what the CPU utilization, memory utilization, and run time of your workload would have been on recommended AWS resource options. This helps you understand how your workload would have performed on the recommended options before implementing the recommendations.
- **Get started quickly:** With just four simple clicks from the AWS Management Console, Compute Optimizer automatically generates recommendations. These recommendations are based on your current resource utilization data from Amazon CloudWatch metrics and AWS resource metadata. You don't need to invest substantial time and money to set up rules-based thresholds.

117.2. Backup/Restore and Disaster Recovery

This requirement is not applicable for this service. For additional information beyond what is described herein, please refer to http://aws.amazon.com/documentation/.

117.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

117.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/compute-optimizer/
- **Service quotas:** https://docs.aws.amazon.com/general/latest/gr/compute-optimizer.html
- **Service FAQs:** https://aws.amazon.com/compute-optimizer/faqs/

117.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/compute-optimizer/ and the following links for comprehensive technical documentation regarding this service.

- User Guide: Describes key concepts of AWS Compute Optimizer and provides instructions to opt in and view your recommendations.
- API Reference: Describes the API operations for AWS Compute Optimizer. It also provides the basic structure and elements of an API request and response.

118. AWS Config

118.1. Service Overview

AWS Config is a service that enables you to assess, audit, and evaluate the configurations of your AWS resources. Config continuously monitors and records your AWS resource configurations and allows you to automate the evaluation of recorded configurations against desired configurations. With Config, you can review changes in configurations and relationships between AWS resources, dive into detailed resource configuration histories, and determine your overall compliance against the configurations specified in your internal guidelines. This enables

you to simplify compliance auditing, security analysis, change management, and operational troubleshooting.

118.1.1. Features

- **Configuration history of AWS resources:** AWS Config records details of changes to your AWS resources to provide you with a configuration history. You can use the AWS Management Console, API, or CLI to obtain details of what a resource's configuration looked like at any point in the past. AWS Config will also automatically deliver a configuration history file to the Amazon S3 bucket you specify.
- **Configuration history of software:** AWS Config enables you to record software configuration changes within your Amazon EC2 instances and servers running on-premises, as well as servers and Virtual Machines in environments provided by other cloud providers. With AWS Config, you gain visibility into operating system (OS) configurations, system-level updates, installed applications, network configuration and more. AWS Config also provides a history of OS and system-level configuration changes alongside infrastructure configuration changes recorded for EC2 instances.
- **Resource relationships tracking:** AWS Config discovers, maps, and tracks AWS resource relationships in your account. For example, if a new Amazon EC2 security group is associated with an Amazon EC2 instance, AWS Config records the updated configurations of both the Amazon EC2 security group and the Amazon EC2 instance.
- **Configurable and customizable rules:** AWS Config provides you with pre-built rules for evaluating provisioning and configuring of your AWS resources as well as software within managed instances, including Amazon EC2 instances and servers running on-premises. You can customize pre-built rules to evaluate your AWS resource configurations and configuration changes, or create your own custom rules in AWS Lambda that define your internal best practices and guidelines for resource configurations. Using AWS Config, you can assess your resource configurations and resource changes for compliance against the built-in or custom rules.
- **Conformance packs:** Conformance packs help you manage compliance of your AWS resource configuration at scale--from policy definition to auditing and aggregated reporting--using a common framework and packaging model. Conformance packs are integrated with AWS Organizations. Using conformance packs as your compliance framework, you can package a collection of AWS Config rules and remediation actions into a single entity (known as a conformance pack) and deploy it across an entire organization. This is particularly useful if you need to quickly establish a common baseline for resource configuration policies and best practices across multiple accounts in your organization in a scalable and efficient way.
- **Multi-account, multi-region data aggregation:** Multi-account, multi-region data aggregation is a capability in AWS Config that enables centralized auditing and governance. It gives you an enterprise-wide view of your AWS Config rule compliance status, and you can associate your AWS organization to quickly add your accounts. The aggregated dashboard in AWS Config will display the total count of non-compliant rules across your organization, the top five non-compliant rules by number of resources, and the top five AWS accounts that have the most number of non-compliant rules. You can then drill down to view details about the resources that are violating the rule, and the list of rules that are being violated by an account.
- **Extensibility:** AWS Config supports extensibility by allowing you to publish the configuration of third-party resources into AWS Config using our public APIs. Examples

of third-party resources include version control systems such as GitHub, Microsoft Active Directory resources or any on-premises server. AWS Config enables you to view and monitor the resource inventory and configuration history of these third-party resources using the AWS Config console and APIs, like you do for AWS resources. You can also create AWS Config rules or conformance packs to evaluate these third-party resources against best practices, internal policies, and regulatory policies.

- **Configuration snapshots:** AWS Config can provide you with a configuration snapshot — a point-in-time capture of all your resources and their configurations. Configuration snapshots are generated on demand via the AWS CLI or API and delivered to the Amazon S3 bucket you specify.
- **Cloud governance dashboard:** AWS Config provides you a visual dashboard to help you quickly spot non-compliant resources and take appropriate action. IT Administrators, Security Experts, and Compliance Officers can see a shared view of your AWS resources compliance posture.

118.1.2. Benefits

- **Continuous monitoring:** With AWS Config, you are able to continuously monitor and record configuration changes of your AWS resources. Config also enables you to inventory your AWS resources, the configurations of your AWS resources, as well as software configurations within EC2 instances at any point in time. Once change from a previous state is detected, an Amazon Simple Notification Service (SNS) notification can be delivered for you to review and take action.
- **Continuous assessment:** AWS Config allows you to continuously audit and assess the overall compliance of your AWS resource configurations with your organization’s policies and guidelines. AWS Config provides you with the ability to define rules for provisioning and configuring AWS resources. These rules can be provisioned independently or packaged together with compliance remediation actions inside a pack (known as a conformance pack) that can be deployed across your entire organization with a single click. Resource configurations or configuration changes that deviate from your rules automatically trigger Amazon Simple Notification Service (SNS) notifications and Amazon CloudWatch events so that you can be alerted on a continuous basis. You can also take advantage of the visual dashboard to check your overall compliance status and quickly spot non-compliant resources.
- **Change management:** With AWS Config, you are able to track the relationships among resources and review resource dependencies prior to making changes. Once a change occurs, you are able to quickly review the history of the resource's configuration and determine what the resource’s configuration looked like at any point in the past. Config provides you with information to assess how a change to a resource configuration would affect your other resources, which minimizes the impact of change-related incidents.
- **Operational troubleshooting:** With AWS Config, you can capture a comprehensive history of your AWS resource configuration changes to simplify troubleshooting of your operational issues. Config helps you identify the root cause of operational issues through its integration with AWS CloudTrail, a service that records events related to API calls for your account. Config leverages CloudTrail records to correlate configuration changes to particular events in your account. You can obtain the details of the event API call that invoked the change (e.g., who made the request, at what time, and from which IP address) from the CloudTrail logs.

- **Enterprise-wide compliance monitoring:** With multi-account, multi-region data aggregation in AWS Config, you can view compliance status across your enterprise and identify non-compliant accounts. You can dive deeper to view status for a specific region or a specific account across regions. You can view this data from the Config console in a central account, removing the need to retrieve this information individually from each account, and each region.
- **Support for third-party resources:** AWS Config is designed to be your primary tool to perform configuration audit and compliance verification of both your AWS and third-party resources. You can publish the configuration of third-party resources such as GitHub repositories, Microsoft Active Directory resources, or any on-premises server into AWS. You can then view and monitor the resource inventory and configuration history using the AWS Config console and APIs, just like you do for AWS resources. You can also create AWS Config rules or conformance packs to evaluate these third-party resources against best practices, internal policies, and regulatory policies.

118.2. Backup/Restore and Disaster Recovery

This service allows for backup and recovery. The service can back up config logs to S3. Users control this via a manual or scheduled API call. Users schedule and recover backups through a web interface.

118.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

118.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/config/
- **Service quotas:** https://docs.aws.amazon.com/general/latest/gr/awsconfig.html
- **Service FAQs:** https://aws.amazon.com/config/faq/

118.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/config/ and the following links for comprehensive technical documentation regarding this service.

- Developer Guide: Provides a conceptual overview of AWS Config and includes detailed development instructions for using the various features.
- API Reference: Describes all the API operations for AWS Config in detail. Also provides sample requests, responses, and errors for the supported web services protocols.
- CLI Reference: Documents the AWS Config command line interface.

119. AWS Control Tower

119.1. Service Overview

If you have multiple AWS accounts and teams, cloud setup and governance can be complex and time consuming, slowing down the very innovation you're trying to speed up. AWS Control Tower provides the easiest way to set up and govern a secure, multi-account AWS

environment, called a landing zone. It creates your landing zone using AWS Organizations, bringing ongoing account management and governance as well as implementation best practices based on AWS's experience working with thousands of customers as they move to the cloud. Builders can provision new AWS accounts in a few clicks, while you have peace of mind knowing that your accounts conform to company policies. Extend governance into new or existing accounts, and gain visibility into their compliance status quickly. If you are building a new AWS environment, starting out on your journey to AWS, or starting a new cloud initiative, AWS Control Tower will help you get started quickly with built-in governance and best practices.

119.1.1. Features

- **Landing Zone:** A landing zone is a well-architected, multi-account AWS environment based on security and compliance best practices. AWS Control Tower automates the setup of a new landing zone using best-practices blueprints for identity, federated access, and account structure. Examples of blueprints that are automatically implemented in your landing zone include:
  - Create a multi-account environment using AWS Organizations
  - Provide identity management using AWS Single Sign-On (AWS SSO) default directory
  - Provide federated access to accounts using AWS SSO
  - Centralize logging from AWS CloudTrail, and AWS Config stored in Amazon Simple Storage Service (Amazon S3)
  - Enable cross-account security audits using AWS Identity and Access Management (IAM) and AWS SSO

  The landing zone set up by AWS Control Tower is managed using a set of mandatory and strongly recommended guardrails, which customers select through a self-service console to ensure that accounts and configurations comply with your policies.
- **Account Factory:** The account factory automates provisioning of new accounts in your organization. As a configurable account template, it helps you standardize provisioning of new accounts with pre-approved account configurations. You can configure your account factory with pre-approved network configuration and Region selections, enable self-service for your builders to configure and provision new accounts using AWS Service Catalog. Additionally, you can take advantage of Control Tower Solutions like Account Factory for Terraform to automate the provisioning and customization of a Control Tower-managed account that meets your business and security policies, before delivering it to end users.
- **Preventive and Detective Guardrails:** Guardrails are pre-packaged governance rules for security, operations, and compliance that customers can select and apply enterprise-wide or to specific groups of accounts. A guardrail is expressed in plain English, and enforces a specific governance policy for your AWS environment that can be enabled within an AWS Organizations organizational unit (OU). Each guardrail has two dimensions: it can be either preventive or detective, and it can be either mandatory or optional. Preventive guardrails establish intent and prevent deployment of resources that don't conform to your policies (for example, 'Enable AWS CloudTrail in all accounts'). Detective guardrails (for example, 'Detect whether public read access to Amazon S3 buckets is allowed') continuously monitor deployed resources for nonconformance. AWS Control Tower automatically translates guardrails into granular AWS policies by:
  - Establishing a configuration baseline using AWS CloudFormation

  - Preventing configuration changes of the underlying implementation using service control policies (for preventive guardrails)
  - Continuously detecting configuration changes through AWS Config rules (for detective guardrails)
  - Updating guardrail status on the AWS Control Tower dashboard
- **Mandatory and Optional Guardrails:** AWS Control Tower offers a curated set of guardrails based on AWS best practices and common customer policies for governance. You can automatically leverage mandatory guardrails as part of your landing zone setup. Some examples of mandatory guardrails include:
  - Disallow changes to AWS IAM roles set up by AWS Control Tower and AWS CloudFormation
  - Detect public read access setting for log archive
  - Disallow changes to bucket policy for AWS Control Tower created Amazon S3 buckets in log archive
  - Disallow cross-Region networking

  You can also choose to enable optional guardrails at any time. All accounts provisioned under OUs where optional guardrails are enabled will automatically inherit those guardrails. Examples of optional guardrails include:
  - Detect whether public write access to Amazon S3 buckets is allowed
  - Detect whether MFA for the root user is enabled
  - Detect whether encryption is enabled for Amazon EBS volumes attached to Amazon EC2 instances
- **Dashboard:** The AWS Control Tower dashboard gives you continuous visibility into your AWS environment. You can view the number of OUs and accounts provisioned and, the number of guardrails enabled, and check the status of your OUs and accounts against those guardrails. You can also see a list of noncompliant resources with respect to enabled guardrails.
- **Solutions for AWS Control Tower in AWS Marketplace:** AWS Marketplace now offers integrated third-party software solutions for AWS Control Tower. Built by independent software vendors, these solutions help solve infrastructure and operational use cases including security for a multi-account environment, centralized networking, operational intelligence, and Security and Information Event Management (SIEM).

119.1.2. Benefits

- **Quickly set up and configure a new AWS environment:** Automate the setup of your multi-account AWS environment with just a few clicks. The setup employs blueprints, that capture AWS best practices for configuring AWS security and management services to govern your environment. Blueprints are available to provide identity management, federate access to accounts, centralize logging, establish cross-account security audits, define workflows for provisioning accounts, and implement account baselines with network configurations.
- **Automate ongoing policy management:** AWS Control Tower provides mandatory and strongly recommended high-level rules, called guardrails, that help enforce your policies using service control policies (SCPs), or detect policy violations using AWS Config rules. These rules remain in effect as you create new accounts or make changes to existing accounts, and AWS Control Tower provides a summary report of how each account conforms to your enabled policies. For example, you can enable data residency

guardrails so that customer data, the personal data you upload to the AWS services under your AWS account, is not stored or processed outside a specific AWS Region or Regions.

- **View policy-level summaries of your AWS environment:** AWS Control Tower provides an integrated dashboard so you can see a top-level summary of policies applied to your AWS environment. You can view details on the accounts provisioned, the guardrails enabled across your accounts, and account level status for compliance with your guardrails.

119.2. Backup/Restore and Disaster Recovery

This requirement is not applicable for this service. For additional information beyond what is described herein, please refer to http://aws.amazon.com/documentation/.

119.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

119.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/controltower/index.html
- **Service quotas:** https://docs.aws.amazon.com/controltower/latest/userguide/limits.html
- **Service FAQs:** https://aws.amazon.com/controltower/faqs/

119.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/controltower/index.html and the following links for comprehensive technical documentation regarding this service.

- User Guide: Describes key concepts for AWS Control Tower. Provides instructions for setting up and using a landing zone, a secure and compliant multi-account AWS environment for enterprises or organizations at scale.

120. AWS Cost Explorer

120.1. Service Overview

AWS Cost Explorer has an easy-to-use interface that lets you visualize, understand, and manage your AWS costs and usage over time.

Get started quickly by creating custom reports that analyse cost and usage data. Analyze your data at a high level (for example, total costs and usage across all accounts) or dive deeper into your cost and usage data to identify trends, pinpoint cost drivers, and detect anomalies.

120.1.1. Features

- **Monthly Costs by AWS Service:** AWS Cost Explorer includes a default report that helps you visualize the costs and usage associated with your top five cost-accruing AWS services, and gives you a detailed breakdown on all services in the table view. The reports let you adjust the time range to view historical data going back up to twelve months to gain an understanding of your cost trends.

- **Hourly and Resource Level Granularity:** AWS Cost Explorer helps you visualize, understand, and manage your AWS costs and usage over a daily or monthly granularity. The solution also lets you dive deeper using granular filtering and grouping dimensions such as Usage Type and Tags. You can also access your data with further granularity by enabling hourly and resource level granularity.
- **Savings Plans:** Savings Plans is a flexible pricing model that provides savings of up to 72% on your AWS compute usage. This pricing model offers lower prices on Amazon EC2 instances usage, regardless of instance family, size, OS, tenancy or AWS Region, and also applies to AWS Fargate usage. AWS Cost Explorer will help you to choose a Savings Plan, and will guide you through the purchase process.

120.1.2. Benefits

- **Get started quickly:** A set of default reports are included to help you quickly gain insight into your cost drivers and usage trends.
- **Set time interval and granularity:** Set a custom time period, and determine whether you would like to view your data at a monthly or daily level of granularity.
- **Filter/Group your data:** Dig deeper into your data by taking advantage of filtering and grouping functionality, using a variety of available dimensions.
- **Forecast future costs and usage:** Use forecasting to get a better idea of what your costs and usage may look like in the future, so that you can plan ahead.
- **Save your progress:** Once you arrive at a helpful view, save your progress as a new report that you can refer back to in the future.
- **Build custom applications:** Directly access the interactive, ad-hoc analytics engine that powers AWS Cost Explorer.

120.2. Backup/Restore and Disaster Recovery

This requirement is not applicable for this service. For additional information beyond what is described herein, please refer to http://aws.amazon.com/documentation/.

120.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

120.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/account-billing/index.html
- **Service quotas:** https://docs.aws.amazon.com/general/latest/gr/billing.html
- **Service FAQs:** https://aws.amazon.com/aws-cost-management/aws-cost-explorer/faqs/

120.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/account-billing/index.html and the following links for comprehensive technical documentation regarding this service.

- AWS Billing User Guide: Describes key concepts of how to use the Billing console, integrated with the AWS Management Console.

- AWS Cost Management User Guide: Describes key concepts of the Cost Management console, and provides detailed instructions for using the various features.
- Cost and Usage Reports User Guide: Describes how to use the AWS Cost and Usage Reports feature, integrated with the AWS Billing and Cost Management console.
- API Reference for AWS Billing and Cost Management: Describes all the API operations for AWS Billing and Cost Management in detail. Also provides sample requests, responses, and errors for the supported web services protocols.

121. AWS Data Exchange (BYOS)

121.1. Service Overview

AWS Data Exchange makes it easy to find, subscribe to, and use third-party data in the cloud. Once subscribed to a data product, you can use the AWS Data Exchange API to load data directly into Amazon S3 and then analyze it with a wide variety of AWS analytics and machine learning services. For data providers, AWS Data Exchange makes it easy to reach the millions of AWS customers migrating to the cloud by removing the need to build and maintain infrastructure for data storage, delivery, billing, and entitling.

121.1.1. Features

Copy and paste features from service features page

- **Hundreds of commercial data products :** Subscribers can easily browse the AWS Data Exchange catalog to find hundreds of relevant and up-to-date commercial data products covering a wide range of industries, including financial services, healthcare, life sciences, geospatial, consumer, media & entertainment, and more.
- **And hundreds of free and open data products, too!:** The AWS Data Exchange catalog includes a large selection of free and open data products sourced from academic institutions, government entities, research institutions, and private companies.
- **Fast, direct access to data:** Because data providers have already published the data to AWS Data Exchange, subscribers can access this data (and new revisions) quickly in the cloud without the need to build custom tools and infrastructure specific to each data provider.
- **Data products are included in AWS billing:** Any applicable data subscription costs are consolidated on existing AWS invoices.
- **Analyze data immediately with AWS analytics services:** Subscribers can easily go from finding data to analyzing it using the full portfolio of AWS data lake and analytics services. Quickly export data from AWS Data Exchange to Amazon S3 and immediately start analyzing it with Amazon Athena and Amazon Redshift, build machine learning models with Amazon Sagemaker, transform and process data with Amazon EMR and AWS Glue, or build a data lake with AWS Lake Formation.
- **Easily get started as a qualified data provider:** AWS Data Exchange has step-by-step registration wizards to help data providers create a profile page, and complete registration in minutes. Providers will then be contacted by our team for a further qualification. Once qualified, providers will be able to list data products on the AWS Data Exchange catalog.
- **List data products in minutes and reach AWS customers worldwide:** Once qualified, all it takes is a few clicks or API calls to upload and package data residing in Amazon S3 buckets or on-premises environments. Step-by-step wizards help to specify

product information and publish data sets as free or paid products, with 1-month to 36-month subscription durations, using custom agreements or standard agreement templates.

- **Create Private Products available only to specific customers:** Data providers can create Private Products to issue data products exclusively to specific subscribers. These products do not appear on the public catalog and can include data that is either not meant to be generally available or is customized for a particular subscriber.
- **Make private offers of data products to specific subscribers:** Data providers can also create private offers with custom negotiated prices and/or terms to individual AWS customers.
- **Reporting and commerce analytics**: AWS Marketplace provides daily, weekly, and monthly reports detailing product subscriptions, customers, and financials. These reports are e-mailed to providers and can be accessed any time on either the AWS Marketplace Management Portal or via an API.

121.1.2. Benefits

- **Quickly find diverse data in one place:** AWS Data Exchange has hundreds of commercial data products from category-leading data providers across industries such as financial services, healthcare, retail, media & entertainment, and more. AWS Data Exchange includes hundreds of free data sets too, including data collected from popular public sources as well as trials for commercial products, so customers can explore before they subscribe. You can easily find and subscribe to data products in AWS Marketplace and stay up-to-date as providers publish new revisions in a consistent and cloud-native way.
- **Efficiently access data in the cloud:** AWS Data Exchange removes the friction of finding, licensing, and using data sets. Without AWS Data Exchange, you might spend days or even weeks licensing data and moving it where you need it to power your analytics. AWS Data Exchange simplifies access to data — eliminating the need to receive physical media, manage FTP credentials, or integrate with different APIs from multiple providers. When providers publish updates to their data sets, you will receive a notification so you can automatically consume new data as it's published.
- **Easily analyze new data:** AWS Data Exchange enables you to immediately use data you subscribe to with the full portfolio of AWS analytics services, like big data processing with Amazon EMR, data warehousing with Amazon Redshift, ad-hoc query with Amazon Athena, data integration and ETL with AWS Glue, and building data lakes with AWS Lake Formation. You can quickly copy data to Amazon S3 and immediately transform and process it, analyze it, or build machine learning models on it.

121.2. Backup/Restore and Disaster Recovery

This requirement is not applicable for this service. For additional information beyond what is described herein, please refer to http://aws.amazon.com/documentation/.

121.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

121.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/data-exchange/index.html
- **Service quotas:** https://docs.aws.amazon.com/general/latest/gr/dataexchange.html
- **Service FAQs:** https://aws.amazon.com/data-exchange/faqs/?nc=sn&loc=5&refid=ps_a134p000006gb41aae&trkcampaign=acq_paid_search_brand

121.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/data-exchange/index.html and the following links for comprehensive technical documentation regarding this service.

- AWS Data Exchange User Guide: Describes key concepts for AWS Data Exchange and provides instructions for providing, subscribing, and working with data sets. Data products are available to subscribers on AWS Marketplace as well as the AWS Data Exchange console. Data products can be published by providers using the AWS Data Exchange console and updated using AWS Marketplace Catalog API. Customers can use the AWS Data Exchange APIs or console to create, view, manage, and access data sets.
- AWS Data Exchange API Reference: Documents the AWS Data Exchange API used to create, manage, and access data sets.
- AWS Marketplace Catalog API Reference: The AWS Marketplace Catalog API Service provides an API interface for approved providers to programmatically manage their products. This includes the self-service publishing capabilities on AWS Marketplace.

122. AWS Data Pipeline

122.1. Service Overview

AWS Data Pipeline is a web service that helps you reliably process and move data between different AWS compute and storage services, as well as on-premises data sources, at specified intervals. With AWS Data Pipeline, you can regularly access your data where it's stored, transform and process it at scale, and efficiently transfer the results to AWS services such as Amazon S3, Amazon RDS, Amazon DynamoDB, and Amazon EMR.

AWS Data Pipeline helps you easily create complex data processing workloads that are fault tolerant, repeatable, and highly available. You don't have to worry about ensuring resource availability, managing inter-task dependencies, retrying transient failures or timeouts in individual tasks, or creating a failure notification system. AWS Data Pipeline also allows you to move and process data that was previously locked up in on-premises data silos.

122.1.1. Features

- **ETL Data to Amazon Redshift**: Copy RDS or DynamoDB tables to S3, transform data structure, run analytics using SQL queries and load it to Redshift.
- **ETL Unstructured Data**: Analyze unstructured data like clickstream logs using Hive or Pig on EMR, combine it with structured data from RDS and upload it to Redshift for easy querying.
- **Load AWS Log Data to Amazon Redshift**: Load log files such as from the AWS billing logs, or AWS CloudTrail, Amazon CloudFront, and Amazon CloudWatch logs, from Amazon S3 to Redshift.

- **Data Loads and Extracts**: Copy data from your RDS or Redshift table to S3 and vice-versa.
- **Move to Cloud**: Easily copy data from your on-premises data store, like a MySQL database, and move it to an AWS data store, like S3 to make it available to a variety of AWS services such as Amazon EMR, Amazon Redshift, and Amazon RDS.
- **Amazon DynamoDB Backup and Recovery**: Periodically backup your Dynamo DB table to S3 for disaster recovery purposes.

122.1.2. Benefits

- **Reliable**: AWS Data Pipeline is built on a distributed, highly available infrastructure designed for fault tolerant execution of your activities. If failures occur in your activity logic or data sources, AWS Data Pipeline automatically retries the activity. If the failure persists, AWS Data Pipeline sends you failure notifications via Amazon Simple Notification Service (Amazon SNS). You can configure your notifications for successful runs, delays in planned activities, or failures.
- **Easy to Use**: Creating a pipeline is quick and easy via our drag-and-drop console. Common preconditions are built into the service, so you don't need to write any extra logic to use them. For example, you can check for the existence of an Amazon S3 file by simply providing the name of the Amazon S3 bucket and the path of the file that you want to check for, and AWS Data Pipeline does the rest. In addition to its easy visual pipeline creator, AWS Data Pipeline provides a library of pipeline templates. These templates make it simple to create pipelines for a number of more complex use cases, such as regularly processing your log files, archiving data to Amazon S3, or running periodic SQL queries.
- **Flexible**: AWS Data Pipeline allows you to take advantage of a variety of features such as scheduling, dependency tracking, and error handling. You can use activities and preconditions that AWS provides and/or write your own custom ones. This means that you can configure an AWS Data Pipeline to take actions like run Amazon EMR jobs, execute SQL queries directly against databases, or execute custom applications running on Amazon EC2 or in your own datacenter. This allows you to create powerful custom pipelines to analyze and process your data without having to deal with the complexities of reliably scheduling and executing your application logic
- **Scalable**: AWS Data Pipeline makes it equally easy to dispatch work to one machine or many, in serial or parallel. With AWS Data Pipeline's flexible design, processing a million files is as easy as processing a single file.
- **Low Cost**: AWS Data Pipeline is inexpensive to use and is billed at a low monthly rate. You can try it for free under the AWS Free Usage. Learn more.
- **Transparent**: You have full control over the computational resources that execute your business logic, making it easy to enhance or debug your logic. Additionally, full execution logs are automatically delivered to Amazon S3, giving you a persistent, detailed record of what has happened in your pipeline.

122.2. Backup/Restore and Disaster Recovery

This requirement is not applicable for this service. For additional information beyond what is described herein, please refer to http://aws.amazon.com/documentation/.

122.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

122.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/data-pipeline/
- **Service quotas:** https://docs.aws.amazon.com/general/latest/gr/datapipeline.html
- **Service FAQs:** https://aws.amazon.com/datapipeline/faqs/

122.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/data-pipeline/ and the following links for comprehensive technical documentation regarding this service.

- Developer Guide: Provides a conceptual overview of AWS Data Pipeline and includes detailed development instructions for using the various features.
- API Reference: Describes all the API operations for AWS Data Pipeline in detail. Also provides sample requests, responses, and errors for the supported web services protocols.

123. AWS Database Migration Service

123.1. Service Overview

AWS Database Migration Service (AWS DMS) helps you migrate databases to AWS quickly and securely. The source database remains fully operational during the migration, minimizing downtime to applications that rely on the database. The AWS Database Migration Service can migrate your data to and from the most widely used commercial and open-source databases.

AWS Database Migration Service supports homogeneous migrations such as Oracle to Oracle, as well as heterogeneous migrations between different database platforms, such as Oracle or Microsoft SQL Server to Amazon Aurora. With AWS Database Migration Service, you can also continuously replicate data with low latency from any supported source to any supported target. For example, you can replicate from multiple sources to Amazon S3 to build a highly available and scalable data lake solution. You can also consolidate databases into a petabyte-scale data warehouse by streaming data to Amazon Redshift. Learn more about the supported source and target databases.

When migrating databases to Amazon Aurora, Amazon Redshift, Amazon DynamoDB or Amazon DocumentDB (with MongoDB compatibility) you can use AWS DMS free for six months.

123.1.1. Features

- **Simple to use:** AWS Database Migration Service is simple to use. There is no need to install any drivers or applications, and it does not require changes to the source database in most cases. You can begin a database migration with just a few clicks in the AWS Management Console. Once the migration has started, DMS manages all the complexities of the migration process including automatically replicating data changes that occur in the source database during the migration process. You can also use this service for continuous data replication with the same simplicity.

- **Supports widely used databases:** AWS Database Migration Service can migrate your data to and from most of the widely used commercial and open source databases. It supports homogeneous migrations such as Oracle to Oracle, as well as heterogeneous migrations between different database platforms, such as Oracle to Amazon Aurora. Migrations can be from on-premises databases to Amazon RDS or Amazon EC2, databases running on EC2 to RDS, or vice versa, as well as from one RDS database to another RDS database. It can also move data between SQL, NoSQL, and text based targets.
- **Low cost:** AWS Database Migration Service is a low cost service. You only pay for the compute resources used during the migration process and any additional log storage. Migrating a terabyte-size database can be done for as little as $3. This applies to both homogeneous and heterogeneous migrations of any supported databases. This is in stark contrast to conventional database migration methods that can be very expensive.

123.1.2. Benefits

- **Minimal downtime:** AWS Database Migration Service helps you migrate your databases to AWS with virtually no downtime. All data changes to the source database that occur during the migration are continuously replicated to the target, allowing the source database to be fully operational during the migration process. After the database migration is complete, the target database will remain synchronized with the source for as long as you choose, allowing you to switchover the database at a convenient time.
- **On-going replication:** You can set up a DMS task for either one-time migration or on-going replication. An on-going replication task keeps your source and target databases in sync. Once set up, the on-going replication task will continuously apply source changes to the target with minimal latency. All DMS features such as data validation and transformations are available for any replication task.
- **Reliable:** The AWS Database Migration Service is highly resilient and self–healing. It continually monitors source and target databases, network connectivity, and the replication instance. In case of interruption, it automatically restarts the process and continues the migration from where it stopped. Multi-AZ option allows you to have high-availability for database migration and continuous data replication by enabling redundant replication instances.

123.2. Backup/Restore and Disaster Recovery

This requirement is not applicable for this service. For additional information beyond what is described herein, please refer to http://aws.amazon.com/documentation/.

123.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

123.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/dms/
- **Service quotas:** https://docs.aws.amazon.com/general/latest/gr/dms.html

- **Service FAQs:** https://aws.amazon.com/dms/faqs/?refid=ps_a134p000006gb41aae&trkcampaign=acq_paid_search_brand

123.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/dms/ and the following links for comprehensive technical documentation regarding this service.

- User Guide: Describes all AWS Database Migration Service concepts and provides instructions on using the various features with both the console and the command line interface.
- AWS DMS Step-by-Step Database Migration Guide: Provides step-by-step walkthroughs that go through the process of migrating data to AWS. The source or target database must be on an AWS service.
- API Reference: Describes all the API operations for AWS Database Migration Service in detail. Also provides sample requests, responses, and errors for the supported web services protocols.
- AWS DMS section of the AWS CLI Reference: Documents the AWS CLI commands for AWS Database Migration Service.

124. AWS DataSync

124.1. Service Overview

AWS DataSync is a secure online data transfer service that simplifies, automates, and accelerates copying terabytes of data to and from AWS storage services. Easily migrate or replicate large data sets without having to build custom solutions or oversee repetitive tasks. DataSync can copy data between Network File System (NFS) shares, or Server Message Block (SMB) shares, Hadoop Distributed File Systems (HDFS), self-managed object storage, AWS Snowcone, Amazon Simple Storage Service (Amazon S3) buckets, Amazon Elastic File System (Amazon EFS) file systems, Amazon FSx for Windows File Server file systems, and Amazon FSx for Lustre file systems.

124.1.1. Features

- **Purpose-Built Network Protocol:** AWS DataSync employs an AWS-designed transfer protocol—decoupled from the storage protocol—to accelerate data movement. The protocol performs optimizations on how, when, and what data is sent over the network. Network optimizations performed by DataSync include incremental transfers, in-line compression, and sparse file detection, as well as in-line data validation and encryption. Connections between the local DataSync agent and the in-cloud service components are multi-threaded, maximizing performance over your Wide Area Network (WAN). A single DataSync task is capable fully utilizing 10 Gbps over a network link between your on-premises environment and AWS.
- **Automatic Infrastructure Management:** DataSync removes many of the infrastructure and management challenges you face when writing, optimizing, and managing your own copy scripts, or deploying and tuning heavyweight commercial transfer tools. Simplify your infrastructure, stay in control with built-in monitoring, and retry mechanisms to ensure successful data transfers.

- **Bandwidth Optimization and Control:** Transferring hot or cold data should not impede your business. DataSync is equipped with granular controls to optimize bandwidth consumptions. Throttle transfer speeds up to 10 Gbps during off hours and set limits when network availability is needed elsewhere.
- **Data Transfer Scheduling:** DataSync comes with a built-in scheduling mechanism, allowing you to periodically run data transfer tasks to detect and copy changes from your source storage system to the destination. You can schedule your tasks using the AWS DataSync Console or AWS Command Line Interface (CLI) without writing scripts to manage repeated transfers. Task scheduling automatically runs tasks on your configured schedule with hourly, daily, or weekly options provided directly in the AWS Console.
- **Data Encryption and Validation:** All your data is encrypted in transit with Transport Layer Security (TLS). DataSync supports using default encryption for Amazon S3 buckets, Amazon EFS file system encryption of data at rest, Amazon FSx for Windows File Server, and Amazon FSx for Lustre encryption at rest and in transit. DataSync ensures that your data arrives intact. For each transfer, the service performs integrity checks both in transit and at rest. These checks ensure that the data written to your destination matches the data read from your source, validating consistency.
- **File System Integration and Metadata Preservation:** The DataSync agent connects to your existing storage systems using the industry-standard NFS and SMB protocols, to your Hadoop cluster as an HDFS client, or to your self-managed object storage, using the Amazon S3 application programming interface (API). The agent transfers data rapidly and writes it into your designated Amazon S3 bucket, Amazon EFS file system, Amazon FSx for Windows File Server file system, or Amazon FSx for Lustre file system. File permissions and metadata are preserved when copying objects and or data between Amazon S3, Amazon EFS, Amazon FSx for Windows File Server, or Amazon FSx for Lustre file system. When copying data to Amazon S3, DataSync automatically converts each file to a single S3 object in a 1:1 relationship, and preserves POSIX metadata from NFS shares or HDFS as Amazon S3 object metadata. When you copy objects containing file system metadata back to file formats, the original file metadata (that DataSync copied to S3) is restored.
- **Integration with AWS Infrastructure and Management Services:** DataSync works natively with AWS security, monitoring, and audit services to simplify data movement and to provide a consistent management experience for your IT, storage, and DevOps teams. In addition to integrations with Amazon S3, Amazon EFS, Amazon FSx for Windows File Server, and Amazon FSx for Lustre, DataSync supports AWS Virtual Private Cloud (VPC) endpoints (powered by AWS PrivateLink) to move files directly into your Amazon VPC. Like other AWS services, you can use AWS Identity and Access Management (IAM) to securely manage DataSync access. Similarly, you can configure an IAM role to control the services accessing your Amazon S3 bucket.
- **Monitoring and Auditing with Amazon CloudWatch and AWS CloudTrail:** With Amazon CloudWatch, you can monitor the status of any DataSync transfers currently in progress and check previous data transfer history. With CloudWatch Metrics, you can see the number of files and amount of data copied. Consult CloudWatch Logs for information about individual files transferred at a given time, as well as the results of DataSync integrity verification. This simplifies monitoring, reporting, and troubleshooting, enabling you to provide timely updates to stakeholders. In addition, CloudWatch Events are triggered as your transfer tasks complete, enabling automation of dependent

workflows. For audit purposes, you can consult AWS CloudTrail, which logs all actions performed by DataSync.

- **Pay-As-You-Go Pricing:** With AWS DataSync, you pay only for data copied by the service at a flat, per-gigabyte rate. No software licenses, contracts, maintenance fees, development cycles, or hardware are required. This provides a lower total cost of ownership (TCO) compared to manually building, operating, and optimizing your own high-performance scripted transfers, as well as lower total cost than buying and running commercial transfer tools.

124.1.2. Benefits

- **Secure Migration:** Securely migrate your data to AWS with end-to-end security, including data encryption and data integrity validation.
- **Reduce on-premises data movement costs:** Reduce expensive on-premises data movement costs with a fully managed service that seamlessly scales as data loads increase.
- **Easily manage data movement:** Easily manage data movement workloads with bandwidth throttling, migration scheduling, and task filtering.
- **Rapid Migration:** Rapidly migrate file and object data to the cloud for data replication or archival.

124.2. Backup/Restore and Disaster Recovery

This requirement is not applicable for this service. For additional information beyond what is described herein, please refer to http://aws.amazon.com/documentation/.

124.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

124.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/datasync/index.html
- **Service quotas:** https://docs.aws.amazon.com/datasync/latest/userguide/datasync-limits.html
- **Service FAQs:** https://aws.amazon.com/datasync/faqs/

124.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/datasync/index.html and the following links for comprehensive technical documentation regarding this service.

- AWS DataSync User Guide: Describes all AWS DataSync concepts and provides instructions on using the various features with the console, the command line interface, and the AWS DataSync API.

125. AWS DeepRacer

125.1. Service Overview

AWS DeepRacer gives you an interesting and fun way to get started with reinforcement learning (RL). RL is an advanced machine learning (ML) technique that takes a very different approach to training models than other machine learning methods. Its super power is that it learns very complex behaviors without requiring any labeled training data, and can make short term decisions while optimizing for a longer-term goal.

125.1.1. Features

- **Simulator:** Build models in Amazon SageMaker and train, test, and iterate quickly and easily on the track in the AWS DeepRacer 3D racing simulator.
- **Car:** Experience the thrill of the race in the real-world when you deploy your reinforcement learning model onto AWS DeepRacer.
- **League:** Compete in the world's first global, autonomous racing league, to race for prizes and glory and a chance to advance to the Championship Cup.

125.1.2. Benefits

- **A fun way to learn machine learning:** Get started with machine learning quickly with hands-on tutorials that help you learn the basics of machine learning, start training reinforcement learning models and test them in an exciting, autonomous car racing experience.
- **Experiment and grow:** Test these new found skills in the AWS DeepRacer 3D racing simulator. Experiment with multiple sensor inputs, the latest reinforcement learning algorithms, neural network configurations and simulation to-real domain transfer methods.
- **Community and competition:** The AWS DeepRacer League provides an opportunity for you to compete for prizes and meet fellow machine learning enthusiasts, online and in person. Share ideas and insights on how to succeed and create your own private virtual race.

125.2. Backup/Restore and Disaster Recovery

This requirement is not applicable for this service. For additional information beyond what is described herein, please refer to http://aws.amazon.com/documentation/.

125.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

125.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/deepracer/
- **Service quotas:** https://docs.aws.amazon.com/general/latest/gr/deepracer.html
- **Service FAQs:** https://aws.amazon.com/deepracer/faqs/?nc=sn&loc=8&refid=ps_a134p000006gb41aae&trkcampaign=acq_paid_search_brand

125.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/deepracer/ and the following links for comprehensive technical documentation regarding this service.

- Developer Guide: Learn how to train and evaluate a reinforcement learning model for autonomous driving in simulation. The tasks include how define a reward function, customize the action space, set up the training configuration, and creating a training job using Amazon SageMaker and AWS RoboMaker. Find out how to submit a trained model to the AWS DeepRacer League to race against other AWS DeepRacer users in the online Virtual Circuit. Learn how to set up and calibrate your AWS DeepRacer vehicle, upload a trained AWS DeepRacer model to the vehicle, build a physical track for the vehicle to drive autonomously.
- User Guide: AWS DeepRacer Student is a free service for students currently enrolled in high school or a university or community college. Using AWS DeepRacer, you can train a reinforcement learning model that will attempt to drive autonomously around a racetrack. The AWS DeepRacer Student League is a place where you can join up with your friends and compete for prizes. Students also have the opportunity to win Udacity scholarships through the AWS AI & ML Scholarship program.

126. AWS Device Farm

126.1. Service Overview

AWS Device Farm is an application testing service that lets you improve the quality of your web and mobile apps by testing them across an extensive range of desktop browsers and real mobile devices; without having to provision and manage any testing infrastructure. The service enables you to run your tests concurrently on multiple desktop browsers or real devices to speed up the execution of your test suite, and generates videos and logs to help you quickly identify issues with your app.

126.1.1. Features

- **Automated Testing:** Test your app in parallel against a massive collection of physical devices in the AWS Cloud. Use one of our built-in frameworks, to test your applications without having to write or maintain test scripts, or use one of our supported automation testing frameworks.
- **Remote Access:** Gesture, swipe, and interact with devices in real time, directly from your web browser.
- **Testing on desktop browsers:** Run your Selenium tests in parallel on multiple versions of Chrome, Internet Explorer, and Firefox, that are hosted in the AWS Cloud.

126.1.2. Benefits

- **Use the same devices your customers use:** Run tests and interact with a large selection of physical devices. Unlike emulators, physical devices give you a more accurate understanding of the way users interact with your app by taking into account factors like memory, CPU usage, location, and modifications made by manufactures and carriers to the firmware and software. We are always adding devices to the fleet. See the device list.
- **Reproduce and fix issues faster:** Manually reproduce issues and run automated tests in parallel. We collect videos, logs, and performance data so you can dive deep and

solve problems quickly. For automated tests, we'll identify and group issues so you can focus on the most important problems first.

- **Simulate real-world environments:** Fine-tune your test environment by configuring location, language, network connection, application data, and installing prerequisite apps to simulate real-world customer conditions.
- **Choose the tests that work for you:** Run our built-in test suite (no scripting required) or customize your tests by selecting from open-source test frameworks like Appium, Calabash, and Espresso (see supported frameworks). You can also perform manual tests with Remote Access.
- **Integrate with your development workflow:** Use our service plugins and API to automatically initiate tests and get results from IDEs and continuous integration environments like Android Studio and Jenkins.
- **Setup your own private device lab in the cloud:** Our private device lab offering lets you choose iOS and Android devices for your exclusive use. Device Farm provisions these devices with the exact configurations you need, and lets you persist settings between sessions. Since these devices are exclusively for your use, you don't have to wait for other users to finish using them.
- **Execute your tests concurrently on multiple browser instances:** Device Farm's fully managed browser grid scales as needed allowing you to run multiple tests in parallel to speed up the execution of your test suite. With pay-as-you-go pricing you can run multiple tests concurrently without worrying about incurring any additional costs as you scale - you only pay for the total number of minutes your test executes for.
- **Identify and Debug issues quickly:** Use videos, console logs, action logs, and web driver logs generated by Device Farm to identify, analyze, and quickly fix issues with your web app.
- **Test on multiple desktop browsers and browser versions:** Run your tests on multiple desktop browsers, including Chrome, Firefox, and Internet Explorer, to ensure your web app functions as expected in different browser environments.

126.2. Backup/Restore and Disaster Recovery

This requirement is not applicable for this service. For additional information beyond what is described herein, please refer to http://aws.amazon.com/documentation/.

126.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

126.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/devicefarm/index.html
- **Service quotas:** https://docs.aws.amazon.com/general/latest/gr/devicefarm.html
- **Service FAQs:** https://aws.amazon.com/device-farm/faqs/

126.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/devicefarm/index.html and the following links for comprehensive technical documentation regarding this service.

- Developer Guide: Provides a conceptual overview of Device Farm and includes detailed instructions for using the service.
- Guide for Desktop Browser Testing: Provides a conceptual overview and detailed instructions for testing desktop web applications with Device Farm.
- API Reference: Describes all the Device Farm operations in detail. Also provides sample requests, responses, and errors for the supported web services protocols.

127. AWS Digital Investigation and Forensics

127.1. Service Overview

A core mission for Justice and Public Safety customers, and other organisations undertaking investigations is effectively ingesting and managing the scale of digital intelligence and evidence submitted to them or acquired in their workflows.

AWS Digital Investigation and Forensics Storage enables customers to quickly establish a secure landing zone in the cloud and migrate forensic storage and archive data to AWS object storage service that offers industry-leading scalability, data availability, security, performance and durability.

Through a range of AWS services, customers can implement effective tagging and lifecycle management processes to ensure compliance with data retention, deletion and retrieval requirements, whilst ensuring data is managed cost effectively throughout the investigative and archive lifecycle.

127.1.1. Features

- **Elastic, Web-Scale Computing:** Amazon EC2 enables you to increase or decrease capacity within minutes, not hours or days. You can commission one, hundreds, or even thousands of server instances simultaneously.
- **Durable:** Amazon S3 provides durable infrastructure to store important data and is designed for durability of 99.999999999% of objects.
- **Low Cost:** Amazon S3 allows you to store large amounts of data at a very low cost. You pay for what you need, with no minimum commitments or upfront fees.
- **Available:** Amazon S3 is designed for 99.99% availability of objects over a given year.
- **Secure:** Amazon S3 supports data transfer over SSL and automatic encryption of your data once it is uploaded.
- **Landing zone:** A landing zone is a well-architected, multi-account environment that's based on security and compliance best practices. It is the enterprise-wide container that holds all of your organizational units (OUs), accounts, users, and other resources that you want to be subject to compliance regulation. A landing zone can scale to fit the needs of an enterprise of any size.

127.1.2. Benefits

Automate the setup of your multi-account AWS environment with just a few clicks. The setup employs blueprints that capture AWS best practices for configuring AWS security and management services to govern your environment. Blueprints are available to provide identity

management, federate access to accounts, centralize logging, establish cross-account security audits, define workflows for provisioning accounts, and implement account baselines with network configurations.

Object storage service from Amazon Simple Storage Service (Amazon S3) offering industry-leading scalability, data availability, security, and performance. Customers of all sizes and industries can store and protect any amount of data for virtually any use case, such as data lakes, cloud-native applications, and mobile apps. With cost-effective storage classes and easy-to-use management features, you can optimize costs, organize data, and configure fine-tuned access controls to meet specific business, organizational, and compliance requirements.

127.2. Backup/Restore and Disaster Recovery

Amazon S3 offers a highly durable, scalable, and secure solution for backing up and archiving your critical data. You can use Amazon S3's versioning capability to provide even further protection for your stored data. You can also define rules to archive sets of Amazon S3 objects to Amazon Glacier's extremely low-cost storage service based on object lifetimes. As your data ages, these rules enable you to ensure that it is automatically stored on the storage option that is most cost effective for your needs.

127.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

127.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:**
    a. https://docs.aws.amazon.com/ec2/index.html
    b. https://docs.aws.amazon.com/s3/index.html
    c. https://docs.aws.amazon.com/controltower/
- **Service quotas:**
    a. https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ec2-resource-limits.html
    b. https://docs.aws.amazon.com/general/latest/gr/s3.html
    c. https://docs.aws.amazon.com/controltower/latest/userguide/limits.html
- **Service FAQs:**
    a. https://aws.amazon.com/ec2/faqs/
    b. https://aws.amazon.com/s3/faqs/
    c. https://aws.amazon.com/controltower/faqs/

127.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/ec2/index.html, https://docs.aws.amazon.com/s3/index.html, https://docs.aws.amazon.com/controltower/ for comprehensive technical documentation regarding this service.

128. AWS Direct Connect

128.1. Service Overview

The AWS Direct Connect cloud service is the shortest path to your AWS resources. While in transit, your network traffic remains on the AWS global network and never touches the public internet. This reduces the chance of hitting bottlenecks or unexpected increases in latency. When creating a new connection, you can choose a hosted connection provided by an AWS Direct Connect Delivery Partner, or choose a dedicated connection from AWS—and deploy at over 100 AWS Direct Connect locations around the globe. With AWS Direct Connect SiteLink, you can send data between AWS Direct Connect locations to create private network connections between the offices and data centers in your global network.

128.1.1. Features

- **AWS Direct Connect Locations Worldwide:** AWS Direct Connect is available at locations worldwide to ensure you can make connections close to where you need them. Some AWS Direct Connect features, such as MACsec and 100 Gbps connections, are available at select locations.
- **Connection speeds up to 100 Gbps:** AWS Direct Connect is available in speeds starting at 50 Mbps and scaling up to 100 Gbps, so you can find the right connection for you.
- **MACsec and IPsec Encryption options:** Add extra protection to communications between your data centers, branch offices, or colocation facilities with multiple encryption options. Secure your 10 Gbps and 100 Gbps connections with native IEEE 802.1AE (MACsec) point-to-point encryption at select locations. AWS Site-to-Site VPN is also available for secure connections using IPsec (IP security).
- **SiteLink:** AWS Direct Connect SiteLink creates private, end-to-end network connections between the offices, data centers, and colocation facilities in your global network. Once you have made connections at two or more AWS Direct Connect locations, you can turn the SiteLink feature on (or off) with a simple configuration change using the AWS Management Console, WS Command Line Interface (CLI), or APIs. In minutes, a global, reliable, and private network is ready for use.
- **Multiple deployment options:** Dedicated Connections create links to AWS using a 1 Gbps, 10 Gbps, or 100 Gbps Ethernet port. AWS Direct Connect Partners provide Hosted connections using pre-established network links between themselves and AWS, and are available from 50 Mbps up to 10 Gbps.

128.1.2. Benefits

- **High Performance:** Improve application performance by connecting directly to AWS and bypassing the public internet.
- **Secures moving data:** Secure your data as it moves between your network and AWS with multiple encryption options.
- **Reduces costs of networking:** Reduce your networking costs with low data transfer rates out of AWS.
- **Build hybrid networks:** Link your AWS and on-premises networks to build applications that span environments without compromising performance.

- **Extend your existing network:** Once you link your network to AWS Direct Connect, you can use SiteLink to send data between your locations. When using SiteLink, data travels over the shortest path between locations.
- **Manage large datasets:** Ensure smooth and reliable data transfers at massive scale for real-time analysis, rapid data backup, or broadcast media processing.

128.2. Backup/Restore and Disaster Recovery

This requirement is not applicable for this service. For additional information beyond what is described herein, please refer to http://aws.amazon.com/documentation/.

128.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

128.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/directconnect/index.html
- **Service quotas:** https://docs.aws.amazon.com/directconnect/latest/UserGuide/limits.html
- **Service FAQs:** https://aws.amazon.com/directconnect/faqs/?nc=sn&loc=6

128.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/directconnect/index.html and the following links for comprehensive technical documentation regarding this service.

- User Guide: Describes key concepts and provides instructions for using the features of AWS Direct Connect.
- API Reference: Describes the API operations for AWS Direct Connect.
- Direct Connect section of the AWS CLI Reference: Describes the AWS CLI commands for AWS Direct Connect.

129. AWS Directory Service

129.1. Service Overview

AWS Directory Service for Microsoft Active Directory, also known as AWS Managed Microsoft Active Directory (AD), enables your directory-aware workloads and AWS resources to use managed Active Directory (AD) in AWS. AWS Managed Microsoft AD is built on actual Microsoft AD and does not require you to synchronize or replicate data from your existing Active Directory to the cloud. You can use the standard AD administration tools and take advantage of the built-in AD features, such as Group Policy and single sign-on. With AWS Managed Microsoft AD, you can easily join Amazon EC2 and Amazon RDS for SQL Server instances to your domain, and use AWS End User Computing (EUC) services, such as Amazon WorkSpaces, with AD users and groups.

129.1.1. Features

- **Actual Microsoft Active Directory:** AWS Managed Microsoft AD is actual Microsoft Active Directory (AD) running on AWS-managed infrastructure. This enables you to

administer your users and devices in AWS Managed Microsoft AD by using the tools you already know, such as Active Directory Administrative Center and Active Directory Users and Computers.

- **High availability:** Because directories are mission-critical infrastructure, AWS Managed Microsoft AD is deployed in high availability and across multiple Availability Zones. You can also scale out your AWS Managed Microsoft AD directory by deploying additional domain controllers to increase the resiliency of your managed directory for even higher availability.
- **AWS-managed infrastructure:** AWS Managed Microsoft AD runs on AWS managed infrastructure with monitoring that automatically detects and replaces domain controllers that fail. In addition, data replication and automated daily snapshots are configured for you. You do not need to install software, and AWS handles all of the patching and software updates.
- **Multi-region replication:** Multi-region replication enables you to deploy and use a single AWS Managed Microsoft AD directory across multiple AWS Regions. This makes it easier and more cost-effective for you to deploy and manage your Microsoft Windows and Linux workloads globally. With the automated multi-region replication capability, you get higher resiliency, while your applications use a local directory for optimal performance.
- **HIPAA and PCI Eligible:** You can use AWS Managed Microsoft AD to build and run AD–aware cloud applications that are subject to the U.S. Health Insurance Portability and Accountability Act (HIPAA) or Payment Card Industry Data Security Standard (PCI DSS) compliance. AWS Managed Microsoft AD reduces the effort required to deploy compliant AD infrastructure for your cloud applications, as you manage your own HIPAA risk management programs or PCI DSS compliance certification.
- **Trust support:** You can easily integrate AWS Managed Microsoft AD with your existing AD by using AD trust relationships. Using trusts enables you to use your existing Active Directory to control which AD users can access your AWS resources.
- **Group-based policies:** AWS Managed Microsoft AD allows you to manage users and devices using native Active Directory Group Policy objects (GPOs). You can create GPOs with existing tools, such as the Group Policy Management Console (GPMC).
- **Single sign-on (SSO):** AWS Managed Microsoft AD uses the same Kerberos-based authentication as your existing on-premises AD. By integrating your AWS resources with AWS Managed Microsoft AD, your AD users will be able to sign in with SSO to AWS applications and resources using a single set of credentials.
- **Seamless domain join:** AWS Managed Microsoft AD enables you to use seamless domain join for new and existing Amazon EC2 for Windows Server and Amazon EC2 for Linux instances. For new EC2 instances, you can choose which domain to join at launch time by using the AWS Management Console. You can use seamless domain join for existing EC2 instances by using the EC2Config service. Amazon EC2 instances can also join to a single shared directory from any AWS account and any Amazon VPC within a Region.
- **Single directory for all directory-aware workloads:** AWS Managed Microsoft AD enables you to use a single directory for your directory-aware workloads in AWS resources such as Amazon EC2 instances, Amazon RDS for SQL Server instances, and AWS End User Computing services, such as Amazon WorkSpaces. Sharing a directory allows your directory-aware workloads to easily manage Amazon EC2 instances across

multiple AWS accounts and Amazon VPCs within a Region. It also helps avoid the complexity of replicating and synchronizing data across multiple directories.

129.1.2. Benefits

- **Easily migrate directory-aware, on-premises workloads:** AWS Managed Microsoft AD makes it easy to migrate AD-dependent applications and Windows workloads to AWS. With AWS Managed Microsoft AD, you can use Group Policies to manage EC2 instances and run AD-dependent applications in the AWS Cloud without the need to deploy your own AD infrastructure.
- **Use actual Microsoft Active Directory (AD):** Take advantage of actual Microsoft Active Directory to manage your users, groups, and devices. Use familiar AD administration tools and features, such as Group Policy objects (GPOs), domain trusts, fine-grain password policies, group Managed Service Account (gMSA), schema extensions, and Kerberos-based single sign-on. You can also delegate administrative tasks and authorize access using AD security groups.
- **Share a single directory for cloud workloads:** Share a single directory for all your AD-aware Amazon EC2 instances, Amazon RDS for SQL Server instances, and AWS End User Computing services, such as Amazon WorkSpaces. You can also share your AD with multiple accounts. Using AWS Managed Microsoft AD helps avoid the complexity of replicating and synchronizing data across multiple directories.
- **Easily extend existing domains:** AWS Managed Microsoft AD makes it easy to extend your existing Active Directory to AWS. It enables you to leverage your existing on-premises user credentials to access cloud resources such as the AWS Management Console, Amazon Workspaces, Amazon Chime, and Windows workloads in the cloud.
- **Centrally manage application access and devices in AWS:** AWS Managed Microsoft AD provides you the option to administer your on-premises users, groups, applications, and systems without the complexity of running and maintaining an on-premises, highly available AD. You can easily join your existing computers, laptops, and printers to an AWS Managed Microsoft AD domain.
- **Simplify administration with a managed service:** AWS Managed Microsoft AD is built on highly available, AWS-managed infrastructure. Each directory is deployed across multiple Availability Zones, and monitoring automatically detects and replaces domain controllers that fail. In addition, data replication and automated daily snapshots are configured for you. You do not have to install software, and AWS handles all patching and software updates.

129.2. Backup/Restore and Disaster Recovery

This service allows for backup and recovery. The service can back up directory snapshots. Users control this via a manual or scheduled API call. Users schedule and recover backups through a web interface.

129.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace

129.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/directory-service/
- **Service quotas:** https://docs.aws.amazon.com/general/latest/gr/ds_region.html
- **Service FAQs:** https://aws.amazon.com/directoryservice/faqs/?nc=sn&loc=5

129.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/directory-service/ and the following links for comprehensive technical documentation regarding this service.

- Directory Service Administration Guide: Describes how to create and manage an AWS Directory Service directory.
- Directory Service API Reference: Describes the API operations for AWS Directory Service.
- AWS Directory Service in the AWS CLI: Describes the Directory Service commands in the AWS Command Line Interface.

130. AWS Elastic Beanstalk

130.1. Service Overview

AWS Elastic Beanstalk is an easy-to-use service for deploying and scaling web applications and services developed with Java, .NET, PHP, Node.js, Python, Ruby, Go, and Docker on familiar servers such as Apache, Nginx, Passenger, and IIS.

You can simply upload your code and Elastic Beanstalk automatically handles the deployment, from capacity provisioning, load balancing, auto-scaling to application health monitoring. At the same time, you retain full control over the AWS resources powering your application and can access the underlying resources at any time.

There is no additional charge for Elastic Beanstalk - you pay only for the AWS resources needed to store and run your applications.

130.1.1. Features

- **Wide selection of application platforms:** AWS Elastic Beanstalk supports web applications written in many popular languages and frameworks. It requires no or minimal code changes to go from development machine to the cloud. Development options for deploying your web applications include Java, .NET, Node.js, PHP, Ruby, Python, Go, and Docker.
- **Variety of application deployment options:** With AWS Elastic Beanstalk, you can deploy your code through the AWS Management Console, Elastic Beanstalk Command Line Interface, Visual Studio, and Eclipse. Multiple deployment policies—all at once, rolling, rolling with an additional batch, immutable, and blue/green—offer choices for the speed and safety of deploying your applications while reducing the administrative burden.
- **Monitoring:** Elastic Beanstalk provides a unified user interface (UI) to monitor and manage the health of your applications.
    - **Application Health**: Elastic Beanstalk collects 40+ key metrics and attributes to determine the health of your applications. With the Elastic Beanstalk Health Dashboard, you can visualize overall application health and customize application health checks, health permissions, and health reporting in one UI.

  - **Monitoring, Logging, and Tracing**: Elastic Beanstalk integration with Amazon CloudWatch and AWS X-Ray means you can use monitoring dashboards to view key performance metrics such as latency, CPU utilization, and response codes. You can also set up CloudWatch alarms to get notified when metrics exceed your chosen thresholds.
- **Updates and management:** You can choose to automatically get the latest platform versions of your Elastic Beanstalk environment and new patches using managed platform updates. An immutable deployment mechanism ensures these updates are implemented safely. For ongoing management, you can also customize application properties, create alarms, and enable e-mail notifications via Amazon Simple Notification Service (Amazon SNS).
- **Scaling:** Elastic Beanstalk uses Elastic Load Balancing and Auto Scaling to automatically scale your application in and out based on its specific needs. Multiple availability zones give you an option to improve application reliability and availability.
- **Customization:** With Elastic Beanstalk, you have the freedom to select the AWS resources, such as Amazon EC2 instance type including Spot instances, that are optimal for your application. You also retain full control over the AWS resources powering your application. If you decide you want to take over some (or all) of the elements of your infrastructure, you can do so seamlessly by using Elastic Beanstalk's management capabilities.
- **Compliance:** Elastic Beanstalk meets the criteria for ISO, PCI, SOC 1, SOC 2, and SOC 3 compliance along with the criteria for HIPAA eligibility. This means applications running on Elastic Beanstalk can process regulated financial data or protected health information (PHI).
- **AWS Graviton support:** AWS Graviton arm64-based processors deliver the best price performance for your cloud workloads running in Amazon EC2. With AWS Graviton on Elastic Beanstalk, you can select EC2 instance types to meet optimization needs of your workloads and benefit from improved price performance over a comparable x86-based processor.

130.1.2. Benefits

- **Fast and simple to begin:** Elastic Beanstalk is the fastest and simplest way to deploy your application on AWS. You simply use the AWS Management Console, a Git repository, or an integrated development environment (IDE) such as Eclipse or Visual Studio to upload your application, and Elastic Beanstalk automatically handles the deployment details of capacity provisioning, load balancing, auto-scaling, and application health monitoring. Within minutes, your application will be ready to use without any infrastructure or resource configuration work on your part.
- **Developer productivity:** Elastic Beanstalk provisions and operates the infrastructure and manages the application stack (platform) for you, so you don't have to spend the time or develop the expertise. It will also keep the underlying platform running your application up-to-date with the latest patches and updates. Instead, you can focus on writing code rather than spending time managing and configuring servers, databases, load balancers, firewalls, and networks.
- **Impossible to outgrow:** Elastic Beanstalk automatically scales your application up and down based on your application's specific need using easily adjustable Auto Scaling settings. For example, you can use CPU utilization metrics to trigger Auto Scaling

actions. With Elastic Beanstalk, your application can handle peaks in workload or traffic while minimizing your costs.

- **Complete resource control:** You have the freedom to select the AWS resources, such as Amazon EC2 instance type and processor type to run the workload on, that are optimal for your application. You also retain full control over the AWS resources powering your application. If you decide you want to take over some (or all) of the elements of your infrastructure, you can do so seamlessly by using Elastic Beanstalk's management capabilities.

130.2. Backup/Restore and Disaster Recovery

This service allows for backup and recovery. The service can back up machine images and volumes. Users control this via a manual or scheduled API call. Users schedule and recover backups through a web interface.

130.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace

130.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/elastic-beanstalk/index.html
- **Service quotas:** https://docs.aws.amazon.com/general/latest/gr/elasticbeanstalk.html
- **Service FAQs:** https://aws.amazon.com/elasticbeanstalk/faqs/

130.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/elastic-beanstalk/index.html and the following links for comprehensive technical documentation regarding this service.

- Developer Guide: Read conceptual and detailed instructions for using AWS Elastic Beanstalk to quickly deploy and manage applications in the AWS cloud without worrying about the infrastructure that runs those applications.
- API Reference: Get formal descriptions of all the API operations for AWS Elastic Beanstalk. In addition, find sample requests, responses, and errors for the supported web services protocols.
- Getting Started Walkthrough: Walks developers through the use of the console to create, view, deploy, and update an application for the first time, as well as the steps for editing and terminating an environment.
- Platforms: Find detailed listings of current AWS Elastic Beanstalk platform versions. Also find lists of historical platform versions and the date ranges they were current.
- Release Notes: Find details about AWS Elastic Beanstalk releases—new features, updates, and fixes related to the service, platform, console, and EB CLI.

131. AWS Elastic Disaster Recovery

131.1. Service Overview

AWS Elastic Disaster Recovery (AWS DRS) minimizes downtime and data loss with fast, reliable recovery of on-premises and cloud-based applications using affordable storage, minimal compute, and point-in-time recovery.

Set up AWS Elastic Disaster Recovery on your source servers to initiate secure data replication. Your data is replicated to a staging area subnet in your AWS account, in the AWS Region you select. The staging area design reduces costs by using affordable storage and minimal compute resources to maintain ongoing replication. You can perform non-disruptive tests to confirm that implementation is complete. During normal operation, maintain readiness by monitoring replication and periodically performing non-disruptive recovery and failback drills. If you need to recover applications, you can launch recovery instances on AWS within minutes, using the most up-to-date server state or a previous point in time. After your applications are running on AWS, you can choose to keep them there, or you can initiate data replication back to your primary site when the issue is resolved. You can fail back to your primary site whenever you're ready.

131.1.1. Features

- **On-premises to AWS:** Quickly recover operations after unexpected events such as software issues or datacenter hardware failures. AWS DRS enables RPOs of seconds and RTOs of minutes.
- **Cloud to AWS:** Help increase resilience and meet compliance requirements using AWS as your recovery site. AWS DRS converts your cloud-based applications to run natively on AWS.
- **AWS Region to AWS Region:** Increase application resilience and help meet availability goals for your AWS-based applications, using AWS DRS to recover applications in a different AWS Region.

131.1.2. Benefits

- **Saves cost:** Save costs by removing idle recovery site resources, and pay for your full disaster recovery site only when needed.
- **Recover applications in Minutes:** Recover your applications within minutes, at their most up-to-date state or from a previous point in time.
- **Easy to use:** Use a unified process to test, recover, and fail back a wide range of applications, without specialized skillsets.
- **Flexible:** Gain flexibility using AWS as your elastic recovery site, with the ability to add or remove replicating servers as needed.

131.2. Backup/Restore and Disaster Recovery

This service allows for backup and recovery. The service can back up on-premises or cloud-based applications running on supported operating systems. Users control this via the AWS Management Console to configure replication and launch settings, monitor data replication, and launch instances for drills or recovery.

131.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

131.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/drs/index.html
- **Service quotas:** https://docs.aws.amazon.com/general/latest/gr/drs.html
- **Service FAQs:** https://aws.amazon.com/disaster-recovery/faqs/?nc=sn&loc=4

131.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/drs/index.html and the following links for comprehensive technical documentation regarding this service.

- User Guide: Learn how to set up and use AWS Elastic Disaster Recovery.
- API Reference: Describes all the API operations for AWS Elastic Disaster Recovery in detail.

132. AWS Elemental MediaConnect

132.1. Service Overview

AWS Elemental MediaConnect is a high-quality transport service for live video. Today, broadcasters and content owners rely on satellite networks or fibre connections to send their high-value content into the cloud or to transmit it to partners for distribution. Both satellite and fibre approaches are expensive, require long lead times to set up, and lack the flexibility to adapt to changing requirements. To be more nimble, some customers have tried to use solutions that transmit live video on top of IP infrastructure, but have struggled with reliability and security.

Now you can get the reliability and security of satellite and fibre combined with the flexibility, agility, and economics of IP-based networks using AWS Elemental MediaConnect. MediaConnect enables you to build mission-critical live video workflows in a fraction of the time and cost of satellite or fibre services. You can use MediaConnect to ingest live video from a remote event site (like a stadium), share video with a partner (like a cable TV distributor), or replicate a video stream for processing (like an over-the-top service). MediaConnect combines reliable video transport, highly secure stream sharing, and real-time network traffic and video monitoring that allow you to focus on your content, not your transport infrastructure.

For the highest-quality, lowest-latency workflows, MediaConnect supports AWS Cloud Digital Interface (AWS CDI) for transporting uncompressed video between applications in AWS. You can also use JPEG XS encode and decode in MediaConnect to transport high-quality video to and from the AWS Cloud. For hybrid broadcast control, live production, and other uncompressed live video workflows, you can build cloud-based applications that work seamlessly with on-premises infrastructure, including AWS Elemental Live for JPEG XS encoding and decoding.

132.1.1. Features

- **Comprehensive range of video industry standard protocols:** MediaConnect supports a range of protocols for video delivery including the Zixi protocol, Reliable Internet Stream Transport (RIST), Secure Reliable Transport (SRT), Real-Time Transport Protocol (RTP), and RTP with forward error correction (FEC). Using these protocols, MediaConnect enables a quality-of-service layer over IP network transport,

which maintains stream integrity through packet recovery, for reliable live video contribution into the AWS Cloud using the public internet or AWS Direct Connect.

- **Transport uncompressed and visually-lossless video with AWS Cloud Digital Interface (AWS CDI) and JPEG XS:** MediaConnect can route video, audio, and metadata streams using AWS CDI to provide uncompressed, low-latency video to CDI-enabled applications. MediaConnect can also receive and decode JPEG XS for contributing live video to applications in the AWS Cloud, and can encode and send AWS CDI content as JPEG XS for return-feed applications back to on-premises infrastructure.
- **Built-in security:** MediaConnect lets you protect your content using industry-standard, end-to-end AES encryption, and you can enable whitelisting to limit access only to trusted sources. MediaConnect is fully integrated with AWS Secrets Manager for key storage and retrieval. You can even configure unique encryption keys for each destination as your video leaves MediaConnect, giving you control over your content's security.
- **High-quality video sharing:** Create entitlements in MediaConnect to grant another account access to your content, so you can share your video streams with your customers and partners. Instead of hosting distribution hubs in your data center to share live video, you can now reach a larger number of customers quickly and easily using the AWS Cloud. Content sharing is tracked using entitlements, and you can revoke access at any time. Accounts that are granted access can create their own video workflows using MediaConnect and other AWS services.
- **Monitor in real time:** Using AWS monitoring services, such as Amazon CloudWatch, MediaConnect provides at-a-glance network performance metrics. Broadcast-standard alerts identify issues with transport streams, so you can adjust settings to maximize the quality of your video workflows. The service monitors stream performance and recovers automatically without operator intervention.
- **Use independently or with other AWS services:** You can choose to use MediaConnect as a standalone service or integrate it with other AWS services such as AWS Elemental MediaLive to prepare, process, and deliver your content. Whether you manage live 24x7 channels, stream live events, or need disaster recovery capabilities, MediaConnect gives you the tools to build your applications in the AWS Cloud.

132.1.2. Benefits

- **Reliably transport video:** MediaConnect adds a video-specific quality-of-service layer over standard IP transport, enabling uninterrupted, high-quality delivery. This makes your live video workflows more reliable, even over less dependable networks.
- **Securely share live video streams:** MediaConnect lets you secure your live video using industry-standard encryption, so only customers you authorize can access the content. This gives you control over the distribution of your video.
- **Easily manage high-value live broadcasts:** MediaConnect provides real-time visibility into more than 15 critical metrics of your video stream quality, automatically adjusting settings to optimize performance. Broadcast-standard alerts identify issues with transport streams, so you can maintain confidence that your video is delivered without disruption.
- **Connect remote production applications in the cloud:** MediaConnect supports remote production architectures by using AWS CDI flows to connect components such as production switchers, graphics engines, multiviewers, and playout servers running on different Amazon Elastic Compute Cloud (Amazon EC2) instances. MediaConnect

bridges these components together without sacrificing latency or video quality. MediaConnect also supports JPEG XS so you can connect your on-premises cameras and monitor walls to your Amazon EC2 remote production architecture.

- **Build for high availability:** MediaConnect provides resilience at the signal level through technologies such as automatic repeat request (ARQ) and at the flow level by employing SMPTE 2022-7 and hot/hot failover. MediaConnect also supports placement at the Availability Zone (AZ) level so it's easy to build live video workflows without any single points of failure.
- **Lower your costs by 30 percent or more:** MediaConnect uses your IP infrastructure or AWS Direct Connect to enable cost-effective, high-quality live video transport between on-premises and the AWS Cloud. Instead of long-term commitments to satellite transponders and fixed fiber networks, MediaConnect provides a pay-as-you-go transport alternative, saving you 30% or more compared to a typical satellite primary distribution use case with 70 destinations.

132.2. Backup/Restore and Disaster Recovery

This requirement is not applicable for this service. For additional information beyond what is described herein, please refer to http://aws.amazon.com/documentation/.

132.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

132.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/mediaconnect/index.html
- **Service quotas:** https://docs.aws.amazon.com/mediaconnect/latest/ug/quotas.html
- **Service FAQs:** https://aws.amazon.com/mediaconnect/faqs/

132.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/mediaconnect/index.html and the following links for comprehensive technical documentation regarding this service.

- User Guide: Describes the components and features that provides and how to use them.
- API Reference: Describes the AWS Elemental MediaConnect API operations in detail. Also provides sample requests, responses, and errors for the supported web services protocols.

133. AWS Elemental MediaConvert

133.1. Service Overview

AWS Elemental MediaConvert is a file-based video transcoding service with broadcast-grade features. It allows you to easily create video-on-demand (VOD) content for broadcast and multiscreen delivery at scale. The service combines advanced video and audio capabilities with a simple web services interface and pay-as-you-go pricing. With AWS Elemental MediaConvert, you can focus on delivering compelling media experiences without having to worry about the complexity of building and operating your own video processing infrastructure.

133.1.1. Features

- **Broadcast Capabilities for Video-on-Demand Content:** Configure a broadcast-grade video-on-demand experience: Built with technology from AWS Elemental that has been proven over many years by leading broadcast and internet video providers, AWS Elemental MediaConvert gives you a comprehensive set of features to create engaging viewing experiences, including graphic overlays, content protection, multi-language audio, closed captioning support, and professional broadcast formats.
- **Comprehensive Input and Output Support:** Create streams using standard video formats: AWS Elemental MediaConvert supports a broad range of video input and output formats, including those for broadcast as well as formats for delivery over the internet. The service supports the AVC, HEVC, AV1, Apple ProRes, and MPEG-2 compression standards, including support for advanced color sampling (10-bit 4:2:2). It supports a broad range of adaptive bitrate packaging formats including CMAF, Apple HLS, DASH ISO, and Microsoft Smooth Streaming. AWS Elemental MediaConvert also supports processing and conversion of 4K and 8K resolution sources, and high dynamic range (HDR) video content including Dolby Vision.
- **Automated Resource Provisioning:** Automate the work of creating and managing video infrastructure: AWS Elemental MediaConvert eliminates the burden of managing video processing infrastructure by automating the key aspects of workload provisioning and management. Simply use the AWS Management Console or the API to launch video processing workloads configured to your specifications; AWS Elemental MediaConvert handles resource provisioning and optimization, service orchestration, scaling, healing, resiliency failover, monitoring, and reporting. AWS Elemental MediaConvert supports a broad range of video input and output formats, including those for broadcast as well as formats for delivery over the internet.
- **Built-in Reliability:** Redundancy and automatic scalability for video-on-demand workflows: Each job you create with AWS Elemental MediaConvert runs on redundant infrastructure distributed across physically separated Availability Zones. The service monitors resources for health and automatically replaces any degraded components without disrupting your jobs. AWS Elemental MediaConvert scales elastically to handle peak workloads without reducing turnaround time or performance, automatically provisioning the right compute resources in step with demand for your video content. Incorporating features like error handling is straightforward, allowing you to monitor the entire media workflow and take actions, such as issuing notifications, in real-time.

133.1.2. Benefits

- **Broadcast-grade capabilities:** AWS Elemental MediaConvert lets you use a wide range of internet and professional media formats to produce high-quality video outputs that look great on any device. With support for ultra-high definition resolutions, high dynamic range video, graphic overlays, advanced audio features, content protection, and closed captioning, AWS Elemental MediaConvert offers a full set of tools to deliver high-quality viewing experiences.
- **Reliable and easy to manage:** AWS Elemental MediaConvert does not require any set up, management, or maintenance of underlying infrastructure. Simply submit jobs with the video processing settings you want and get started without spending time or resources managing transcoding infrastructure. Not only does AWS Elemental MediaConvert provision the required resources to process your jobs, but it also monitors them automatically, so you don’t need to worry about reliability.

- **Simple, predictable pricing:** AWS Elemental MediaConvert lets customers create high-quality, end-to-end video processing workflows in the cloud without upfront investment or capital expenditures for video processing infrastructure. You simply pay based on the duration of video that is processed and the features you use.

133.2. Backup/Restore and Disaster Recovery

This requirement is not applicable for this service. For additional information beyond what is described herein, please refer to http://aws.amazon.com/documentation/.

133.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

133.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/mediaconvert/
- **Service quotas:** https://docs.aws.amazon.com/general/latest/gr/mediaconvert.html
- **Service FAQs:** https://aws.amazon.com/mediaconvert/faqs/

133.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/mediaconvert/ and the following links for comprehensive technical documentation regarding this service.

- User Guide: Describes the components and features that AWS Elemental MediaConvert provides and how to use them.
- API Reference: Describes basic operations of AWS Elemental MediaConvert. Provides schema structure for job settings and detailed descriptions of encoding settings. Includes sample job requests.
- SPEKE Documentation: Describes the Secure Packager and Encoder Key Exchange (SPEKE) DRM Interface. DRM platform key providers and encryptors and packagers of media content use SPEKE for content encryption.

134. AWS Elemental MediaLive

134.1. Service Overview

AWS Elemental MediaLive is a broadcast-grade live video processing service. It lets you create high-quality video streams for delivery to broadcast televisions and internet-connected multiscreen devices, like connected TVs, tablets, smart phones, and set-top boxes. The service works by encoding your live video streams in real-time, taking a larger-sized live video source and compressing it into smaller versions for distribution to your viewers. With AWS Elemental MediaLive, you can easily set up streams for both live events and 24x7 channels with advanced broadcasting features, high availability, and pay-as-you-go pricing. AWS Elemental MediaLive lets you focus on creating compelling live video experiences for your viewers without the complexity of building and operating broadcast-grade video processing infrastructure.

134.1.1. Features

- **Comprehensive Video Standards Support:** Encode live streams using common video formats: AWS Elemental MediaLive supports a broad range of video industry standards

used to input, output, and archive live video. It includes support for the latest codecs - the compression standards used for video, like h.264/AVC and h.265/HEVC - and media communication protocols - the standards used to send video over the internet, like Real Time Protocol (RTP), HTTP Live Streaming (HLS) or Real-Time Streaming Protocol (RTMP). You can read about all the standards AWS Elemental MediaLive supports in our documentation.

- **Broadcast Capabilities for Live Video Streams:** Configure a broadcast-grade live viewing experience: AWS Elemental MediaLive uses technology from AWS Elemental that has been proven over many years by leading broadcast and internet video providers. AWS Elemental MediaLive supports broadcast features like ad markers, closed captions, multiple language audio tracks, audio descriptors, and FCC-mandated loudness correction. MediaLive also works natively with AWS Elemental MediaConnect, providing secure and reliable transport of video to use as inputs to live channels.
- **Statistical Multiplexing (Statmux):** Manage broadcast distribution in the cloud: Statistical Multiplexing (Statmux) lets you process and originate live content and share it with distribution partners for delivery over satellite, cable, or terrestrial networks. Launch broadcast distribution workloads quickly, optimize video quality and available network bandwidth, and meet resiliency goals with built-in high availability.
- **Automated Resource Provisioning:** Easy configuration of live video channels: AWS Elemental MediaLive manages the encoding resources needed to deliver high-availability live video streams, so you can focus on your content, not your encoding infrastructure. It automatically deploys encoding resources and manages scaling, healing, resiliency failover, monitoring, and reporting. With a few clicks in the AWS Elemental MediaLive management console, you can launch fully configured live video channels in minutes.
- **Automated High Availability:** Redundancy and automatic scalability for live video: Each live video channel you create with AWS Elemental MediaLive runs on redundant infrastructure distributed across physically separated Availability Zones. The service monitors encoding resources for health and automatically replaces any degraded components without disrupting your channels. Resources scale elastically with demand, assuring a consistent service for your viewers.
- **Flexible Workflows:** Flexible functionality for one-time events or around-the-clock live streams: AWS Elemental MediaLive lets you build flexible 24x7 linear workflows or event-based live streams. The service makes it simple to output a mix of live video streams at different resolutions and bitrates to meet the requirements of the devices in the hands of your viewers.
- **Use Independently or with other AWS Media Services:** Built-in integration with other AWS Media Services, or deploy on its own: As one of the AWS Media Services, you can choose to use AWS Elemental MediaLive as a standalone service or integrate it with AWS services for video transport, video-on-demand (VOD) processing, just-in-time packaging, ad personalization and monetization, or media-optimized storage. A typical integration would be to tie AWS Elemental MediaLive to our video transport, packaging and server-side ad insertion services for the core video workflow. Additionally, other AWS services, such as the Amazon CloudFront CDN, offer seamless interoperability with AWS Elemental Media Services.

134.1.2. Benefits

- **Simple deployment and management:** AWS Elemental MediaLive automates the provisioning and management of all the infrastructure used for video encoding, letting you deploy a simple live channel in minutes. The service transparently provisions resources and manages all the scaling, failover, monitoring, and reporting needed to power a live video stream. This lets you focus on your live content, not your encoding infrastructure. You can also use AWS Elemental Link, an on-premises device that makes it easy to connect your live video source to MediaLive.
- **Broadcast-grade capabilities:** AWS Elemental MediaLive makes it easy for anyone to produce broadcast-quality live streaming video. The service includes support for advanced capabilities such as statistical multiplexing, ad marker support, audio features including audio normalization and Dolby audio, and multiple caption standards. MediaLive works natively with AWS Elemental MediaConnect, providing secure and reliable transport of video to use as inputs to live channels.
- **Highly available:** AWS Elemental MediaLive provides built-in reliability and resiliency. The service transparently manages resources across multiple Availability Zones, and automatically monitors their health, so that any potential issues are detected and resolved without disrupting live channels. With AWS Elemental MediaLive, you can exceed the reliability of infrastructure typically used for broadcast workloads with a straightforward pay-as-you-go model based on the hours of content processed.
- **Increased efficiency and reduced cost:** With AWS Elemental MediaLive, you only pay for the service as you use it, with no upfront investment in encoding infrastructure and no operational overhead devoted to managing physical resources. Pricing for AWS Elemental MediaLive uses a straightforward model based on the hours of content processed and delivered.

134.2. Backup/Restore and Disaster Recovery

This requirement is not applicable for this service. For additional information beyond what is described herein, please refer to http://aws.amazon.com/documentation/.

134.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

134.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/medialive/index.html
- **Service quotas:** https://docs.aws.amazon.com/medialive/latest/ug/limits.html
- **Service FAQs:** https://aws.amazon.com/medialive/faqs/

134.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/medialive/index.html and the following links for comprehensive technical documentation regarding this service.

User Guide: Describes the components and features that AWS Elemental MediaLive provides and how to use them.

API Reference: Describes all the AWS Elemental MediaLive API operations. Also provides sample requests, responses, and errors for the supported web services protocols**.**

135. AWS Elemental MediaPackage

135.1. Service Overview

AWS Elemental MediaPackage reliably prepares and protects your video for delivery over the Internet. From a single video input, AWS Elemental MediaPackage creates video streams formatted to play on connected TVs, mobile phones, computers, tablets, and game consoles. It makes it easy to implement popular video features for viewers (start-over, pause, rewind, etc.), like those commonly found on DVRs. AWS Elemental MediaPackage can also protect your content using Digital Rights Management (DRM). AWS Elemental MediaPackage scales automatically in response to load, so your viewers will always get a great experience without you having to accurately predict in advance the capacity you'll need.

135.1.1. Features

- **Comprehensive Output Formats:** Deliver video content to the maximum number of viewers on multiple playback devices: AWS Elemental MediaPackage supports the standards and formats commonly used to stream video, including a range of MPEG-DASH implemtations, Smooth Streaming (MSS), and HTTP Live Streaming (HLS) with MPEG-2 Transport streams or Common Media Application Format (CMAF) fragmented MP4, to reach a maximum number of devices and viewers.
- **Flexible Video Content Protection:** Apply just-in-time content protection to secure your video assets: AWS Elemental MediaPackage lets you protect your streams by integrating with multiple Digital Rights Management (DRM) technologies. The right content protection is selected based on the capabilities of each playback device. Protection capabilities are standards-based, including support for Apple FairPlay, Widevine, and Microsoft PlayReady using AES-128 encryption.
- **High-Availability Architecture:** Built-in resiliency and automatic scalability for video workflows: AWS Elemental MediaPackage automatically scales based on the incoming stream requests you receive. MediaPackage has a built-in origin shield and cache for reliable performance without errors when accessed directly by one or more CDNs, and is designed for high reliability with distributed resources across multiple AWS Availability Zones. Integrated monitoring continuously tracks metrics (such as bandwidth, number of concurrent contacts or instance resources), and new instances are launched automatically to scale with increased workload as needed.
- **Use Independently or with AWS Media Services:** Use with other AWS Media Services, or deploy on its own: As one of the AWS Media Services, you can choose to use AWS Elemental MediaPackage as a standalone service or integrate it with AWS services for live video encoding, VOD processing, ad personalization and monetization, or media-optimized storage. Additionally, other AWS services, such as the Amazon CloudFront CDN, offer seamless interoperability with AWS Elemental MediaPackage.

135.1.2. Benefits

- **Reach a wide range of connected devices:** AWS Elemental MediaPackage makes it easy to package and distribute your content to a broad range of video playback devices, including web players, smart phones, game consoles, tablets, and connected TVs.
- **Advanced video experiences and content protection:** AWS Elemental MediaPackage lets you configure a DVR-like experience for viewers of your live stream.

It offers support for a range of digital rights management (DRM) providers, supports advanced audio features, and multi-language subtitle tracks.

- **Built-in scalability and reliability:** AWS Elemental MediaPackage scales automatically as the audience for your video grows. It automatically manages resources across multiple Availability Zones, and monitors their health, so that any potential issues are detected and resolved without disrupting your live video stream.
- **Easy integration with AWS cloud services:** AWS Elemental MediaPackage is built to work with Amazon CloudFront CDN for global distribution and with AWS Elemental MediaLive for live encoding to form a complete solution for live video processing and delivery. Integration with Amazon CloudWatch gives you real-time monitoring and notifications.

135.2. Backup/Restore and Disaster Recovery

This requirement is not applicable for this service. For additional information beyond what is described herein, please refer to http://aws.amazon.com/documentation/.

135.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

135.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/medialive/index.html
- **Service quotas:** https://docs.aws.amazon.com/medialive/latest/ug/limits.html
- **Service FAQs:** https://aws.amazon.com/mediapackage/faqs/

135.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/medialive/index.html and the following links for comprehensive technical documentation regarding this service.

- User Guide: Describes the components and features that AWS Elemental MediaLive provides and how to use them.
- API Reference: Describes all the AWS Elemental MediaLive API operations. Also provides sample requests, responses, and errors for the supported web services protocols.

136. AWS Elemental MediaStore

136.1. Service Overview

AWS Elemental MediaStore is an AWS storage service optimized for media. It gives you the performance, consistency, and low latency required to deliver live streaming video content. AWS Elemental MediaStore acts as the origin store in your video workflow. Its high performance capabilities meet the needs of the most demanding media delivery workloads, combined with long-term, cost-effective storage.

136.1.1. Features

- **Performance when you need it:** Low, predictable latencies for any sized audience: When you create or update a video file in AWS Elemental MediaStore, it is automatically held in a replicated cache for the first few minutes. This gives predictable latency and great performance, regardless of how many viewers watch your stream.
- **Consistent, low latency reads, writes, and updates:** Reduce buffering and end-to-end latency: Consistent read-after-write, and read-after-update performance is required for HTTP-based streaming video protocols that use manifests to tell a player what objects to download. If the content is not available immediately when requested, buffering or playback failures will occur, and if content manifests aren't current, devices will simply stop playback. AWS Elemental MediaStore is built for low-latency reads and writes, and high volumes of requests, which allows you to deliver consistent quality-of-service to viewers.
- **Access Control Management:** Control access to your video content: Integrated security management tools let you create Identity and Access Management (IAM) roles, improving security and restricting access to content to specific roles. As an example, you can restrict access to entire storage containers, or just particular content paths, exposing only the content required to CDNs and end viewers. Regardless of your workload, AWS security management tools give you the control, whether its managing content read access, or delegating write access to a particular set of users.
- **Use Independently or with AWS Media Services:** Use with other AWS Media Services, or deploy on its own: As one of the AWS Media Services, you can choose to use AWS Elemental MediaStore as a standalone service or integrate it with AWS services for live video encoding, VOD processing, ad personalization and monetization, or media-optimized storage. Additionally, other AWS services, such as the Amazon CloudFront CDN, offer seamless interoperability with AWS Elemental MediaPackage.

136.1.2. Benefits

- **High performance, optimized for video:** AWS Elemental MediaStore is optimized to deliver performance to meet the unique requirements of high-scale, high-quality media workloads (delivering low-latency reads and writes concurrently). This means you can deliver consistent quality-of-service to your viewers, lowering the risk of buffering video and reducing end-to-end latency.
- **Scale with your audience:** AWS Elemental MediaStore scales automatically with the volume of requests you receive. Increases in load don't negatively impact the quality of the viewing experience. This automatic scalability eliminates the expense and complexity of pre-provisioning capacity. If your audience size goes down, there is no need to de-scale - the service automatically sizes to varied capacity while you pay only for what you use.
- **Familiar management tools for access control:** AWS Elemental MediaStore integrates with AWS features for access control, using AWS Identity and Access Management (IAM) policies and roles, with support for resource policies, allowing you to specify granular access controls.

136.2. Backup/Restore and Disaster Recovery

This requirement is not applicable for this service. For additional information beyond what is described herein, please refer to http://aws.amazon.com/documentation/.

136.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

136.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/mediastore/index.html
- **Service quotas:** https://docs.aws.amazon.com/mediastore/latest/ug/quotas.html
- **Service FAQs:** https://aws.amazon.com/mediastore/faqs/

136.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/mediastore/index.html and the following links for comprehensive technical documentation regarding this service.

- User Guide: Describes the components and features that AWS Elemental MediaStore provides and how to use them.
- API Reference: Describes all the AWS Elemental MediaStore API operations in detail. Also provides sample requests, responses, and errors for the supported web services protocols.

137. AWS Elemental MediaTailor

137.1. Service Overview

AWS Elemental MediaTailor is a channel assembly and personalized ad insertion service for video providers to create linear OTT (internet delivered) channels using existing video content and monetize those channels, or other live streams and VOD content, with personalized advertising. With MediaTailor, virtual linear channels are created without the expense, complexity, and management of real-time live encoding, and live streams maintain a TV-like experience across multiscreen video applications. Advertisements are seamlessly stitched into the content and can be tailored to individual viewers, maximizing monetization opportunities for every ad break and mitigating ad blocking. The service works with any content delivery network to provide dependable, scalable channel assembly and ad personalization.

137.1.1. Features

- **Server-side Ad Insertion:** Insert advertising content upstream, prior to video delivery: AWS Elemental MediaTailor allows you to insert advertising content at the location of the start of a requested stream, prior to delivery. This eliminates the need to build and maintain unique configurations for every type of client device in order to insert personalized ads during video playback. Instead, the service generates and maintains a unique “manifest file” for each viewer, which is used to deliver ad placements that are personalized to the individual. Advertising content is seamlessly inserted into the primary content stream and can be played from the same source location to reduce the risk of buffering caused by high format and bitrate variability during video playback. This also reduces the effects of ad blocking software by making it difficult to distinguish ads from other content.
- **Media Manifest Manipulation:** Assemble linear channels using existing content: All content is delivered with a consolidated manifest that includes content and personalized

ads in a continuous stream to give viewers a seamless, TV-like viewing experience without buffering between program content and ad breaks. Support for HTTP Live Streaming (HLS) and Dynamic Adaptive Streaming over HTTP (DASH) standard manifests and playlists including Common Media Application Format (CMAF) means your live stream can be viewed on a broad range of devices and players.

- **Hybrid Measurement and Reporting:** Capture accurate, granular measurement of ad impressions: Accurate measurement and reporting of internet-delivered video advertising is required for advertisers and video providers to be compensated for every ad placement. AWS Elemental MediaTailor achieves the Interactive Advertising Bureau (IAB) level of playback metrics by implementing advertising measurement and reporting from the client side through playback observation APIs deployed on the viewing device. In addition, MediaTailor reports server-side ad metrics for legacy set-top boxes and other devices where changes to the viewing device is not possible, in order to comply with IAB specifications.
- **Automatic Scaling:** Scale resources up or down with viewership: AWS Elemental MediaTailor automatically scales with the number of concurrent viewers, maintaining consistent performance and quality of service for your internet-delivered video content as viewership goes up or down.
- **Your Choice of Video Workflow Components:** Use the vendors and solutions you choose: With AWS Elemental MediaTailor, you are not limited to specific vendors or solutions, including those components of the video workflow that operate directly with the service: the content delivery network (CDN), ad decision server, and origin server. The service works with most standard CDN or ad decision servers, and works with origin servers accessible over HTTP that can be configured using common video streaming protocol and proper ad markers.
- **Use Independently or with AWS Media Services:** Make use of built-in integration or deploy on its own: As one of the AWS Media Services, you can choose to use AWS Elemental MediaTailor as a standalone service or integrate it with AWS services for live video encoding, VOD processing, just-in-time packaging, or media-optimized storage. Additionally, other AWS services, such as the Amazon CloudFront CDN, offer seamless interoperability with AWS Elemental MediaTailor.

137.1.2. Benefits

- **Efficient linear channel creation and distribution:** Create live streams with low running cost by using existing encoded content. Provide viewers with better quality video using advanced codecs for higher quality with less bits to reduce distribution costs.
- **Engage audiences with relevant content:** Increase engagement with your audiences by delivering content and ads that are relevant with personalized ad insertions. Serve individual ads or groups of ads based on your specific business needs.
- **Seamless experience across different devices:** Give viewers a seamless, TV-like viewing experience without buffering between program content and ad breaks. Deliver consistent quality playback for your viewers with ad content transcoded to match the program content, so the quality and resolution of your streams remain smooth and predictable.
- **Increase effectiveness of monetization:** Get a faster and better return on investment (ROI) with assembled live channels that have a low running cost, and personalized ads that provide a higher return for each ad impression. Your viewers are also more likely to watch ads to the end if they are relevant to their interests.

- **Accurate ad view reporting:** Get accurate reporting from player and/or server side logs to capture ad view metrics without losing revenue from missing data or lost impressions. All content for streams, including ad content is delivered from one location, which reduces the effectiveness of ad-blockers.
- **Flexible standards-based solution:** Build exactly what you need using standards-based live streams and personalized ad-insertions that works with any content delivery networks (CDN) and a range of protocols and players.

137.2. Backup/Restore and Disaster Recovery

This requirement is not applicable for this service. For additional information beyond what is described herein, please refer to http://aws.amazon.com/documentation/.

137.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

137.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/mediatailor/index.html
- **Service quotas:** https://docs.aws.amazon.com/mediatailor/latest/ug/quotas.html
- **Service FAQs:** https://aws.amazon.com/mediatailor/faqs/

137.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/mediatailor/index.html and the following links for comprehensive technical documentation regarding this service.

- User Guide: Describes the components and features that AWS Elemental MediaTailor provides and how to use them.
- API Reference: Describes all the AWS Elemental MediaTailor API operations in detail. Also provides sample requests, responses, and errors for the supported web services protocols.

138. AWS Fargate

138.1. Service Overview

AWS Fargate is a serverless, pay-as-you-go compute engine that lets you focus on building applications without managing servers. AWS Fargate is compatible with both Amazon Elastic Container Service (ECS) and Amazon Elastic Kubernetes Service (EKS).

138.1.1. Features

- **Web apps, APIs, and microservices:** Build and deploy your applications, APIs, and microservices architectures with the speed and immutability of containers. Fargate removes the need to own, run, and manage the lifecycle of a compute infrastructure so that you can focus on what matters most: your applications.
- **Run and scale container workloads:** Use Fargate with Amazon ECS or Amazon EKS to easily run and scale your containerized data processing workloads. Fargate also

enables you to migrate and run your Amazon ECS Windows containers without refactoring or rearchitecting your legacy applications.

- **Support AI and ML training applications:** Create an AI and ML development environment that is flexible and portable. With Fargate, achieve the scalability you need to boost server capacity without over-provisioning—to train, test, and deploy your machine learning (ML) models.
- **Optimize Costs:** With AWS Fargate there are no upfront expenses, pay for only the resources used. Further optimize with Compute Savings Plans and Fargate Spot, then use Graviton2 powered Fargate for up to 40% price performance improvements.

138.1.2. Benefits

- **Application management, not infrastructure management:** Deploy and manage your applications, not infrastructure. Fargate removes the operational overhead of scaling, patching, securing, and managing servers.
- **Monitoring:** Monitor your applications via built-in integrations with AWS services like Amazon CloudWatch Container Insights. Gather metrics and logs with third-party tools.
- **Improved security:** Improve security through workload isolation by design. Amazon ECS tasks and Amazon EKS pods run in their own dedicated runtime environment.
- **Pay for what you use:** Only pay for what you use. Fargate scales the compute to closely match your specified resource requirements. With Fargate, there is no over-provisioning and paying for additional servers.

138.2. Backup/Restore and Disaster Recovery

This service allows for backup and recovery. The service can back up machine images. Users control this via a manual or scheduled API call. Users schedule and recover backups through a web interface.

138.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

138.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/ecs/index.html
- **Service quotas:** https://docs.aws.amazon.com/general/latest/gr/ecs-service.html
- **Service FAQs:** https://aws.amazon.com/fargate/faqs/?nc=sn&loc=4

138.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/ecs/index.html and the following links for comprehensive technical documentation regarding this service.

- User Guide for AWS Fargate: Describes key concepts of Amazon ECS on AWS Fargate and provides instructions for launching containers on the serverless infrastructure provided by Fargate.

139. AWS Fault Injection Simulator

139.1. Service Overview

AWS Fault Injection Simulator is a fully managed service for running fault injection experiments on AWS that makes it easier to improve an application's performance, observability, and resiliency. Fault injection experiments are used in chaos engineering, which is the practice of stressing an application in testing or production environments by creating disruptive events, such as sudden increase in CPU or memory consumption, observing how the system responds, and implementing improvements. Fault injection experiment helps teams create the real-world conditions needed to uncover the hidden bugs, monitoring blind spots, and performance bottlenecks that are difficult to find in distributed systems.

Fault Injection Simulator simplifies the process of setting up and running controlled fault injection experiments across a range of AWS services so teams can build confidence in their application behaviour. With Fault Injection Simulator, teams can quickly set up experiments using pre-built templates that generate the desired disruptions. Fault Injection Simulator provides the controls and guardrails that teams need to run experiments in production, such as automatically rolling back or stopping the experiment if specific conditions are met. With a few clicks in the console, teams can run complex scenarios with common distributed system failures happening in parallel or building sequentially over time, enabling them to create the real world conditions necessary to find hidden weaknesses.

139.1.1. Features

- **Simple setup:** AWS Fault Injection Simulator supports best practice chaos engineering parameters to make it easy to get started building and running fault injection experiments, without needing to install any agents. Sample experiments are available to use as a starting point. Fully managed fault injection actions are used to define actions such as stopping an instance, throttling an API, and failing over a database. Fault Injection Simulator supports Amazon CloudWatch so that you can use your existing metrics to monitor Fault Injection Simulator experiments.
- **Run real-world scenarios:** Simplistic scenarios can be insufficient to create the real-world conditions that cause failure so AWS Fault Injection Simulator supports gradually and simultaneously impairing performance of different types of resources, APIs, services, and geographic locations. Affected resources can be randomized, and custom fault types can be created using AWS Systems Manager to further increase complexity.
- **Fine grained safety controls:** When running experiments in live environments, there's a risk of unintended impact. To provide guardrails and keep your fault injection experiments under control, AWS Fault Injection Simulator allow you to target based on environments, application, and other dimensions using tags. For example, you could increase CPU utilization on 10% of your instances with the tag "environment":"prod". Fault Injection Simulator also has the option to set rules based on Amazon CloudWatch Alarms or other tools to stop an experiment. For example, an experiment can be set to stop before completion if a web page response time decreases below an acceptable level.
- **Integrated security model:** AWS Fault Injection Simulator is integrated with AWS Identity and Access Management (IAM) so that you can control which users and resources have permission to access and run Fault Injection Simulator experiments, and which resources and services can be affected.

- **Visibility throughout an experiment:** AWS Fault Injection Simulator provides visibility throughout every stage of an experiment via the console and APIs. As an experiment is running you can observe what actions have executed. After an experiment has completed you can see details on what actions were run, if stop conditions were triggered, how metrics compared to your expected steady state, and more. To support accurate operational metrics and effective troubleshooting, you can also identify what resources and APIs are affected by a Fault Injection Simulator experiment.
- **Console and programmatic access:** You can use AWS Fault Injection Simulator with the AWS Management Console, AWS CLI, and AWS SDKs. The Fault Injection Simulator APIs allow you to programmatically access the service so that you can integrate fault injection testing into your continuous integration and continuous delivery (or CI/CD) pipeline, and custom tooling.

139.1.2. Benefits

- **Improve application performance, resiliency, and observability:** AWS Fault Injection Simulator makes it easy for teams to run and observe their experiments from end-to-end, making it easier to find their monitoring blind spots, performance bottlenecks, or other “unknown” weaknesses missed by traditional software tests.
- **Validate how your application performs on AWS:** AWS Fault Injection Simulator supports creating disruptive events across a range of AWS services, such as Amazon EC2, Amazon EKS, Amazon ECS, and Amazon RDS. Teams can run GameDay scenarios or stress test their most critical applications on AWS at scale, helping them ensure their application will behave as expected.
- **Safeguard fault injection experiments:** AWS Fault Injection Simulator provides the fine-grained controls that teams need to define the specific conditions under which they want to stop an experiment or roll back to the pre-experiment state.
- **A fast and easy way to get started with fault injection experiments:** AWS Fault Injection Simulator provides prebuilt templates that enable teams to set up and run high quality experiments in minutes. Fault Injection Simulator structures the experiment process so that teams can quickly run fault injection experiments by following the step-by-step process in the console and selecting from a predefined list of actions.
- **Get superior insights by generating real-world failure conditions:** AWS Fault Injection Simulator is designed to run disruptive real-world scenarios on AWS that are very difficult for teams to accomplish on their own. With Fault Injection Simulator, teams can take actions such as gradually or simultaneously impairing the performance of different resources in a production environment at scale, enabling them to better validate their application behavior.

139.2. Backup/Restore and Disaster Recovery

Experiment Templates (json) can be created and exported; Experiment Results can be downloaded/stored.

139.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

139.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/fis/?id=docs_gateway
- **Service quotas:** https://docs.aws.amazon.com/fis/latest/userguide/fis-quotas.html
- **Service FAQs:** https://aws.amazon.com/fis/faqs/

139.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/fis/?id=docs_gateway and the following links for comprehensive technical documentation regarding this service.

- User Guide: Describes the key concepts and tasks for AWS FIS.
- API Reference: Describes the API operations for AWS FIS.
- AWS FIS section of the AWS CLI Reference: Describes the AWS CLI commands for AWS FIS.

140. AWS Firewall Manager

140.1. Service Overview

AWS Firewall Manager is a security management service which allows you to centrally configure and manage firewall rules across your accounts and applications in AWS Organizations. As new applications are created, Firewall Manager makes it easy to bring new applications and resources into compliance by enforcing a common set of security rules. Now you have a single service to build firewall rules, create security policies, and enforce them in a consistent, hierarchical manner across your entire infrastructure, from a central administrator account.

Using AWS Firewall Manager, you can easily roll out AWS WAF rules for your Application Load Balancers, API Gateways, and Amazon CloudFront distributions. You can create AWS Shield Advanced protections for your Application Load Balancers, ELB Classic Load Balancers, Elastic IP Addresses and CloudFront distributions. You can also configure new Amazon Virtual Private Cloud (VPC) security groups and audit any existing VPC security groups for your Amazon EC2, Application Load Balancer (ALB) and ENI resource types. You can deploy AWS Network Firewalls across accounts and VPCs in your organization. Finally, with AWS Firewall Manager, you can also associate your VPCs with Amazon Route 53 Resolvers DNS Firewall rules.

140.1.1. Features

- **Centrally deploy AWS Network Firewall across VPCs:** Using Firewall Manager, your security administrator can deploy firewall rules for AWS Network Firewall to control traffic leaving and entering your network across accounts and Amazon VPCs, from a single place. Any changes to the centrally configured set of rules are automatically deployed to your accounts and VPCs. This enables security administrators to consistently enforce centrally mandated firewall rules across an organization, even as new accounts and VPCs are created in your organization. At the same time, Firewall Manager also reports non-compliant issues including any VPCs and accounts that are missing Network Firewall protections.
- **Automatically deploy Amazon VPC security groups, AWS WAF rules, AWS Shield Advanced protections, AWS Network Firewall rules, and Amazon Route 53 Resolver DNS Firewall rules:** You can automatically enforce policies on AWS resources that currently exist or are created in the future, thereby ensuring compliance with firewall rules across the organization. AWS Firewall Manager gives customers the ability to apply AWS WAF rules, as well as Managed Rules for AWS WAF, on

Application Load Balancers, API Gateways and Amazon CloudFront accounts. You can apply AWS Shield Advanced protections on Application or Classic Load Balances, Elastic IP addresses or CloudFront distributions. Similarly, you can use AWS Firewall Manager to create a common primary security group across your EC2 instances in your VPC. With Firewall Manager you can automatically deploy Network Firewall endpoints and associated rules, for your VPCs. At the same time, Firewall Manager also lets you associate your VPCs with Route 53 Resolver DNS Firewall rules. You can choose to automatically enforce the rule on a newly created resource, or you can choose to be notified when the new resource is created.

- **Multi-account resource groups:** Within AWS Firewall Manager, you are able to group resources by Account, by Resource Type, and by Tag. Your security team can create policies for all resources within a particular group or across accounts in the organization.
- **Cross-account protection policies:** AWS Firewall Manager is integrated with AWS Organizations and will automatically fetch the list of accounts in your AWS organization to enable you to group resources across accounts. First, you build protection policies, which define a group of resources and associate the group with your policy. Then, you specify the scope of the policy to cover a specific set of AWS accounts, or all of your Organizations' accounts. Firewall Manager will deploy the protections only on the resources in the accounts based on the scope of the policy.
- **Hierarchical rule enforcement:** AWS Firewall Manager allows you to apply protection policies in a hierarchical manner, so you can delegate the creation of application-specific rules while retaining the ability to enforce certain rules centrally. Centrally applied rules are constantly monitored for any accidental removal or mishandling, thereby ensuring they are applied consistently.
- **Dashboard with compliance notifications:** AWS Firewall Manager provides a visual dashboard where you can quickly view which AWS resources are protected, identify non-compliant resources, and take appropriate action. You can also get notified when there are changes to your configurations through SNS notification streams.
- **Audit existing and future security groups in your VPCs:** With AWS Firewall Manager, you can create policies to set guardrails that define what security groups are allowed/disallowed across your VPCs. AWS Firewall Manager continuously monitors security groups to detect overly permissive rules, and helps improve firewall posture. You can get notifications of accounts and resources that are non-compliant or allow AWS Firewall Manager to take action directly through auto-remediation.

140.1.2. Benefits

- **Simplify management of firewall rules across your accounts:** AWS Firewall Manager is integrated with AWS Organizations so you can enable AWS WAF rules, AWS Shield Advanced protections, security groups, AWS Network Firewall rules, and Amazon Route 53 Resolver DNS Firewall rules for your Amazon VPC across multiple AWS accounts and resources from a single place. You can group rules, build policies, and centrally apply those policies across your entire infrastructure. For example, you can delegate the creation of application-specific rules within an account while retaining the ability to enforce global security policies across accounts.
- **Ensure compliance of existing and new applications:** AWS Firewall Manager automatically enforces mandatory security policies that you define across existing and newly created resources. The service discovers new resources as they are created across accounts. For example, if you are required to meet US Department of Treasury's

Office of Foreign Assets Control (OFAC) regulations, you can use Firewall Manager to deploy an AWS WAF rule to block traffic from embargoed countries across your Application Load Balancer, API Gateway, and Amazon CloudFront accounts. As new resources are created, they will automatically be brought under the policy scope.

- **Easily deploy managed rules across accounts:** AWS Firewall Manager integrates with Managed Rules for AWS WAF, which gives you an easy way to deploy pre-configured WAF rules on your applications. You can choose a Managed Rule from an AWS Marketplace Seller and deploy it consistently across your Application Load Balancer, API Gateway, and Amazon CloudFront infrastructure with just a few clicks in the console. For example, you can easily protect your entire organization from zero-day vulnerabilities by subscribing to a Managed Rule for WAF from the AWS Marketplace that provides CVE patch updates. For Advanced Shield protections, you can use AWS Firewall Manager to automatically protect against various types of DDoS attacks such as UDP reflection attacks, SYN flood, DNS query flood and HTTP flood attacks across accounts.
- **Centrally deploy protections for your VPCs:** With Firewall Manager, your security administrator can deploy baseline set of VPC security group rules for EC2 instances, Application Load Balancers (ALBs) and Elastic Network Interfaces (ENIs) in your Amazon VPCs. At the same time, you can also audit any existing security groups in your VPCs for over permissive rules and remediate them from a single place. You can leverage Firewall Manager to deploy rules for AWS Network Firewalls across your VPCs in your organization, to control traffic leaving and entering your network. At the same time, with Firewall Manager, you can also associate your VPCs with Route 53 Resolver DNS Firewall rules to block DNS queries made for known malicious domains and to allow queries for trusted domains.

140.2. Backup/Restore and Disaster Recovery

This requirement is not applicable for this service. For additional information beyond what is described herein, please refer to http://aws.amazon.com/documentation/.

140.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

140.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/firewall-manager/
- **Service quotas:** https://docs.aws.amazon.com/waf/latest/developerguide/fms-limits.html
- **Service FAQs:** https://aws.amazon.com/firewall-manager/faqs/

140.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/firewall-manager/ and the following links for comprehensive technical documentation regarding this service.

- Developer Guide: Describes how to get started with AWS Firewall Manager. Explains key concepts, and provides step-by-step instructions that show you how to use the features.

- API Reference: Describes all the API operations for AWS Firewall Manager in detail.

141. AWS Global Accelerator

141.1. Service Overview

AWS Global Accelerator is a networking service that improves the performance of your users' traffic by up to 60% using Amazon Web Services' global network infrastructure. When the internet is congested, AWS Global Accelerator optimizes the path to your application to keep packet loss, jitter, and latency consistently low.

With Global Accelerator, you are provided two global static public IPs that act as a fixed entry point to your application, improving availability. On the back end, add or remove your AWS application endpoints, such as Application Load Balancers, Network Load Balancers, EC2 Instances, and Elastic IPs without making user-facing changes. Global Accelerator automatically re-routes your traffic to your nearest healthy available endpoint to mitigate endpoint failure.

Set up your accelerator on the AWS Management Console in minutes with step-by-step documentation or with one click in the Elastic Load Balancing Console. Learn more by following the self-service workshop and test performance benefits from your location with the AWS Global Accelerator speed comparison tool.

141.1.1. Features

- **Static anycast IP addresses:** AWS Global Accelerator provides you with static IP addresses that serve as a fixed entry point to your applications hosted in one or more AWS Regions. These IP addresses are anycast from AWS edge locations, so they're announced from multiple AWS edge locations at the same time. This enables traffic to ingress onto the AWS global network as close to your users as possible. You can associate these addresses to regional AWS resources or endpoints, such as Application Load Balancers, Network Load Balancers, EC2 instances, and Elastic IP addresses. AWS Global Accelerator's IP addresses serve as the frontend interface of your applications. By using these static IP addresses, you don't need to make any client-facing changes or update DNS records as you modify or replace endpoints. The addresses are assigned to your accelerator for as long as it exists, even if you disable the accelerator and it no longer accepts or routes traffic.
- **Fault tolerance using network zones:** AWS Global Accelerator has a fault-isolating design that increases the availability of your applications. When you create an accelerator, AWS Global Accelerator allocates two static IPv4 addresses for you that are serviced by independent network zones. Similar to Availability Zones, these network zones are isolated units with their own set of physical infrastructure and service IP addresses from a unique IP subnet. If one IP address from a network zone becomes unavailable, due to network disruptions or IP address blocking by certain client networks, your client applications can retry using the healthy static IP address from the other isolated network zone.
- **Global performance-based routing:** AWS Global Accelerator uses the vast, congestion-free AWS global network to route TCP and UDP traffic to a healthy application endpoint in the closest AWS Region to the user. If there's an application failure, AWS Global Accelerator provides instant failover to the next best endpoint. You can test the performance benefits from your location with the speed comparison tool.

- **TCP Termination at the Edge:** Normally, a TCP connection is established via a three-way handshake (i.e., three messages) between the client on the internet and the application endpoint in the AWS Region. With TCP termination at the Edge, AWS Global Accelerator reduces the initial setup time by establishing a TCP connection between the client and the AWS edge location closest to the client. Almost concurrently, a second TCP connection is made between the edge location and the application endpoint in the AWS Region. Because of this process, the client gets a faster response from the Global Accelerator edge location, and the upstream connection from the edge location to the application endpoint in the Region is optimized to run over the AWS global network.
- **Bring your own IP (BYOIP):** AWS Global Accelerator allows you to bring your own IP addresses (BYOIP) and use them as a fixed entry point to your application endpoints. You can bring up to two /24 IPv4 address ranges and choose which /32 IP addresses to use when you create your accelerator. If you only bring one /24 IP address range, when you create an accelerator, Global Accelerator will assign a second /32 IP address from the Amazon IP address pool as the other static IP for your accelerator.
- **Fine-grained traffic control:** AWS Global Accelerator gives you the option to dial up or dial down traffic to a specific AWS Region by using traffic dials. For each Region (or endpoint group), you can set a traffic dial to control the percentage of traffic that is directed to that Region. The percentage is applied only to traffic that is already directed to that Region, based on proximity and health of the endpoints. The traffic dial lets you easily do performance testing or blue/green deployment testing for new releases across different AWS Regions, for example. If an endpoint fails, AWS Global Accelerator assigns your user traffic to the other endpoints, to maintain high availability. By default, traffic dials are set to 100% across all endpoint groups so that AWS Global Accelerator can select the best endpoint for your applications.
- **Continuous availability monitoring:** AWS Global Accelerator continuously monitors the health of your application endpoints by using TCP, HTTP, and HTTPS health checks. It instantly reacts to changes in the health or configuration of your endpoints, and redirects user traffic to healthy endpoints that deliver the best performance and availability to your users.
- **Client affinity:** AWS Global Accelerator enables you to build applications that require maintaining state. For stateful applications where you need to consistently route users to the same endpoint, you can choose to direct all requests from a user to the same endpoint, regardless of the port and protocol.
- **Distributed denial of service (DDoS) resiliency at the edge:** By default, AWS Global Accelerator is protected by AWS Shield Standard, which minimizes application downtime and latency from denial of service attacks by using always-on network flow monitoring and automated in-line mitigation. You can also enable AWS Shield Advanced for automated resource-specific enhanced detection and mitigation, as well as 24x7 access to the AWS DDoS Response Team (DRT) for manual mitigations of sophisticated DDoS attacks. AWS Shield Advanced also provides complete visibility into DDoS attacks and DDoS cost protection for scaling. This ensures scalable, reliable, and cost-efficient DDoS protection at the edge for your applications.

141.1.2. Benefits

- **Accelerate latency-sensitive applications:** Your network latency is driven by the number of networks your user data needs to hop and the bandwidth available along the path to your AWS application endpoints. These network variables create opportunities for internet congestion to delay connections and lose data. AWS Global Accelerator

combines advanced networking features with the dedicated AWS Global Network to improve your application network performance by up to 60%. TCP connections are terminated at the AWS Edge location closest to your users, instead of at your endpoint, accelerating data transfers globally. Once on the AWS network, automated routing directs your user traffic to the most performant AWS endpoints in Regions and/or Availability Zones. For UDP workloads, the AWS network provides the global capacity needed to avoid packet loss and jitter during traffic spikes.

- **Improve resiliency and availability:** You need to build your architecture with resiliency and availability in mind. This can mean running your application in a single AWS Region across multiple Availability Zones or across multiple AWS Regions. Wherever you route your traffic on the AWS network, with Global Accelerator, failover between application endpoints happens automatically and within seconds. If Global Accelerator detects a failure of your application endpoint it instantly triggers traffic re-routing to the next available, closest endpoint in another AZ or AWS Region. Your users are redirected without needing new IP addresses or updates to their DNS cache.
- **Simplified global traffic management:** As your application grows, the number of endpoints and IP addresses that you need to manage increases and becomes burdensome. As you update your application, to add or remove endpoints, you risk lowering availability of your application due to firewalls, hardcoded devices and allow-lists not having the latest information. AWS Global Accelerator simplifies global traffic management by providing 2 static anycast IP addresses that only need to be configured by users once. Behind these IP address you can add or remove AWS origins, opening up uses such as endpoint failover, scaling, or testing without any user-side changes. For A/B testing or blue green deployment, use traffic dials or endpoint weights to customize how much traffic is going to each endpoint. You can also bring your own IP addresses (BYOIP) to AWS Global Accelerator or use static IP addresses from the Amazon IP address pool.
- **Protect your applications:** Exposing your application built on AWS, through services such as Application Load Balancers or EC2 instances, to public internet traffic creates an opportunity for malicious attack. AWS Global accelerator decreases the risk of attack by masking your application behind two static entry points. These entry points are protected by default from Distributed Denial of Service (DDoS) attacks with AWS Shield. AWS Global Accelerator creates a peering connection with your Amazon Virtual Private Cloud using private IP addresses, keeping connections to your internal Application Load Balancer or private EC2 instance off the public internet.

141.2. Backup/Restore and Disaster Recovery

This requirement is not applicable for this service. For additional information beyond what is described herein, please refer to http://aws.amazon.com/documentation/.

141.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

141.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/global-accelerator/index.html

- **Service quotas:** https://docs.aws.amazon.com/global-accelerator/latest/dg/limits-global-accelerator.html
- **Service FAQs:** https://aws.amazon.com/global-accelerator/faqs/

141.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/global-accelerator/index.html and the following links for comprehensive technical documentation regarding this service.

- Developer Guide: Provides an overview of AWS Global Accelerator, detailed feature descriptions, procedures for using the console, and an explanation of the API.
- API Reference: Describes all the API operations for AWS Global Accelerator in detail. Also provides sample requests, responses, and errors for the supported web services protocols.

142. AWS Glue

142.1. Service Overview

AWS Glue is a serverless data integration service that makes it easy to discover, prepare, and combine data for analytics, machine learning, and application development. AWS Glue provides all the capabilities needed for data integration so that you can start analysing your data and putting it to use in minutes instead of months.

Data integration is the process of preparing and combining data for analytics, machine learning, and application development. It involves multiple tasks, such as discovering and extracting data from various sources; enriching, cleaning, normalizing, and combining data; and loading and organizing data in databases, data warehouses, and data lakes. These tasks are often handled by different types of users that each use different products.

AWS Glue provides both visual and code-based interfaces to make data integration easier. Users can easily find and access data using the AWS Glue Data Catalog. Data engineers and ETL (extract, transform, and load) developers can visually create, run, and monitor ETL workflows with a few clicks in AWS Glue Studio. Data analysts and data scientists can use AWS Glue DataBrew to visually enrich, clean, and normalize data without writing code. With AWS Glue Elastic Views, application developers can use familiar Structured Query Language (SQL) to combine and replicate data across different data stores.

142.1.1. Features

- **Discover and search across all your AWS data sets:** The AWS Glue Data Catalog is your persistent metadata store for all your data assets, regardless of where they are located. The Data Catalog contains table definitions, job definitions, schemas, and other control information to help you manage your AWS Glue environment. It automatically computes statistics and registers partitions to make queries against your data efficient and cost-effective. It also maintains a comprehensive schema version history so you can understand how your data has changed over time.
- **Automatic schema discovery:** AWS Glue crawlers connect to your source or target data store, progresses through a prioritized list of classifiers to determine the schema for your data, and then creates metadata in your AWS Glue Data Catalog. The metadata is stored in tables in your data catalog and used in the authoring process of your ETL jobs. You can run crawlers on a schedule, on-demand, or trigger them based on an event to ensure that your metadata is up-to-date.

- **Manage and enforce schemas for data streams:** AWS Glue Schema Registry, a serverless feature of AWS Glue, enables you to validate and control the evolution of streaming data using registered Apache Avro schemas, at no additional charge. Through Apache-licensed serializers and deserializers, the Schema Registry integrates with Java applications developed for Apache Kafka, Amazon Managed Streaming for Apache Kafka (MSK), Amazon Kinesis Data Streams, Apache Flink, Amazon Kinesis Data Analytics for Apache Flink, and AWS Lambda. When data streaming applications are integrated with the Schema Registry, you can improve data quality and safeguard against unexpected changes using compatibility checks that govern schema evolution. Additionally, you can create or update AWS Glue tables and partitions using schemas stored within the registry.
- **Visually transform data with a drag-and-drop interface:** AWS Glue Studio allows you to author highly scalable ETL jobs for distributed processing without becoming an Apache Spark expert. Define your ETL process in the drag-and-drop job editor and AWS Glue automatically generates the code to extract, transform, and load your data. The code is generated in Scala or Python and written for Apache Spark.
- **Build complex ETL pipelines with simple job scheduling:** AWS Glue jobs can be invoked on a schedule, on-demand, or based on an event. You can start multiple jobs in parallel or specify dependencies across jobs to build complex ETL pipelines. AWS Glue will handle all inter-job dependencies, filter bad data, and retry jobs if they fail. All logs and notifications are pushed to Amazon CloudWatch so you can monitor and get alerts from a central service.
- **Clean and transform streaming data in-flight:** Serverless streaming ETL jobs in AWS Glue continuously consume data from streaming sources including Amazon Kinesis and Amazon MSK, clean and transform it in-flight, and make it available for analysis in seconds in your target data store. Use this feature to process event data like IoT event streams, clickstreams, and network logs. AWS Glue streaming ETL jobs can enrich and aggregate data, join batch and streaming sources, and run a variety of complex analytics and machine learning operations.
- **Combine and replicate data across multiple data stores using SQL:** AWS Glue Elastic Views enables you to create views over data stored in multiple types of AWS data stores, and materialize the views in a target data store of your choice. You can use AWS Glue Elastic Views to create materialized views by writing queries in PartiQL. PartiQL is an open source SQL-compatible query language that you can use to query and manipulate data, regardless of whether the data has a tabular or a flexible, document-like structure. You can interactively write PartiQL queries using the query editor in the AWS Management Console or issue queries through the API or CLI. AWS Glue Elastic Views supports Amazon DynamoDB as a source (with support for Amazon Aurora and Amazon RDS to follow), and Amazon Redshift, Amazon OpenSearch Service (successor to Amazon Elasticsearch Service), and Amazon S3 as targets (with support for Amazon Aurora, Amazon RDS, and Amazon DynamoDB to follow). You can speed up development time by sharing your materialized views with other users for use in their applications. AWS Glue Elastic Views monitors for changes to data in your source data stores continuously, and provides updates to your target data stores automatically. Learn more about AWS Glue Elastic Views.
- **Deduplicate and cleanse data with built-in machine learning:** AWS Glue helps clean and prepare your data for analysis without becoming a machine learning expert. Its FindMatches feature deduplicates and finds records that are imperfect matches of each

other. For example, use FindMatches to find duplicate records in your database of restaurants, such as when one record lists “Joe's Pizza” at “121 Main St.” and another shows a “Joseph's Pizzeria” at “121 Main”. FindMatches will just ask you to label sets of records as either “matching” or “not matching.” The system will then learn your criteria for calling a pair of records a “match” and will build an ETL job that you can use to find duplicate records within a database or matching records across two databases.

- **Edit, debug, and test ETL code with developer endpoints:** If you choose to interactively develop your ETL code, AWS Glue provides development endpoints for you to edit, debug, and test the code it generates for you. You can use your favorite IDE or notebook. You can write custom readers, writers, or transformations and import them into your AWS Glue ETL jobs as custom libraries. You can also use and share code with other developers in our GitHub repository.
- **Normalize data without code using a visual interface:** AWS Glue DataBrew provides an interactive, point-and-click visual interface for users like data analysts and data scientists to clean and normalize data without writing code. You can easily visualize, clean, and normalize data directly from your data lake, data warehouses, and databases, including Amazon S3, Amazon Redshift, Amazon Aurora, and Amazon RDS. You can choose from over 250 built-in transformations to combine, pivot, and transpose the data, and automate data preparation tasks by applying saved transformations directly to the new incoming data.

142.1.2. Benefits

- **Faster data integration:** Different groups across your organization can use AWS Glue to work together on data integration tasks, including extraction, cleaning, normalization, combining, loading, and running scalable ETL workflows. This way, you reduce the time it takes to analyze your data and put it to use from months to minutes.
- **Automate your data integration at scale:** AWS Glue automates much of the effort required for data integration. AWS Glue crawls your data sources, identifies data formats, and suggests schemas to store your data. It automatically generates the code to run your data transformations and loading processes. You can use AWS Glue to easily run and manage thousands of ETL jobs or to combine and replicate data across multiple data stores using SQL.
- **No servers to manage:** AWS Glue runs in a serverless environment. There is no infrastructure to manage, and AWS Glue provisions, configures, and scales the resources required to run your data integration jobs. You pay only for the resources your jobs use while running.

142.2. Backup/Restore and Disaster Recovery

This requirement is not applicable for this service. For additional information beyond what is described herein, please refer to http://aws.amazon.com/documentation/.

142.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

142.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/glue/

- **Service quotas:** https://docs.aws.amazon.com/general/latest/gr/glue.html
- **Service FAQs:** https://aws.amazon.com/glue/faqs/

142.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/glue/ and the following links for comprehensive technical documentation regarding this service.

- Developer Guide: Provides a conceptual overview of AWS Glue, detailed instructions for using the various features, and a complete API reference for developers
- AWS Glue Studio User Guide: Describes how to use the AWS Glue Studio console and the visual job editor interface to build and monitor ETL jobs.
- AWS Glue section of the AWS CLI Reference: Describes the AWS CLI commands that you can use with AWS Glue.
- AWS Glue DataBrew Developer Guide: Describes how to prepare data visually with ready-made data transformations for analytics and machine learning. Also provides an API reference complete with instructions, syntax, and examples.

143. AWS Identity and Access Management (IAM)

143.1. Service Overview

AWS Identity and Access Management (IAM) provides fine-grained access control across all of AWS. With IAM, you can specify who can access which services and resources, and under which conditions. With IAM policies, you manage permissions to your workforce and systems to ensure least-privilege permissions.

143.1.1. Features

- **Fine-grained access control:** Permissions let you specify and control access to AWS services and resources. To grant permissions to IAM roles, you can attach a policy that specifies the type of access, the actions that can be performed, and the resources on which the actions can be performed. Using IAM policies, you grant access to specific AWS service APIs and resources. You also can define specific conditions in which access is granted, such as granting access to identities from a specific AWS organization or access through a specific AWS service.
- **Delegate access by using IAM roles:** With IAM roles you delegate access to users or AWS services to operate within your AWS account. Users from your identity provider or AWS services can assume a role to obtain temporary security credentials that can be used to make an AWS request in the account of the IAM role. Consequently, IAM roles provide a way to rely on short-term credentials for users, workloads, and AWS services that need to perform actions in your AWS accounts.
- **IAM Access Analyzer:** Achieving least privilege is a continuous cycle to grant the right fine-grained permissions as your requirements evolve. IAM Access Analyzer helps you streamline permissions management as you set, verify, and refine permissions.
- **Permissions guardrails:** With AWS Organizations, you can use service control policies (SCPs) to establish permissions guardrails that all IAM users and roles in an organization's accounts adhere to. Whether you're just getting started with SCPs or have existing SCPs, you can use IAM access advisor to help you restrict permissions confidently across your AWS organization.

- **Attribute-based access control:** Attribute-based access control (ABAC) is an authorization strategy you can use to create fine-grained permissions based on user attributes, such as department, job role, and team name. Using ABAC, you can reduce the number of distinct permissions that you need for creating fine-grained controls in your AWS account.

143.1.2. Benefits

- **Apply fine-grained access control:** Grant access to specific AWS service APIs and resources by using IAM policies. You also can define specific conditions in which access is granted, such as granting access to identities from a specific AWS organization or access through a specific AWS service.
- **Establish permissions guardrails and data perimeters across your AWS organization:** With AWS Organizations, you can use service control policies (SCPs) to establish permissions guardrails that all IAM users and roles in an organization's accounts adhere to. Whether you're just getting started with SCPs or have existing SCPs, you can use IAM access advisor to help you restrict permissions confidently.
- **Achieve least-privilege permissions with IAM Access Analyzer:** Achieving least privilege is a continuous cycle to grant the right fine-grained permissions as your requirements evolve. IAM Access Analyzer helps you streamline permissions management as you set, verify, and refine permissions.
- **Automatically scale fine-grained permissions with ABAC:** Attribute-based access control (ABAC) is an authorization strategy for creating fine-grained permissions based on user attributes, such as department, job role, and team name. With ABAC, you can reduce the number of distinct permissions you need for creating fine-grained controls in your AWS account.

143.2. Backup/Restore and Disaster Recovery

This requirement is not applicable for this service. For additional information beyond what is described herein, please refer to http://aws.amazon.com/documentation/.

143.3. Pricing Overview

IAM is a feature of your AWS account and is offered at no additional charge.

143.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/iam/index.html
- **Service quotas:** https://docs.aws.amazon.com/IAM/latest/UserGuide/reference_iam-quotas.html
- **Service FAQs:** https://aws.amazon.com/iam/faqs/

143.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/iam/index.html and the following links for comprehensive technical documentation regarding this service.

- IAM User Guide: Introduces you to AWS Identity and Access Management, helps you set up users and groups, and shows you how to protect your resources with access

control policies. Also shows how to connect to other identity services to grant external users access to your AWS resources.

- Service Authorization Reference: Provides a list of the actions, resources, and condition keys supported by each AWS service that can be used in an IAM policy.

144. AWS IoT Analytics

144.1. Service Overview

AWS IoT Analytics is a fully-managed service that makes it easy to run and operationalize sophisticated analytics on massive volumes of IoT data without having to worry about the cost and complexity typically required to build an IoT analytics platform. It is the easiest way to run analytics on IoT data and get insights to make better and more accurate decisions for IoT applications and machine learning use cases.

AWS IoT Analytics automates each of the difficult steps that are required to analyze data from IoT devices. AWS IoT Analytics filters, transforms, and enriches IoT data before storing it in a time-series data store for analysis. You can setup the service to collect only the data you need from your devices, apply mathematical transforms to process the data, and enrich the data with device-specific metadata such as device type and location before storing the processed data. Then, you can analyze your data by running ad hoc or scheduled queries using the built-in SQL query engine, or perform more complex analytics and machine learning inference. AWS IoT Analytics makes it easy to get started with machine learning by including pre-built models for common IoT use cases.

You can also use your own custom analysis, packaged in a container, to execute on AWS IoT Analytics. AWS IoT Analytics automates the execution of your custom analyses created in Jupyter Notebook or your own tools (such as Matlab, Octave, etc.) to be executed on your schedule. AWS IoT Analytics is a fully managed service that operationalizes analyses and scales automatically to support up to petabytes of IoT data. With AWS IoT Analytics, you can analyse data from millions of devices and build fast, responsive IoT applications without managing hardware or infrastructure.

144.1.1. Features

- **Ingest data from any source including AWS IoT Core**: Ingest data directly from AWS IoT Core to AWS IoT Analytics. Or, use the BatchPutMessage API to send your data to AWS IoT Analytics from Amazon S3, Amazon Kinesis or any other source. With AWS IoT Analytics' full integration with AWS IoT Core and the API, it is easy to receive messages from connected devices as they stream in.
- **Collect only the data you want to store and analyse**: You use the AWS IoT Analytics console to configure AWS IoT Analytics to receive messages from devices through MQTT topic filters in various formats and frequencies. AWS IoT Analytics validates that the data is within specific parameters you define and creates channels. Then the service routes the channels to appropriate pipelines for message processing, transformation, and enrichment.
- **Cleanse and filter**: AWS IoT Analytics let you define AWS Lambda functions that can be triggers on when AWS IoT Analytics detects missing data, so you can run code to estimate and fill gaps. You can also define max/min filters and percentile thresholds to remove outliers in your data.

- **Transform**: AWS IoT Analytics can transform messages using mathematical or conditional logic you define, so you can perform common calculations like Celsius into Fahrenheit conversion.
- **Enrich**: AWS IoT Analytics can enrich data with external data sources such as a weather forecast information, and then route the data to the AWS IoT Analytics data store.
- **Reprocess**: AWS IoT Analytics can reprocess raw data from the Channel connected to the Pipeline. Reprocessing your raw data gives you the flexibility to create a new pipeline or revisit an older pipeline so you can capture new and historical data, make changes to your pipeline, or simply process your data in a different way. This capability is often needed to gain deeper insights or test hypothesis. Simply connect the Pipeline to the appropriate Channel to reprocess.
- **Time-Series Data Store**: AWS IoT Analytics stores the device data in an IoT optimized time-series data store for analysis. You can manage access permissions, implement data retention policies and export your data to external access points.
- **Store Processed and Raw Data**: AWS IoT Analytics stores the processed data and also automatically stores the raw ingested data so you can process it at a later time.
- **Run Ad hoc or Scheduled SQL Queries**: AWS IoT Analytics provides a built-in SQL query engine so you can run ad hoc or scheduled queries and get results quickly. For example, you may want to run a quick query to find out how many monthly active users there are for each device in your fleet.
- **Time-Series Analysis**: AWS IoT Analytics supports time-series analysis so you can analyse the performance of devices over time and understand how and where they are being used, continuously monitor device data to predict maintenance issues, and monitor sensors to predict and react to environmental conditions.
- **Hosted Notebooks for Sophisticated Analytics and Machine Learning**: AWS IoT Analytics includes support for hosted Jupyter Notebooks for statistical analysis and machine learning. The service includes a set of pre-built notebook templates that contain AWS-authored machine learning models and visualizations to help you get started with IoT use cases related to device failure profiling, forecasting events such as low usage that might signal the customer will abandon the product, or segmenting devices by customer usage levels (for example heavy users, weekend users) or device health.
- **Bring Your Custom Container**: WS IoT Analytics will import your custom authored code containers, built in AWS IoT Analytics or a third party, such as Matlab, or Octave, etc., giving you more time to focus on what sets you apart from your competition. No need to recreate your existing analyses created in third party tools. Simply import your analyses container on AWS IoT Analytics and execute it as needed. If you are using Jupyter Notebooks, simply create an executable container image of your Jupyter Notebook code with just a click of a button and visualize your container analysis on the AWS IoT Analytics console.
- **Automate Container Execution**: AWS IoT Analytics lets you automate the execution of containers hosting custom authored analytical code or Jupyter Notebooks to perform continuous analysis. You can schedule execution of your custom analysis on the recurring schedule that best meets the need of your business.

- **Incremental Data Capture with Customizable Time Windows**: AWS IoT Analytics enables users to perform analysis on new incremental data captured since the last analysis. You can improve analysis efficiency and lower costs by precisely scanning just your new data. No matter when you ran your last analysis, customizable time windows will capture the new data for you since your last analysis.
- **QuickSight Integration**: AWS IoT Analytics provides a connector to Amazon QuickSight so you can visualize your data sets in a QuickSight dashboard. You can also visualize the results or your ad-hoc analysis in the embedded Jupyter Notebooks within the AWS IoT Analytics' console.

144.1.2. Benefits

- **Operationalize your analytical workflows:** You supply the analysis, AWS IoT Analytics automates the execution of your analysis when and where you need it. AWS IoT Analytics will import your custom authored code containers, built in external tools such as Matlab, Octave, etc, and execute them on your schedule to generate operational insights, giving you more time to focus on what you do best.
- **Easily run queries on IoT data:** With AWS IoT Analytics, you can run simple, ad-hoc queries using the built-in SQL query engine. Using standard SQL queries to extract data from the data store, you can calculate the average distance travelled of a fleet of vehicles or number of doors locked in a smart building, for example. Also, AWS IoT Analytics provides a series of non-overlapping, contiguous time windows to perform analysis on new, incremental data. You can improve analysis efficiency and lower costs by scanning only the data you need.
- **Data storage optimized for IoT:** AWS IoT Analytics stores the processed device data in a time-series data store that is optimized to deliver fast response times on IoT queries. The raw data is also automatically stored for later processing or reprocessing for another use case.
- **Prepares your IoT data for analysis:** AWS IoT Analytics includes data preparation techniques that make it easy to prepare and process your data for analysis. AWS IoT Analytics also supports time-series analyses so you can analyse the performance of devices over time and understand how and where they are being used, continuously monitor device data to predict maintenance issues, and monitor sensors to predict and react to environmental conditions. AWS IoT Analytics is integrated with AWS IoT Core to easily ingest device data directly from connected devices. It cleans false readings, fills gaps in the data, and performs mathematical transformations of message data. As the data is ingested, AWS IoT Analytics can process it using conditional statements, filter data to collect just the data you want to analyse, and enrich it with information from the AWS IoT registry. You can also use AWS Lambda functions to enrich your device data from external sources like the Weather Service, HERE Maps, Salesforce, or Amazon DynamoDB.
- **Tools for machine learning:** AWS IoT Analytics makes it easy to apply machine learning to your IoT data with hosted Jupyter notebooks. You can directly connect your IoT data to the notebook and build, train, and execute models right from the AWS IoT Analytics console without having to manage any of the underlying infrastructure. Using AWS IoT Analytics, you can apply machine learning algorithms to your device data to produce a health score for each device in your fleet. For example, an auto manufacturer can detect which of their customers have worn brake pads and alert them to seek

maintenance for their vehicles. With just the click of a button, you can also package your Jupyter Notebook code into an executable container image and execute that container on AWS IoT Analytics as needed.

- **Automated scaling with pay as you go pricing:** AWS IoT Analytics is a fully managed and pay-as-you go service that scales automatically to support up to petabytes of IoT data. With IoT Analytics, you can analyse your entire fleet of connected devices without managing hardware or infrastructure. As your needs change, compute power and the data store automatically scale up or down so you always have the right capacity for your IoT applications and you only pay for the resources that you use.

144.2. Backup/Restore and Disaster Recovery

This requirement is not applicable for this service. For additional information beyond what is described herein, please refer to http://aws.amazon.com/documentation/.

144.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

144.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/iotanalytics/
- **Service quotas:** https://docs.aws.amazon.com/iotanalytics/latest/userguide/limits.html
- **Service FAQs:** https://aws.amazon.com/iot-analytics/faq/

144.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/iotanalytics/ and the following links for comprehensive technical documentation regarding this service.

- User Guide: Walks through how to set up AWS IoT Analytics and integrate it with other AWS services.

145. AWS IoT Core

145.1. Service Overview

AWS IoT Core enables you to connect devices to AWS Services and other devices, secure data and interactions, process and act upon device data, enables applications to interact with devices even when they are offline and that allows you to produce low-cost Alexa built-in devices.

145.1.1. Features

- **AWS IoT Device SDK:** The AWS IoT Device SDK helps you easily and quickly connect your hardware device or your mobile application to AWS IoT Core. The AWS IoT Device SDK enables your devices to connect, authenticate, and exchange messages with AWS IoT Core using the MQTT, HTTP, or WebSockets protocols. The AWS IoT Device SDK supports C, JavaScript, and Arduino, and includes the client libraries, the developer guide, and the porting guide for manufacturers. You can also use an open source alternative or write your own SDK.

- **Device Advisor:** Device Advisor is a fully managed cloud-based test capability for validating IoT devices during development. It provides pre-built tests that helps developers to validate their IoT devices for reliable and secure connectivity with AWS IoT Core. By using Device Advisor, developers can test if their IoT devices can reliably interoperate with AWS IoT Core and follow security best practices.
- **Device Gateway:** The Device Gateway serves as the entry point for IoT devices connecting to AWS. The Device Gateway manages all active device connections and implements semantics for multiple protocols to ensure that devices are able to securely and efficiently communicate with AWS IoT Core. Currently the Device Gateway supports the MQTT, WebSockets, and HTTP 1.1 protocols.
- **Message Broker:** The Message Broker is a high throughput pub/sub message broker that securely transmits messages to and from all of your IoT devices and applications with low latency. The flexible nature of the Message Broker's topic structure allows you to send messages to, or receive messages from, as many devices as you would like.
- **Authentication and Authorization:** AWS IoT Core provides mutual authentication and encryption at all points of connection, so that data is never exchanged between devices and AWS IoT Core without a proven identity. AWS IoT Core supports the AWS method of authentication (called 'SigV4'), X.509 certificate based authentication, and customer created token based authentication (through custom authorizers.) Connections using HTTP can use any of these methods, while connections using MQTT use certificate based authentication, and connections using WebSockets can use SigV4 or custom authorizers.
- **Registry:** The Registry establishes an identity for devices and tracks metadata such as the devices' attributes and capabilities. The Registry assigns a unique identity to each device that is consistently formatted regardless of the type of device or how it connects. It also supports metadata that describes the capabilities of a device, for example whether a sensor reports temperature, and if the data are Fahrenheit or Celsius.
- **Device Shadow:** With AWS IoT Core, you can create a persistent, virtual version, or Device Shadow, of each device that includes the device's latest state so that applications or other devices can read messages and interact with the device. The Device Shadow persists the last reported state and desired future state of each device even when the device is offline. You can retrieve the last reported state of a device or set a desired future state through the API or using the rules engine.
- **Rules Engine:** The Rules Engine makes it possible to build IoT applications that gather, process, analyse and act on data generated by connected devices at global scale without having to manage any infrastructure. The Rules Engine evaluates inbound messages published into AWS IoT Core and transforms and delivers them to another device or a cloud service, based on business rules you define. A rule can apply to data from one or many devices, and it can take one or many actions in parallel.
- **Alexa Voice Service (AVS) Integration:** Alexa Built-in is a category of devices created with the Alexa Voice Service (AVS) that have a microphone and speaker. You can talk to these products directly with the wake word "Alexa," and receive voice responses and content instantly.
- **AWS IoT Core for LoRaWAN:** AWS IoT Core for LoRaWAN enables customers to connect wireless devices that use low-power, long-range wide area network (LoRaWAN) technology. Using AWS IoT Core, customers can now setup a private LoRaWAN

network by connecting their own LoRaWAN devices and gateways to the AWS Cloud - without developing or operating a LoRaWAN Network Server (LNS). This eliminates the undifferentiated development work and operational burden of managing an LNS and associated infrastructure, accelerating the network set-up time.

- **Amazon Sidewalk Integration:** Amazon Sidewalk is a shared network that helps connected devices work better through improved connectivity options. Operated by Amazon at no charge to customers, Sidewalk can help simplify new device setup, extend the low-bandwidth working range of devices, and help devices stay online even if they are outside the range of their home Wi-Fi.

145.1.2. Benefits

- **Easy to use:** Connect, manage, and scale your device fleets easily and reliably without provisioning or managing servers.
- **Range of protocols:** Choose your preferred communication protocol, including MQTT, HTTPS, MQTT over WSS, and LoRaWAN.
- **Secure:** Secure device connections and data with mutual authentication and end-to-end encryption.
- **Data management:** Filter, transform, and act upon device data on the fly, based on your defined business rules.
- **Monitor and manage industrial operations:** Build industrial IoT applications for predictive quality, maintenance, and remote operation monitoring.
- **Build differentiated consumer products:** Create connected applications for home automation, home security and monitoring, and home networking.
- **Innovate with automotive data:** Develop solutions for connected, autonomous, shared, and electric vehicle (EV) applications.
- **Develop safety products:** Design commercial applications for traffic monitoring, public safety, and health monitoring.

145.2. Backup/Restore and Disaster Recovery

This requirement is not applicable for this service. For additional information beyond what is described herein, please refer to http://aws.amazon.com/documentation/.

145.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

145.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/iot/
- **Service quotas:** https://docs.aws.amazon.com/general/latest/gr/iot-core.html
- **Service FAQs:** https://aws.amazon.com/iot-core/faqs/

145.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/iot/ and the following links for comprehensive technical documentation regarding this service.

- Developer Guide: Walks through how to set up AWS IoT and integrate it with other services.
- AWS IoT Core for LoRaWAN: Describes how to configure wireless devices for AWS IoT Core for LoRaWAN.

146. AWS IoT Device Defender

146.1. Service Overview

AWS IoT Device Defender is a fully managed service for auditing and monitoring devices connected to AWS IoT. It assesses the cloud configuration of your IoT device fleet, provides ongoing monitoring of device activities via rule-based and ML-based Detect capabilities, triggers an alarm when an audit violation or behaviour anomaly is identified, and enables you to address issues quickly with built-in mitigation actions.

146.1.1. Features

- **Audit:** AWS IoT Device Defender audits your device-related resources (such as X.509 certificates, IoT policies, and Client IDs) against AWS IoT security best practices (for example, the principle of least privilege or unique identity per device). AWS IoT Device Defender reports configurations that are out of compliance with security best practices, such as multiple devices using the same identity, or overly permissive policies that can allow one device to read and update data for many other devices.
- **Rules Detect:** AWS IoT Device Defender detects unusual device behaviours that may be indicative of a compromise by continuously monitoring high-value security metrics from the device and AWS IoT Core (e.g., the number of listening TCP ports on your devices or authorization failure counts). You can specify normal device behaviour for a group of devices by setting up behaviours (rules) for these metrics. AWS IoT Device Defender monitors and evaluates each datapoint reported for these metrics against user-defined behaviour (rules) and alerts you if an anomaly is detected.
- **ML Detect:** AWS IoT Device Defender monitors and identifies anomalous datapoints for six cloud-side metrics (e.g., authorization failure counts, message sent counts) and seven device-side metrics (e.g., packets out, listening TCP port counts) with machine learning (ML) models and triggers an alarm if an anomaly is detected. AWS IoT Device Defender removes the need to define accurate behaviours of your devices and automatically sets them with ML models using your device data from a trailing 14-day period. It then retrains the models each day (as long as it has sufficient amount of data to retrain on) to refresh the expected device behaviours based on the latest trailing 14 days. ML Detect makes getting started with monitoring easy.
- **Mitigation actions:** AWS IoT Device Defender enables you to use built-in mitigation actions to perform steps on Audit and Detect alarms such as adding things to a thing group, replacing default policy version and updating device certificate.
- **Alerting:** AWS IoT Device Defender publishes alarms to the AWS IoT console, AWS IoT Device Defender API, Amazon CloudWatch, and Amazon SNS if you configured SNS topics to receive Device Defender alarms.

146.1.2. Benefits

- **Audit device configurations for security vulnerabilities:** AWS IoT Device Defender audits IoT configurations associated with your devices against a set of defined IoT security best practices so you know exactly where you have security gaps. You can run audits on a continuous or ad-hoc basis. AWS IoT Device Defender comes with security best practices that you can select and run as part of the audit. For example, you can create an audit to check for identity certificates that are inactive, revoked, expiring, or pending transfer in less than 7 days. Audits make it possible for you to receive alerts as your IoT configuration is updated.
- **Continuously monitor device behaviour to identify anomalies:** AWS IoT Device Defender detects anomalies in device behaviour that may indicate a compromised device by monitoring high-value security metrics from the cloud and AWS IoT Core and comparing them against expected device behaviour that you define. For example, AWS IoT Device Defender lets you define how many ports are open on the device, who the device can talk to, where it is connecting from, and how much data it sends or receives. AWS IoT Device Defender also allows you to use machine learning models to set device normal behaviour, for example, the number of times your devices connect with AWS IoT cloud every five minutes. Then, it monitors the device communication and traffic and alerts you if something looks wrong according to defined behaviours or ML models, like traffic from devices to a known malicious IP or a spike in connection attempts.
- **Receive alerts and take action:** AWS IoT Device Defender publishes security alarms to the AWS IoT Console, Amazon CloudWatch, and Amazon SNS when an audit fails or when behaviour anomalies are detected so you can investigate and determine the root cause. For example, AWS IoT Device Defender can alert you when device identities are accessing sensitive APIs. AWS IoT Device Defender also provides built-in mitigation actions you can take to minimize the impact of security issues such as adding a thing to a thing group (for example, quarantine), updating a device certificate, replacing default policy version and enabling IoT logging.

146.2. Backup/Restore and Disaster Recovery

This requirement is not applicable for this service. For additional information beyond what is described herein, please refer to http://aws.amazon.com/documentation/.

146.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

146.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/iot/latest/developerguide/device-defender.html
- **Service quotas:** https://docs.aws.amazon.com/general/latest/gr/iot_device_defender.html
- **Service FAQs:** https://aws.amazon.com/iot-device-defender/faq/

146.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/iot/latest/developerguide/device-defender.html for comprehensive technical documentation regarding this service.

147. AWS IoT Device Management

147.1. Service Overview

AWS IoT Device Management makes it easy to securely register, organize, monitor, and remotely manage IoT devices at scale. With AWS IoT Device Management, you can register your connected devices individually or in bulk, and easily manage permissions so that devices remain secure. You can also organize your devices, monitor and troubleshoot device functionality, query the state of any IoT device in your fleet, and send firmware updates over-the-air (OTA)—all through a fully managed web application. AWS IoT Device Management is agnostic to device type and OS, so you can manage devices from constrained microcontrollers to connected cars all with the same service. AWS IoT Device Management allows you to scale your fleets and reduce the cost and effort of managing large and diverse IoT device deployments.

147.1.1. Features

- **Easily Register Connected Devices in Bulk:** AWS IoT Device Management helps you register new devices by using the IoT management console or API to upload templates that you populate with information like device manufacturer and serial number, X.509 identity certificates, or security policies. Then, you can configure the entire fleet of devices with this information with a few clicks in the management console.
- **Organize Connected Devices into Groups:** With AWS IoT Device Management, you can group your device fleet into a hierarchical structure based on function, security requirements, or any other category. You can group one device in a room, group devices together that operate on the same floor, or group all the devices that operate within a building. Then, you can use these groups to manage access policies, view operational metrics, or perform actions on your devices across the entire group. You can also automate organization of your devices with dynamic thing groups. Your dynamic thing groups will automatically add devices that meet your specified criteria and remove the devices that no longer match the criteria.
- **Fleet Indexing and Search:** AWS IoT Device Management makes it easy to query a group of devices and aggregate statistics on device records based on any combination of device attribute, state and connectivity indexing so that you can better organize and understand your fleet. For example, you can search for a group of connected temperature sensors in a manufacturing facility, count the number of sensors with a specific firmware version, and find the average temperature reading for those sensors.
- **Fine-Grained Device Logging:** AWS IoT Device Management lets you collect device logs so that in the event of a problem you can query the log data to figure out what went wrong. You can configure the logs to include only the metrics that are critical to device performance so you can identify issues quickly. For example, you can include device metrics like error codes that indicate download failures or device restart counters, and quickly identify and troubleshoot issues on devices within the device group.
- **Remotely Manage Connected Devices:** AWS IoT Device Management allows you to push software and firmware to devices in the field to patch security vulnerabilities and

improve device functionality. You can execute bulk updates, control deployment velocity, set failure thresholds, and define continuous jobs to update device software automatically so they are always running the latest version. You can send actions such as device reboots or factory resets remotely to fix software issues in the device or restore the device to its original settings. You can also digitally sign files that you send to your devices, helping to ensure your devices are not compromised.

- **Secure Tunneling:** AWS IoT Device Management supports the creation of a device tunnel - a secure remote communications session to a device. This provides secure connectivity to individual devices, which you can then use to diagnose issues and act to solve in just a few clicks. You can also make multiple, concurrent client connections over a single secure tunnel, enabling you to perform more advanced device troubleshooting, such as issuing remote shell commands to a device while simultaneously debugging a web application on the same device.
- **Fleet Hub:** AWS IoT Device Management includes the ability to create no-code, fully-managed web applications using Fleet Hub to visualize and interact with your device fleets connected to AWS IoT. With Fleet Hub, you can search across your large and diverse fleets and view device state and health data, in near real time—such as connection status, firmware version, country code, or battery level.

147.1.2. Benefits

- **Fast device registration:** With AWS IoT Device Management, you can securely add device attributes like device name, type and manufacturing year, certificates and access policies to the IoT Registry in bulk, assign them to devices, and put large fleets of connected devices into service quickly.
- **Simple IoT device organization:** AWS IoT Device Management lets you organize your devices into groups and manage access policies for these groups. This makes it easy to track, operate, and manage your devices according to business and security requirements, such as deploying a firmware update for all devices in a building or defining how devices communicate with each other. You can create a hierarchy for your groups such as grouping multiple sensors within a single vehicle and grouping multiple vehicles in a fleet. Then, your devices will inherit access policies based on the group hierarchy.
- **Locate connected devices quickly:** AWS IoT Device Management lets you quickly search and find any IoT device across your entire device fleet in near real-time. You can easily find devices based on a combination of attributes like device ID, device state, and type, so that you can take action or troubleshoot.
- **Easy remote device management:** AWS IoT Device Management makes it easy for you to maintain the health of your device fleet. With AWS IoT Device Management, you can remotely update the software running on your devices after they have been deployed in the field – allowing you to ensure that devices are always running on the latest software. You can also remotely execute fleet-wide actions such as reboots, factory resets, and security patches.

147.2. Backup/Restore and Disaster Recovery

This requirement is not applicable for this service. For additional information beyond what is described herein, please refer to http://aws.amazon.com/documentation/.

147.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

147.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/iot-device-management/
- **Service quotas:** https://docs.aws.amazon.com/general/latest/gr/iot_device_management.html
- **Service FAQs:** https://aws.amazon.com/iot-device-management/faq/

147.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/iot-device-management/ and the following links for comprehensive technical documentation regarding this service.

- Developer Guide: Describes device provisioning, thing groups, and jobs that are sent to and run on one or more devices connected to AWS IoT Core.
- User Guide: You can use Fleet Hub for AWS IoT Device Management to build standalone web applications for monitoring the health of your device fleets.

148. AWS IoT Events

148.1. Service Overview

AWS IoT Events is a fully managed service that makes it easy to detect and respond to events from IoT sensors and applications. Events are patterns of data identifying more complicated circumstances than expected, such as changes in equipment when a belt is stuck or motion detectors using movement signals to activate lights and security cameras. Before IoT Events, you had to build costly, custom applications to collect data, apply decision logic to detect an event, and then trigger another application to react to the event.

Using IoT Events, it's simple to detect events across thousands of IoT sensors sending different telemetry data, such as temperature from a freezer, humidity from respiratory equipment, and belt speed on a motor. You simply select the relevant data sources to ingest, define the logic for each event using simple 'if-then-else' statements, and select the alert or custom action to trigger when an event occurs. IoT Events continuously monitors data from multiple IoT sensors and applications, and it integrates with other services, such as AWS IoT Core and AWS IoT Analytics, to enable early detection and unique insights into events. IoT Events automatically triggers alerts and actions in response to events based on the logic you define to resolve issues quickly, reduce maintenance costs, and increase operational efficiency.

148.1.1. Features

- **Telemetry data ingestion:** AWS IoT Events helps you evaluate telemetry data to detect events in equipment or a process. IoT Events ingests raw data from any device connected to AWS IoT, processed data from AWS IoT Analytics, and data from third party applications via IoT Events' direct ingest APIs.
- **Event detection:** AWS IoT Events evaluates multiple telemetry inputs to detect events and derive the state of processes, equipment, or products by applying user-defined,

conditional logic. Events are patterns of data identifying more complicated circumstances than expected, such as changes in equipment when a belt is stuck or motion detectors using movement signals to activate lights and security cameras. You can schedule maintenance and send alarms or alerts prior to device failure.

- **Integration with analytics tools and other AWS services:** AWS IoT Events can leverage output from advanced analytics services to make better decisions. With easy integrations to and from other AWS services, you can further optimize operations. You can complete an event detector setup in AWS IoT Events, write your event logic using simple 'if-then-else' statements, and select the alert or custom action to trigger when the event occurs. Such conditional statements can receive inputs from either raw data or data first processed by AWS IoT Analytics (where you can apply machine learning models to aid in detection of more complex events).
- **Scalability:** AWS IoT Events easily scales when you are connecting many devices. Define a model once for a specific device, and the service will automatically scale and manage all devices of that model that connect to IoT Events. If an event is detected on one device or many, IoT Events can trigger the appropriate reaction or alert.
- **Alarms:** AWS IoT Events can evaluate equipment behaviour or identify equipment performance issues based on industrial data in the cloud. For an asset data property that you want to monitor, you can define an alarm rule to apply (e.g. rotations per minute is greater than a user defined value), select the severity for this alarm definition (e.g. severity values of 1, 2, 3 and 4), and configure the notifications to send when an alarm is triggered (e.g. Email and SMS). Once an alarm has been defined, operators can manage the alarm workflow by taking actions to acknowledge, snooze or disable the alarm.

148.1.2. Benefits

- **Easily ingest operations data:** With AWS IoT Events, you can easily evaluate multiple sources of telemetry data to detect the state of processes, equipment, or products quickly and schedule maintenance and send alarms or alerts to support teams and trigger actions, such as shutting down malfunctioning equipment before more damage is done. For example, you could use IoT Events to quickly build a food spoilage notification system that notifies technicians before spoilage occurs if a freezer is malfunctioning. Faster identification of events can prevent food spoilage and waste, saving thousands of dollars in potential lost revenue.
- **Trigger a range of actions:** In IoT Events, you can combine multiple sources of telemetry data like belt speed, motor voltage, amperage, and noise levels, then define conditional logic to apply to that data, to gain full insight into your equipment and processes. This visibility helps you better understand events, such as when a motor might be stuck. You can also select a pre-built action to trigger, such as sending a message to the motor to shut down, before putting equipment at greater risk.
- **Easily build rules:** You can write event logic, using simple 'if-then-else' statements, to identity critical events using sensor attributes, such as temperature and pressure. These attributes can trigger automatic alerts and pre-defined responses from IoT Events. For example, you can specify events related to a welding robot, such as an arm becoming misaligned. When IoT Events detects the incident, it automatically issues an alert and triggers the appropriate response.

148.2. Backup/Restore and Disaster Recovery

This requirement is not applicable for this service. For additional information beyond what is described herein, please refer to http://aws.amazon.com/documentation/.

148.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

148.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/iotevents/
- **Service quotas:** https://docs.aws.amazon.com/iotevents/latest/developerguide/iotevents-quotas.html
- **Service FAQs:** https://aws.amazon.com/iot-events/faqs/

148.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/iotevents/ and the following links for comprehensive technical documentation regarding this service.

- Developer Guide: Describes key concepts of AWS IoT Events and provides instructions for using the service.

149. AWS IoT Greengrass

149.1. Service Overview

AWS IoT Greengrass is an open-source edge runtime and cloud service for building, deploying, and managing device software.

149.1.1. Features

- **Local processing for AWS Lambda:** AWS IoT Greengrass includes support for AWS Lambda. With AWS IoT Greengrass, you can run AWS Lambda functions on the device to respond quickly to local events, interact with local resources, and process data to minimize the cost of transmitting data to the cloud.
- **Local support for containers:** You can deploy, run, and manage Docker containers on AWS IoT Greengrass devices. Your Docker images can be stored in Docker container registries, such as Amazon Elastic Container Registry (Amazon ECR), Docker Hub, or private Docker Trusted Registries (DTRs).
- **Local support for AWS IoT Device Shadows:** AWS IoT Greengrass also includes the functionality of AWS IoT Device Shadows. The Device Shadow caches the state of your device, like a virtual version or “shadow,” of each device that tracks the device’s current versus desired state and synchronizes that state with the cloud when connectivity is available.
- **Local messaging:** AWS IoT Greengrass enables messaging between the AWS IoT Greengrass Core and devices using the AWS IoT Device SDK on a local network, facilitating communication even when there is no connection to AWS. With AWS IoT

Greengrass, your devices can process messages and deliver them to another device or to the cloud based on business rules you define.

- **Local resource access:** AWS Lambda functions deployed on an AWS IoT Greengrass Core can access local resources that are attached to the device. This allows you to use serial ports, peripherals such as add-on security devices, sensors and actuators, on-board GPUs, or the local file system to quickly access and process local data.
- **Local development:** AWS IoT Greengrass lets you rapidly develop and debug code on a test device before using the cloud to deploy to your production devices. You can use the AWS IoT Greengrass command-line interface (CLI) to locally develop and debug applications on your device, and the local debug console to help you visually debug applications.
- **AWS IoT Greengrass ML Inference:** AWS IoT Greengrass ML Inference is a feature of AWS IoT Greengrass that makes it easy to perform machine learning inference locally on AWS IoT Greengrass devices using models that are built and trained in the cloud. This means you won't incur data transfer costs or increased latency for applications that use machine learning inference. To learn more about the ML Inference feature, click here.
- **Stream Manager for AWS IoT Greengrass:** You can use AWS IoT Greengrass to collect, process, and export data streams from IoT devices and manage the life cycle of that data on the device to minimize development time. AWS IoT Greengrass provides a standard mechanism to process data streams, manage local data-retention policies, and transmit device data to AWS cloud services such as Amazon Simple Storage Service (Amazon S3), Amazon Kinesis, AWS IoT Core, and AWS IoT Analytics.
- **AWS IoT Greengrass components:** AWS IoT Greengrass provides pre-built components for common use cases so you can discover and import, configure, and deploy applications and services at the edge without the need to understand different device protocols, manage credentials, or interact with external APIs. You can also create your own components or simply re-use common business logic from one AWS IoT Greengrass device to another.
- **Manage IoT applications at scale:** AWS IoT Greengrass makes it easy to remotely deploy and manage device software on millions of devices. You can organize your devices in groups and deploy and manage device software and configuration to a subset of devices or to all devices at once. AWS IoT thing groups allow you to group multiple AWS IoT Greengrass devices, view deployment history, and start or stop deployments.
- **Over the air updates:** AWS IoT Greengrass provides the ability to update the AWS IoT Greengrass Core software on AWS IoT Greengrass devices. You can use the AWS IoT Greengrass console, APIs, or command-line interface to update the version of AWS IoT Greengrass Cores or components running on your devices in order to deploy security updates, bug fixes, and new AWS IoT Greengrass features.
- **AWS IoT Greengrass Secrets Manager:** AWS IoT Greengrass Secrets Manager allows you to securely store, access, rotate, and manage secrets – credentials, keys, endpoints, and configurations – at the edge. With AWS IoT Greengrass components integration, if an AWS IoT Greengrass component needs a secret to authenticate with an application or service, you can select and deploy a secret to the AWS IoT Greengrass Core as part of the component configuration.

- **Hardware Security Integration:** AWS IoT Greengrass offers customers the option to store their device private key on a hardware secure element. You can store sensitive device information at the edge with AWS IoT Greengrass Secrets Manager and encrypt your secrets using private keys for root of trust security. For a list of eligible hardware partners, visit the AWS Partner Device Catalog.
- **AWS IoT Device Tester for AWS IoT Greengrass:** AWS IoT Device Tester for AWS IoT Greengrass is a test automation tool that helps you validate if your device meets the software and hardware requirements to run AWS IoT Greengrass. It supports configuration and dependency checks and end-to-end tests to validate if a device can support specific AWS IoT Greengrass features such as Machine Learning Inference. Additionally, hardware partners can download signed qualification reports from Device Tester and submit these reports to AWS Partner Central to qualify and list devices in the AWS Partner Device Catalog.

149.1.2. Benefits

- **Run at the edge:** AWS IoT Greengrass makes it easy to bring intelligence to edge devices, such as for anomaly detection in precision agriculture or powering autonomous devices.
- **Manage apps:** Deploy new or legacy apps across fleets using any language, packaging technology, or runtime.
- **Control fleets:** Manage and operate device fleets in the field locally or remotely using MQTT or other protocols.
- **Process locally:** Collect, aggregate, filter, and send data locally. Manage and control what data goes to the cloud for optimized analytics and storage.

149.2. Backup/Restore and Disaster Recovery

This requirement is not applicable for this service. For additional information beyond what is described herein, please refer to http://aws.amazon.com/documentation/.

149.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

149.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/greengrass/
- **Service quotas:** https://docs.aws.amazon.com/general/latest/gr/greengrassv2.html
- **Service FAQs:** https://aws.amazon.com/greengrass/faqs/

149.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/greengrass/ and the following links for comprehensive technical documentation regarding this service.

- Developer Guide: Walks through how to set up AWS IoT Greengrass V2 and integrate it with other services.

150. AWS IoT SiteWise

150.1. Service Overview

AWS IoT SiteWise is a managed service that makes it easy to collect, organize, and analyse data from industrial equipment at scale.

150.1.1. Features

- **Time series storage integrated with your industrial data lake:** Use AWS IoT SiteWise to store industrial data generated from your equipment in a fast and scalable time series data store. AWS IoT SiteWise storage supports two tiers for equipment data: a hot tier optimized for real-time applications and a cold tier optimized for analytical applications. SiteWise helps you to reduce storage cost by keeping recent data in hot tier and moving historical data to a cost optimized storage tier based upon policies.
- **Asset modelling:** Use AWS IoT SiteWise to build models of your physical operations that represent your assets, processes, and facilities, which will help you understand industrial data in the context of your equipment. Once your models are created, you can define an asset hierarchy to accurately represent relationships between devices and equipment within a single facility or across multiple facilities.
- **Asset metrics:** Map data streams and define static or computed equipment and process properties across all facilities so they're readily available for analysis. Using a built-in library of operators and functions, you can create two types of custom computations: transforms and metrics. You can define transforms that trigger when equipment data arrives and metrics computed at user-defined intervals that can be configured for an asset or rolled up from a group of assets.
- **SiteWise Edge:** AWS IoT SiteWise includes AWS IoT SiteWise Edge, on-premises software used to collect, organize, process, and monitor equipment data locally before sending it to AWS. SiteWise Edge runs on local hardware such as third-party industrial gateways and computers, or on AWS Outposts and AWS Snow Family compute devices. SiteWise Edge uses AWS IoT Greengrass, which provides a local software runtime environment for edge devices to help build, deploy, and manage applications.
- **Data ingestion:** In addition to ingesting data with AWS IoT SiteWise Edge, AWS IoT SiteWise supports other data ingestion methods including MQ Telemetry Transport (MQTT) protocol integration with AWS IoT Core. Use the AWS IoT Message Broker to subscribe to a topic that is publishing messages from your industrial equipment, then use the AWS IoT Core Rules Engine to route messages to AWS IoT SiteWise.
- **Gateway management:** Configure and monitor edge gateways across all facilities, and view a consolidated list of active gateways through the console or APIs. Monitor gateway health remotely to view the status of all production lines from one place. Using the Amazon CloudWatch Metrics console, you can also view gateway metrics to monitor the health, status, and performance of your gateway resources.
- **AWS IoT SiteWise Monitor:** AWS IoT SiteWise allows you to create no-code, fully managed web applications using AWS IoT SiteWise Monitor. With this feature, you can visualize and interact with operational data from devices and equipment connected to AWS IoT services.

- **Alarms:** To assess equipment behaviour or identify equipment performance issues, you can define and update alarms, and set alarm notifications using the AWS IoT SiteWise console, AWS IoT SiteWise Monitor, or AWS IoT SiteWise software development kit (SDK).
- **Extensibility:** Custom edge and cloud applications can use query APIs to easily retrieve asset data and computed metrics from the AWS IoT SiteWise time series data store, or a publish/subscribe interface to consume a near real-time stream of structured IoT data.

150.1.2. Benefits

- **Ready to use:** Collect, manage, and visualize data from all your industrial equipment sources without developing additional software.
- **Remote monitoring:** Identify and resolve issues faster through remote equipment performance monitoring.
- **Optimise:** Optimize processes across your facility portfolio with insights from automatic, customizable data visualizations.
- **Seamless:** Collect and process industrial data locally, and build hybrid industrial applications that work seamlessly across the edge and cloud.

150.2. Backup/Restore and Disaster Recovery

This requirement is not applicable for this service. For additional information beyond what is described herein, please refer to http://aws.amazon.com/documentation/.

150.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

150.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/iot-sitewise/
- **Service quotas:** https://docs.aws.amazon.com/iot-sitewise/latest/userguide/quotas.html
- **Service FAQs:** https://aws.amazon.com/iot-sitewise/faqs/

150.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/iot-sitewise/ and the following links for comprehensive technical documentation regarding this service.

- User Guide: Describes key concepts of AWS IoT SiteWise and provides instructions for using the features of AWS IoT SiteWise.
- Application Guide: Describes key concepts of the AWS IoT SiteWise Monitor portal application and provides instructions to use its features.

151. AWS IoT Things Graph

151.1. Service Overview

AWS IoT Things Graph is a service that makes it easy to visually connect different devices and web services to build IoT applications. AWS IoT Things Graph provides a visual drag-and-drop interface for connecting and coordinating interactions between devices and web services, so you can build IoT applications quickly. For example, in a commercial agriculture application, you can define interactions between humidity, temperature, and sprinkler sensors with weather data services in the cloud to automate watering. You represent devices and services using pre-built reusable components, called models, that hide low-level details, such as protocols and interfaces, and are easy to integrate to create sophisticated workflows.

You can get started with AWS IoT Things Graph using these pre-built models for popular device types, such as cameras, motion sensors, and switches, as well as web services such as Amazon Simple Storage Service (S3) or Amazon Rekognition, or create your own custom models. You can deploy and run your IoT applications to the AWS Cloud or AWS IoT Greengrass-enabled devices such as edge gateways and cable set-top boxes, in just a few clicks. AWS IoT Greengrass is software that provides local compute and secure cloud connectivity so devices can respond quickly to local events even without internet connectivity, and runs on a huge range of devices from a Raspberry Pi to a server-level appliance.

151.1.1. Features

- **Models:** AWS IoT Things Graph helps you build IoT applications faster by reducing the time spent understanding low-level device details and writing code to make devices and web services work together. AWS IoT Things Graph makes it easy to work with devices and web services by allowing you to represent them as models. A model is an abstraction that represents a device as a set of actions (inputs), events (outputs), and states (attributes). Models separate the device interface from its underlying implementation. For example, a switch can be represented as a set of attributes (status, dimmable), events (daylight saving time ends), and actions (turn on).
- **Model Repository:** AWS IoT Things Graph makes it easy to reuse models so you don’t need to duplicate code for every IoT application deployment. You can use the model editor in the AWS IoT Things Graph console to build your own model using AWS IoT Things Graph's GraphQL-based schema modelling language, or choose from models for common devices such as light switches and temperature sensors. Once created, models are saved in your model repository where they can be accessed and reused across your applications. As a result, you get reusable building blocks for your IoT applications.
- **Mappings Library:** AWS IoT Things Graph eliminates the need to write code to convert the output of one device into the input of another using a mapping library. For example, a ZigBee based motion sensor cannot talk to a Z-Wave based camera due to differences in device details such as APIs, protocols, and message syntax. Mappings transform low-level device details from one device into a format understood by another device, allowing them to interact without needing any software changes. AWS IoT Things Graph’s built-in mapping library provides hundreds of common concepts for common IoT applications in industrial and the connected home, such as brightness, colour, and volume, or you can build your own.
- **Workflows:** AWS IoT Things Graph simplifies application development by providing a drag-and-drop interface in the AWS IoT Things Graph console. In the drag-and-drop

interface, you can visually build applications by connecting models, defining interactions between them, and building a workflow. Workflows are made up of flows, which consists of multiple things (devices and services) connected in a sequence of steps. The order of the steps can be changed, and new devices and business logic can be added to evolve an application without revising the entire IoT application. Workflows are triggered by telemetry from a device. Once triggered, AWS IoT Things Graph executes each step of the workflow. AWS IoT Things Graph tracks the state of each step and retries if something goes wrong.

- **Run at the Edge:** AWS IoT Things Graph applications can run in the AWS Cloud or at the edge, such as on AWS IoT Greengrass-enabled devices, so they can respond to local events quickly, even without an internet connection. AWS IoT Greengrass is software that lets you securely run local compute, messaging, data caching, sync, and machine inference capabilities. Deployment is easy and can be initiated with just a few clicks from the AWS IoT Things Graph console. AWS IoT Things Graph bundles models along with the run-time, and pushes it to your IoT Greengrass device where it listens to messages and coordinates interactions.
- **Application Monitoring:** AWS IoT Things Graph gives you visibility into how your application is performing so you easily tune your application and fix any defects. AWS IoT Things Graph will emit success, failure, and execution time metrics so you can monitor and manage your application from the console. AWS IoT Things Graph stores the entire IoT application execution history in a data store, and exposes APIs so you know exactly what happened in your application.

151.1.2. Benefits

- **Build IoT applications faster:** AWS IoT Things Graph provides reusable models that represent devices and web services and bridge differences in low-level details such as communication protocols and proprietary interfaces. It's easy to combine models together to create IoT applications using a visual interface. You can use a library of pre-built models for popular devices types such as cameras, motion sensors, and switches, as well as web services such as Amazon Simple Storage Service (S3), Amazon Rekognition, or AWS Lambda or create your own custom models.
- **Easily create sophisticated workflows:** AWS IoT Things Graph provides a visual way to represent complex processes, such as welding car frames on a manufacturing line, automating shutdowns of production lines when anomalies are detected, and implementing building lockdowns when suspicious behaviour is identified, as a visual workflow.
- **Easy to manage and monitor:** With just a few clicks, AWS IoT Things Graph packages and deploys your IoT application to the AWS Cloud or AWS IoT Greengrass-enabled devices. AWS IoT Things Graph coordinates interactions between devices and web services and retries any failed steps to keep your workflow running smoothly. Once deployed, you can use AWS CloudWatch to monitor your flows by collecting and processing workflow data, as well as set alarms and actions around flow performance thresholds.

151.2. Backup/Restore and Disaster Recovery

This requirement is not applicable for this service. For additional information beyond what is described herein, please refer to http://aws.amazon.com/documentation/.

151.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

151.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/thingsgraph/
- **Service quotas:** https://docs.aws.amazon.com/general/latest/gr/iot_thingsgraph.html
- **Service FAQs:** https://aws.amazon.com/iot-things-graph/faqs/

151.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/thingsgraph/ and the following links for comprehensive technical documentation regarding this service.

- User Guide: Walks through how to set up the service and integrate with other AWS services.

152. AWS Key Management Service

152.1. Service Overview

AWS Key Management Service (KMS) makes it easy for you to create and manage cryptographic keys and control their use across a wide range of AWS services and in your applications. AWS KMS is a secure and resilient service that uses hardware security modules that have been validated under FIPS 140-2, or are in the process of being validated, to protect your keys. AWS KMS is integrated with AWS CloudTrail to provide you with logs of all key usage to help meet your regulatory and compliance needs.

152.1.1. Features

- **Centralized Key Management:** AWS KMS provides you with centralized control over the lifecycle and permissions of your keys. You can create new keys whenever you wish, and you can control who can manage keys separately from who can use them. As an alternative to using keys generated by AWS KMS, you can import keys from your own key management infrastructure, or use keys stored in your AWS CloudHSM cluster. You can choose automatic rotation of root keys generated in AWS KMS once per year without the need to re-encrypt previously encrypted data. The service automatically keeps older versions of the root key available to decrypt previously encrypted data. You can manage your root keys and audit their usage from the AWS Management Console or by using the AWS SDK or AWS Command Line Interface (CLI).
- **AWS Service Integration:** AWS KMS integrates with AWS services to encrypt data at rest, or to facilitate signing and verification using an AWS KMS key. To protect data at rest, integrated AWS services use envelope encryption, where a data key is used to encrypt data, and is itself encrypted under a KMS key stored in AWS KMS. For signing and verification, integrated AWS services use a key pair from an asymmetric KMS key in AWS KMS. For more details about how an integrated service uses AWS KMS, see the documentation for your AWS service.

- **Audit Capabilities:** If you have AWS CloudTrail enabled for your AWS account, each request you make to AWS KMS is recorded in a log file that is delivered to the Amazon S3 bucket that you specified when you enabled AWS CloudTrail. The information recorded includes details of the user, time, date, API action and, when relevant, the key used.
- **Scalability, Durability, and High Availability:** AWS KMS is a fully managed service. As your use of encryption grows, the service automatically scales to meet your needs. It enables you to manage thousands of KMS keys in your account and to use them whenever you want. It defines default limits for number of keys and request rates, but you can request increased limits if necessary.
- **Secure:** AWS KMS is designed so that no one, including AWS employees, can retrieve your plaintext keys from the service. The service uses hardware security modules (HSMs) that have been validated under FIPS 140-2, or are in the process of being validated, to protect the confidentiality and integrity of your keys. Your plaintext keys are never written to disk and only ever used in volatile memory of the HSMs for the time needed to perform your requested cryptographic operation.
- **Custom Key Store:** AWS KMS provides the option for you to create your own key store using HSMs that you control. Each custom key store is backed by an AWS CloudHSM cluster. When you create a KMS key in a custom key store, the service generates and stores key material for the KMS key in an AWS CloudHSM cluster that you own and manage. When you use a KMS key in a custom key store, the cryptographic operations under that key are performed in your AWS CloudHSM cluster.
- **Asymmetric Keys:** AWS KMS provides you the capability to create and use asymmetric KMS keys and data key pairs. You can designate a KMS key for use as a signing key pair or an encryption key pair. Key pair generation and asymmetric cryptographic operations using these KMS keys are performed inside HSMs. You can request the public portion of the asymmetric KMS key for use in your local applications, while the private portion never leaves the service.

152.1.2. Benefits

- **Fully managed:** You control access to your encrypted data by defining permissions to use keys while AWS KMS enforces your permissions and handles the durability and physical security of your keys.
- **Centralized key management:** AWS KMS presents a single control point to manage keys and define policies consistently across integrated AWS services and your own applications. You can easily create, import, rotate, delete, and manage permissions on keys from the AWS Management Console or by using the AWS SDK or CLI.
- **Manage encryption for AWS services:** AWS KMS is integrated with AWS services to simplify using your keys to encrypt data across your AWS workloads. You choose the level of access control that you need, including the ability to share encrypted resources between accounts and services. KMS logs all use of keys to AWS CloudTrail to give you an independent view of who accessed your encrypted data, including AWS services using them on your behalf.
- **Encrypt data in your applications:** AWS KMS is integrated with the AWS Encryption SDK to enable you to used KMS-protected data encryption keys to encrypt locally within

your applications. Using simple APIs you can also build encryption and key management into your own applications wherever they run.

- **Digitally sign data:** AWS KMS enables you to perform digital signing operations using asymmetric key pairs to ensure the integrity of your data. Recipients of digitally signed data can verify the signatures whether they have an AWS account or not.
- **Low cost:** There is no commitment and no upfront charges to use AWS KMS. You only pay US $1/month to store any key that you create. AWS managed keys that are created on your behalf by AWS services are free to store. You are charged per-request when you use or manage your keys beyond the free tier.
- **Secure:** AWS KMS uses hardware security modules (HSMs) that have been validated under FIPS 140-2, or are in the process of being validated, to generate and protect keys. Your keys are only used inside these devices and can never leave them unencrypted. KMS keys are never shared outside the AWS region in which they were created.
- **Compliance:** The security and quality controls in AWS KMS have been certified under multiple compliance schemes to simplify your own compliance obligations. AWS KMS provides the option to store your keys in single-tenant HSMs in AWS CloudHSM instances that you control.
- **Built-in auditing:** AWS KMS is integrated with AWS CloudTrail to record all API requests, including key management actions and usage of your keys. Logging API requests helps you manage risk, meet compliance requirements and conduct forensic analysis.

152.2. Backup/Restore and Disaster Recovery

This requirement is not applicable for this service. For additional information beyond what is described herein, please refer to http://aws.amazon.com/documentation/.

152.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

152.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/kms/
- **Service quotas:** https://docs.aws.amazon.com/kms/latest/developerguide/limits.html
- **Service FAQs:** https://aws.amazon.com/kms/faqs/

152.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/kms/ and the following links for comprehensive technical documentation regarding this service.

- Developer Guide: Provides conceptual overviews of AWS Key Management Service and explains how to use it to protect data in your own applications that use AWS.
- AWS KMS Cryptographic Details: Learn about the cryptographic operations that are run within AWS when you use AWS KMS.

153. AWS Lake Formation

153.1. Service Overview

AWS Lake Formation is a service that makes it easy to set up a secure data lake in days. A data lake is a centralized, curated, and secured repository that stores all your data, both in its original form and prepared for analysis. A data lake lets you break down data silos and combine different types of analytics to gain insights and guide better business decisions.

Creating a data lake with Lake Formation is as simple as defining data sources and what access and security policies you want to apply. Lake Formation then helps you collect and catalogue data from databases and object storage, move the data into your new Amazon Simple Storage Service (S3) data lake, clean and classify your data using ML algorithms, and secure access to your sensitive data using granular controls at the column, row, and cell-levels. Your users can access a centralized data catalogue that describes available datasets and their appropriate usage. They then use these datasets with their choice of analytics and ML services, such as Amazon Redshift, Amazon Athena, Amazon EMR for Apache Spark, and Amazon QuickSight. Lake Formation builds on the capabilities available in AWS Glue.

153.1.1. Features

- **Import data from databases already in AWS:** Once you specify where your existing databases are and provide your access credentials, AWS Lake Formation reads the data and its metadata (schema) to understand the contents of the data source. It then imports the data to your new data lake and records the metadata in a central catalogue. With Lake Formation, you can import data from MySQL, PostgreSQL, SQL Server, MariaDB, and Oracle databases running in Amazon Relational Database Service (RDS) or hosted in Amazon Elastic Compute Cloud (EC2). Both bulk and incremental data loading are supported.
- **Import data from other external sources:** You can use Lake Formation to move data from on-premises databases by connecting with Java Database Connectivity (JDBC). Identify your target sources and provide access credentials in the console, and Lake Formation reads and loads your data into the data lake. To import data from databases other than the ones listed above, you can create custom ETL jobs with AWS Glue.
- **Import data from other AWS services:** Using Lake Formation, you can also pull in semi-structured and unstructured data from other Amazon Simple Storage Service (S3) data sources. You can identify existing Amazon S3 buckets containing data to copy into your data lake. Once you specify the S3 path to register your data sources and authorize access, Lake Formation reads the data and its schema. Lake Formation can collect and organize datasets, such as logs from AWS CloudTrail, AWS CloudFront, Detailed Billing Reports, and AWS Elastic Load Balancing (ELB). You can also load your data into the data lake with Amazon Kinesis or Amazon DynamoDB using custom jobs.
- **Catalog and label your data:** Lake Formation crawls and reads your data sources to extract technical metadata (such as schema definitions) and creates a searchable catalogue to describe this information for users so they can discover available datasets. You can also add your own custom labels to your data (at the table and column level) to define attributes, such as “sensitive information” and “European sales data.” Lake Formation provides a text-based search over this metadata so your users can quickly find the data they need to analyse.

- **Transform data:** Lake Formation can perform transformations on your data, such as rewriting various date formats for consistency, to ensure that the data is stored in an analytics-friendly fashion. Lake Formation creates transformation templates and schedules jobs to prepare your data for analysis. Your data is transformed with AWS Glue and written in columnar formats, such as Parquet and ORC, for better performance.
- **Clean and deduplicate data:** Lake Formation helps clean and prepare your data for analysis by providing a Machine Learning (ML) Transform called FindMatches for deduplication and finding matching records.
- **Optimize partitions:** Lake Formation also optimizes the partitioning of data in Amazon S3 to improve performance and reduce costs. Raw data that is loaded may be in partitions that are too small (requiring extra reads) or too large (reading more data than needed.) With Lake Formation, your data is organized by size, time period, and/or relevant keys. This enables both fast scans and parallel, distributed reads for the most commonly used queries.
- **Row and Cell-level security:** Lake Formation provides data filters that allow you to restrict access to a combination of columns and rows. Use row and cell-level security to protect sensitive data like Personal Identifiable Information (PII).
- **Enforce encryption:** Lake Formation uses the encryption capabilities of Amazon S3 for data in your data lake. This approach provides automatic server-side encryption with keys managed by the AWS Key Management Service (KMS). S3 encrypts data in transit when replicating across Regions and lets you use separate accounts for source and destination Regions to protect against malicious insider deletions. These encryption capabilities provide a secure foundation for all data in your data lake.
- **Define and manage access controls:** Lake Formation provides a single place to manage access controls for data in your data lake. You can define security policies that restrict access to data at the database, table, column, row, and cell levels. These policies apply to AWS Identity and Access Management (IAM) users and roles, and to users and groups when federating through an external identity provider. You can use fine-grained controls to access data secured by Lake Formation within Amazon Redshift Spectrum, Amazon Athena, AWS Glue ETL, and Amazon EMR for Apache Spark.
- **Implement audit logging:** Lake Formation provides comprehensive audit logs with CloudTrail to monitor access and show compliance with centrally defined policies. You can audit data access history across analytics and ML services that read the data in your data lake via Lake Formation. This lets you see which users or roles have attempted to access what data, with which services, and when. You can access audit logs in the same way you access any other CloudTrail logs using the CloudTrail APIs and console.
- **Governed tables:** Use ACID (atomic, consistent, isolated, and durable) transactions to allow multiple users and jobs to reliably and consistently insert data, across multiple tables on Amazon S3. Transactions for Governed Tables automatically manage conflicts and errors and ensures consistent views for all users. You can query Governed Tables using transactions from Amazon Redshift, Amazon Athena, and AWS Glue.
- **Label your data with business metadata:** With Lake Formation, you can designate data owners, such as data stewards and business units, by adding a field in table properties as custom attributes. Your owners can augment the technical metadata with business metadata that further defines appropriate uses for the data. You can specify

appropriate use cases and label the sensitivity of your data for enforcement by using Lake Formation security and access controls.

- **Enable self-service access:** Lake Formation facilitates requesting and vending access to datasets to give your users self-service access to the data lake for a variety of analytics use cases. You can specify, grant, and revoke permissions on tables defined in the central data catalogue. The same data catalogue is available for multiple accounts, groups, and services.
- **Discover relevant data for analysis:** With Lake Formation, your users enjoy online, text-based search and filtering of datasets recorded in the central data catalogue. They can search for relevant data by name, contents, sensitivity, or any other custom labels you have defined.
- **Combine analytics approaches for more insights:** With Lake Formation, you can give your analytics users the ability to directly query datasets with Athena for SQL, Redshift for data warehousing, AWS Glue for data integration and preparation and EMR for Apache Spark–based big data processing and ML (Zeppelin notebooks). Once you point these services to Lake Formation, the datasets available are shown in the catalogue and access controls are enforced consistently, allowing your users to readily combine analytics approaches on the same data.

153.1.2. Benefits

- **Build data lakes quickly:** With Lake Formation, you can move, store, catalogue, and clean your data faster. You simply point Lake Formation at your data sources, and it crawls those sources and moves the data into your new Amazon S3 data lake. Lake Formation organizes data in S3 around frequently used query terms and into right-sized chunks to increase efficiency. It also changes data into formats such as Apache Parquet and ORC for faster analytics. In addition, Lake Formation has built-in ML to deduplicate and find matching records (two entries that refer to the same thing) to increase data quality.
- **Simplify security management:** Lake Formation provides a single place to define and enforce access controls that operate at the table, column, row, and cell-level for all the users and services that access your data. Your policies are consistently implemented, eliminating the need to manually configure them across security services such as AWS Identity and Access Management (IAM) and AWS Key Management Service (KMS), storage services such as S3, and analytics and ML services such as Redshift, Athena, AWS Glue, and EMR for Apache Spark. This reduces the effort in configuring policies across services and provides consistent enforcement and compliance.
- **Provide self-service access to data:** With Lake Formation, you build a data catalogue that describes the different datasets available, along with which groups of users have access to each. This makes your users more productive by helping them find the right dataset to analyse. By providing a catalogue of your data with consistent security enforcement, Lake Formation makes it easier for your analysts and data scientists to use their preferred analytics service. They can use EMR for Apache Spark, Redshift, Athena, AWS Glue, and Amazon QuickSight on diverse datasets now housed in a single data lake. Users can also combine these services without having to move data between silos.

153.2. Backup/Restore and Disaster Recovery

This requirement is not applicable for this service. For additional information beyond what is described herein, please refer to http://aws.amazon.com/documentation/.

153.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

153.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/lake-formation/
- **Service quotas:** https://docs.aws.amazon.com/general/latest/gr/lake-formation.html
- **Service FAQs:** https://aws.amazon.com/lake-formation/faqs/

153.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/lake-formation/ and the following links for comprehensive technical documentation regarding this service.

- Developer Guide: Provides a conceptual overview of AWS Lake Formation, a tutorial for setting up a data lake, and an API reference for developers.

154. AWS Lambda

154.1. Service Overview

AWS Lambda is a serverless compute service that runs your code in response to events and automatically manages the underlying compute resources for you. These events may include changes in state or an update, such as a user placing an item in a shopping cart on an ecommerce website. You can use AWS Lambda to extend other AWS services with custom logic, or create your own backend services that operate at AWS scale, performance, and security. AWS Lambda automatically runs code in response to multiple events, such as HTTP requests via Amazon API Gateway, modifications to objects in Amazon Simple Storage Service (Amazon S3) buckets, table updates in Amazon DynamoDB, and state transitions in AWS Step Functions.

Lambda runs your code on high availability compute infrastructure and performs all the administration of your compute resources. This includes server and operating system maintenance, capacity provisioning and automatic scaling, code and security patch deployment, and code monitoring and logging. All you need to do is supply the code.

154.1.1. Features

- **Extend other AWS services with custom logic:** AWS Lambda allows you to add custom logic to AWS resources such as Amazon S3 buckets and Amazon DynamoDB tables, so you can easily apply compute to data as it enters or moves through the cloud.
- **Build custom backend services:** You can use AWS Lambda to create new backend application services triggered on demand using the Lambda application programming interface (API) or custom API endpoints built using Amazon API Gateway. Lambda

processes custom events instead of servicing these on the client, helping you avoid client platform variations, reduce battery drain, and enable easier updates.

- **Bring your own code:** With AWS Lambda, there are no new languages, tools, or frameworks to learn. You can use any third- party library, even native ones. You can also package any code (frameworks, SDKs, libraries, and more) as a Lambda Layer, and manage and share them easily across multiple functions. Lambda natively supports Java, Go, PowerShell, Node.js, C#, Python, and Ruby code, and provides a Runtime API allowing you to use any additional programming languages to author your functions.
- **Completely automated administration:** AWS Lambda manages all the infrastructure to run your code on highly available, fault tolerant infrastructure, freeing you to focus on building differentiated backend services. With Lambda, you never have to update the underlying operating system (OS) when a patch is released, or worry about resizing or adding new servers as your usage grows. AWS Lambda seamlessly deploys your code, handles all the administration, maintenance, and security patches, and provides built-in logging and monitoring through Amazon CloudWatch.
- **Built-in fault tolerance:** AWS Lambda maintains compute capacity across multiple Availability Zones (AZs) in each AWS Region to help protect your code against individual machine or data centre facility failures. Both AWS Lambda and the functions running on the service deliver predictable and reliable operational performance. AWS Lambda is designed to provide high availability for both the service itself and the functions it operates. There are no maintenance windows or scheduled downtimes.
- **Package and deploy functions as container images:** AWS Lambda supports function packaging and deployment as container images, making it easy for customers to build Lambda-based applications using familiar container image tooling, workflows, and dependencies. Customers also benefit from Lambda's operational simplicity, automatic scaling with sub-second startup times, high availability, pay-for-use billing model, and native integrations with over 200 AWS services and software-as-a service (SaaS) applications.
- **Connect to relational databases:** Use Amazon RDS Proxy to take advantage of fully managed connection pools for relational databases. RDS Proxy efficiently manages thousands of concurrent connections to relational databases, making it easy to build highly scalable, secure Lambda-based serverless applications interacting with relational databases. Currently, RDS Proxy offers support for MySQL and Aurora. You can use RDS Proxy for your serverless applications through the Amazon RDS console or AWS Lambda console.
- **Fine-grained control over performance:** Provisioned Concurrency gives you greater control over your serverless application performance. When turned on, Provisioned Concurrency keeps functions initialized and hyper-ready to respond in double-digit milliseconds. Provisioned Concurrency is ideal for any AWS Lambda application requiring greater control over function start time.
- **Run code in response to Amazon CloudFront requests:** With Lambda@Edge, AWS Lambda can run your code across AWS locations globally in response to Amazon CloudFront events, such as content requests to or from origin servers and viewers. This makes it easier to deliver richer, more personalized content to your end users with lower latency. Learn more »

- **Orchestrate multiple functions:** Build AWS Step Functions workflows to coordinate multiple AWS Lambda functions for complex or long-running tasks. Step Functions lets you define workflows that trigger a collection of Lambda functions using sequential, parallel, branching, and error-handling steps. With Step Functions and Lambda, you can build stateful, long-running processes for applications and backends.

154.1.2. Benefits

- **Integrate Lambda with your favourite operational tools:** AWS Lambda extensions enable easy integration with your favourite monitoring, observability, security, and governance tools. Lambda invokes your function in an execution environment, which provides a secure and isolated runtime where your function code is executed. Lambda extensions run within Lambda's execution environment, alongside your function code. With Lambda extensions, you can capture fine grained diagnostic information and send function logs, metrics, and traces to a location of your choice. You can also integrate security agents within Lambda's execution environment, all with no operational overhead and minimal impact to your function performance.
- **Connect to shared file systems:** With Amazon Elastic File System (EFS) for AWS Lambda, you can securely read, write, and persist large volumes of data at low latency, at any scale. You don't need to write code and download data to temporary storage in order to process it. This saves time and simplifies the code, so you can focus on your business logic. EFS for Lambda is ideal for a range of use cases including processing or backing up large data amounts, and loading large reference files or models. You can also share files between serverless instances or container-based applications, and even run machine learning (ML) inference by using EFS for AWS Lambda
- **Achieve better price performance with functions powered by Graviton2:** AWS Lambda functions running on Graviton2, using an Arm-based processor architecture designed by AWS, deliver up to 34% better price performance compared to functions running on x86 processors. This applies to a variety of serverless workloads, such as web and mobile backends, data, and media processing. With lower latency, up to 19% better performance, a 20% lower cost, and the highest power-efficiency currently available at AWS, Graviton2 functions can be used to power mission critical serverless applications.
- **Automatic scaling:** AWS Lambda invokes your code only when needed, and automatically scales to support the rate of incoming requests without any manual configuration. There is no limit to the number of requests your code can handle. AWS Lambda typically starts running your code within milliseconds of an event. Since Lambda scales automatically, the performance remains consistently high as the event frequency increases. Since your code is stateless, Lambda can start as many instances as needed without lengthy deployment and configuration delays.
- **Integrated security model:** AWS Lambda's built-in software development kit (SDK) integrates with AWS Identity and Access Management (IAM) to ensure secure code access to other AWS services. AWS Lambda runs your code within an Amazon Virtual Private Cloud (VPC) by default. Optionally, you can configure AWS Lambda resource access behind your own VPC in order to leverage custom security groups and network access control lists. This provides secure Lambda function access to your resources within a VPC. AWS Lambda is SOC, HIPAA, PCI, and ISO-compliant. For the latest in Lambda certification and compliance readiness, please see the full services in scope.

- **Trust and integrity controls:** Code Signing for AWS Lambda allows you to verify that only unaltered code published by approved developers is deployed in your Lambda functions. You simply create digitally signed code artefacts and configure your Lambda functions to verify the signatures at deployment. This increases the speed and agility of your application development, even within large teams, while enforcing high security standards.
- **Only pay for what you use:** With AWS Lambda, you pay for execution duration rather than server unit. When using Lambda functions, you only pay for requests served and the compute time required to run your code. Billing is metered in increments of one millisecond, enabling easy and cost-effective automatic scaling from a few requests per day to thousands per second. With Provisioned Concurrency, you pay for the amount of concurrency you configure and the duration that you configure it. When Provisioned Concurrency is enabled and your function is executed, you also pay for requests and execution duration.
- **Flexible resource model:** Choose the amount of memory you want to allocate to your functions, and AWS Lambda allocates proportional CPU power, network bandwidth, and disk input/output (I/O).

154.2. Backup/Restore and Disaster Recovery

This service allows for backup and recovery. The service can back up code. Users control this via a manual or scheduled API call. Users schedule and recover backups through a web interface.

154.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

154.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/lambda/
- **Service quotas:** https://docs.aws.amazon.com/lambda/latest/dg/gettingstarted-limits.html
- **Service FAQs:** https://aws.amazon.com/lambda/faqs/

154.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/lambda/ and the following links for comprehensive technical documentation regarding this service.

- Developer Guide: Provides a conceptual overview of AWS Lambda, detailed instructions for using the various features, and a complete API reference for developers.
- Operator Guide: Provides opinionated guidance on creating, securing, and monitoring AWS Lambda-based applications. Learn about event-driven architectures, application design, debugging, and more.

155. AWS License Manager

155.1. Service Overview

AWS License Manager makes it easier to manage your software licenses from vendors such as Microsoft, SAP, Oracle, and IBM across AWS and on-premises environments. AWS License Manager lets administrators create customized licensing rules that mirror the terms of their licensing agreements. Administrators can use these rules to help prevent licensing violations, such as using more licenses than an agreement stipulates. Rules in AWS License Manager help prevent a licensing breach by stopping the instance from launching or by notifying administrators about the infringement. Administrators gain control and visibility of all their licenses with the AWS License Manager dashboard and reduce the risk of non-compliance, misreporting, and additional costs due to licensing overages. Independent software vendors (ISVs) can also use AWS License Manager to easily distribute and track licenses.

AWS License Manager also simplifies the management of your software licenses that require Amazon EC2 Dedicated Hosts. In AWS License Manager, administrators can specify their Dedicated Host management preferences for host allocation and host capacity utilization. Once set up, AWS License Manager takes care of these administrative tasks on your behalf, so that you can seamlessly launch instances just like you would launch an EC2 instance with AWS-provided licenses.

155.1.1. Features

- **Set license terms as rules:** With AWS License Manager, you can enable a centralized team in your organization to manage software licensing agreements and create rules. These rules can then be used across the organization to govern license usage.
- **License tracking enforcement:** When a new EC2 instance gets launched, the rules created with AWS License Manager are attached using the console, CLI, or API. Once rules are attached, end users in your organization can launch instances and these can be tracked from dashboards in the AWS License Manager console. Licenses and usage can be tracked throughout the lifecycle of an instance. AWS License Manager also tracks any violation of the licensing rules and proactively sends an alert to end users and license administrators. When an instance is stopped or terminated, licenses are released and are available for reuse.
- **Limit non-compliance proactively:** Set hard or soft limits to control license usage and prevent the launch of a new, non-compliant instance. A hard limit blocks the launch of an out-of-compliance instance. A soft limit permits out-of-compliance launches but sends an alert when one occurs. AWS License Manager evaluates these limits during instance launches or while attaching licensing rules to existing instances. When license usage exceeds soft limits, AWS License Manager sends notifications to license administrators and end users with Amazon Simple Notification Service. These notifications can be in the form of emails, text messages, or alerts to inform administrators that an instance is non-compliant. For hard limits, AWS License Manager blocks new instances from being launched using AWS License Manager’s built-in integration with EC2.
- **Automate discovery of existing licenses:** AWS License Manager provides a mechanism to automatically discover software running on existing EC2 instances using AWS Systems Manager. Rules can then be attached and validated in EC2 instances allowing the licenses to be tracked using AWS License Manager’s central dashboard.

- **Switch Licenses easily:** Easily change license type of EC2 instances for Microsoft Windows Server and SQL Server workloads. License switching allows customers to switch between AWS provided licenses (license included) to bring-your-own-license (BYOL) with their own licensed media and take advantage of their existing investments, or switch from BYOL to purchasing license included from AWS to benefit from a flexible pay-as-you go licensing model. Customers can easily change the license type associated with their instance to make the switch while retaining the application, instance, and networking configuration associated with the workload.
- **Centralize license management and reporting:** Get a centralized view of license usage across AWS and on-premises environments based on your licensing rules. This makes it easy to manage incremental licensing purchases, compliance, and vendor audits across your organization. By sharing licensing rules across AWS accounts, a single team can be made responsible for creating, modifying, and deleting licensing rules centrally in one AWS account.
- **Automate management tasks for licenses requiring Dedicated Hosts:** To simplify management of licenses that require Dedicated Hosts, AWS License Manager allows administrators to specify Dedicated Host management preferences for host allocation and host capacity utilization. AWS License Manager then takes care of these tasks on your behalf, so that developers can seamlessly launch instances without performing upfront host allocation or managing capacity utilization.
- **Use managed entitlements to track licenses across multiple organizations:** AWS License Manager managed entitlements enable administrators to distribute, activate, and track third-party software procured in AWS Marketplace across multiple AWS accounts for end users and workloads. AWS License Manager managed entitlements also provides built-in controls so that independent software vendors (ISVs) and administrators can assign licenses to approved users and workloads.
- **Built-in AWS integration:** Seamlessly track license usage throughout the lifecycle of your AWS resources. AWS License Manager is integrated with Amazon EC2, AWS Systems Manager, AWS Organizations, AWS Service Catalog, and AWS Marketplace. License administrators can add rules in AWS Service Catalog, which allows them to create and manage catalogues of IT services that are approved for use on all their AWS accounts. Through seamless integration with AWS Systems Manager and AWS Organizations, administrators can manage licenses across all the AWS accounts in an organization and on-premises environments. AWS Marketplace buyers can also use AWS License Manager to track bring your own license (BYOL) software obtained from AWS Marketplace and keep a consolidated view of all their licenses.
- **Leverage dashboard to track usage:** With AWS License Manager, you can track licenses used across AWS and on-premises environments from a central dashboard. You can easily view license allocations, consumption, and alerts that need your action. This makes it easy for you to manage licensing purchases, compliance, and vendor audits.

155.1.2. Benefits

- **Gain control over license usage:** The way organizations manage licenses can vary from using simple spreadsheets to highly customized solutions. Often, these approaches require manual and ad-hoc reporting that can be inaccurate and quickly outdated. With AWS License Manager, administrators can create custom licensing rules, provision, and

track licenses across multiple accounts on AWS and on-premises environments. AWS License Manager centralizes license usage, providing organizations with greater visibility and control over how software licenses are used and can prevent misuse before it happens.

- **Reduce costs:** AWS License Manager provides a centralized view of license usage, so that administrators can determine the right number of licenses required, and not purchase more licenses than needed. With this improved visibility, you can also control overages and avoid penalties from licensing audits. AWS License Manager is easy to use, and helps reduce the time and cost for tracking and managing licenses.
- **Reduce the risk of non-compliance:** AWS License Manager gives administrators the ability to set limits for license usage. When license usage exceeds these limits, AWS License Manager sends an alert to administrators. Administrators also have the option to enforce these limits, and block the launch of new instances that require additional licenses. AWS License Manager also helps reduce the risk of non-compliance by providing independent software vendors (ISV) with a centralized AWS account and built-in controls to ensure only approved users and workloads can consume licenses.

155.2. Backup/Restore and Disaster Recovery

This requirement is not applicable for this service. For additional information beyond what is described herein, please refer to http://aws.amazon.com/documentation/.

155.3. Pricing Overview

There is no additional charge for AWS License Manager. You pay for AWS resources (e.g. EC2 instances) you create to run your application. You only pay for what you use, as you use it; there are no minimum fees and no upfront commitments.

155.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/license-manager/
- **Service quotas:** https://docs.aws.amazon.com/general/latest/gr/licensemanager.html
- **Service FAQs:** https://aws.amazon.com/license-manager/faqs/

155.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/license-manager/ and the following links for comprehensive technical documentation regarding this service.

- User Guide: Describes key concepts of AWS License Manager and provides instructions for using the features of AWS License Manager.

156. AWS Marketplace – BYOL

This is the service definition for the AWS Marketplace Bring Your Own License (BYOL) service offering as provided by Amazon Web Services EMEA SARL, UK Branch (AWS) in the G-Cloud 13 framework catalogue.

This AWS Marketplace BYOL Service Definition describes the key features for the AWS Marketplace BYOL Service available to Customers on G-Cloud 13. The underlying physical infrastructure that Customers provision through an AWS Marketplace BYOL Service will be

provided by AWS in accordance with the corresponding Digital Marketplace Service ID listing. Notwithstanding that AWS has combined its service descriptions into a consolidated set of documents for ease of review by Customers, in order to receive the AWS cloud compute infrastructure 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 is covered under their G-Cloud 13 Call-Off Contract.

AWS would recommend that Buyers list all AWS Digital Marketplace Service ID's in its Call-Off Contract to cover the various underlying cloud compute infrastructure services that may be provisioned through the Customers use of the AWS Marketplace BYOL Service.

156.1. Service Overview - AWS Marketplace

156.1.1. AWS Marketplace General Overview

AWS Marketplace is a curated digital catalogue that enables customers to find, deploy, and manage third-party software, AI/ML models and algorithms and the data sets they need to build solutions and run their organisations. AWS Marketplace includes products under popular categories such as security, networking, storage, machine learning, business intelligence, database, and DevOps. Products can take many forms, such as an Amazon Machine Image (AMI) that is instantiated using a customer's AWS account. Others can be configured to use AWS CloudFormation templates for delivery to the customer or as software as a service (SaaS).

156.1.2. AWS Marketplace (BYOL) on the G-Cloud 13 Framework

This Service Definition applies to the 'Bring Your Own Licence' (BYOL) service that is available through the AWS Marketplace under the G-Cloud 13 Framework.

156.2. Licence BYOL on the AWS Marketplace

The BYOL Service is designed to help Buyers migrate on-premises workloads to AWS, leveraging software for which the Buyer has already purchased a licence, but wishes to now run it in the cloud. The BYOL Service comprises of the following components provided by AWS:

- Over 600 BYOL products for fast provisioning and metered deployment of products such as security, network, storage, BI, and database, to name a few. Delivery methods for deployment of the software include Amazon Machine Image (AMI), CloudFormation and Containers. Details of which delivery method is available for which BYOL offering, is explained on the relevant BYOL product page on AWS Marketplace. Buyers who choose the BYOL Service need to have an active and valid AWS account in order to activate the license.
- Cloud compute infrastructure provided by AWS in accordance with the relevant Service Definition and Pricing Document listed on the Digital Marketplace. The Buyer activates the AWS infrastructure it requires to use their BYOL software.

156.3. Buyer Responsibilities

The Buyer is responsible (and AWS has no responsibility) for the following elements of the BYOL Service:

- Procuring its own software licence from the third-party provider in a separate contractual relationship directly between the Buyer and the third-party-provider.
- Managing its own compliance with applicable public procurement regulations to ensure any software activated via the AWS Marketplace has been lawfully procured.

- Managing its own ongoing compliance to the software licence terms of the third-party provider.
- Ensuring the security and integrity of its own BYOL software licence keys.
- Triage and resolution of issues with the third-party provider's software directly with the provider.
- Activating appropriate AWS infrastructure to enable the BYOL Service to operate efficiently (though AWS can provide Buyers with assistance and support with this where requested).
- Selecting a software service that is listed from a G-Cloud 13 third-party provider if the Buyer wishes to achieve the same legal benefits (i.e. terms and conditions).

The recommended process for the Buyer is to begin with a search and confirmation of the availability of a BYOL version of their existing application in AWS Marketplace. Once confirmed the customer simply selects the "Continue to Subscribe" button in the product page to proceed. The customer will be required to supply a license key to activate and use the product. The license key can be obtained directly from the third-party provider, an approved reseller or distributor.

The third-party provider verifies entitlement and handles licensing enforcement. AWS Marketplace does not charge customers for usage of the software however, customers are responsible for any AWS infrastructure usage fees as identified in the pricing section of product detail page).

156.4. Pricing Overview

Buyers are charged for the underlying AWS cloud computing infrastructure used to deploy and run the BYOL Service. Please see the AWS UK G-Cloud 13 Pricing Document affiliated with this service in the Digital Marketplace.

Any other fees, including for procuring, activating and running the software from the third-party provider, are the responsibility of the Buyer to manage and account for separately in accordance with the terms of its separate licence agreement with the third-party provider. AWS does not monitor and has no responsibility or liability for the ongoing compliance of the Buyer to the terms of the third-party-provider.

AWS Marketplace also provides the Buyer with the Private Offer feature which enables them the ability to further negotiate with the third-party provider on favourable terms for the BYOL software license the Buyer currently possesses. This process requires a signed Custom Transaction Request (CTF) form between the Buyer and third-party provider prior to agreeing to applicable price, terms and payment schedule.

156.5. Service Constraints

Other products existing in the AWS Marketplace are NOT covered under G-Cloud 13 terms and are not included as part of this service offering.

156.6. Service Support

Support for all third-party software products is provided by the third-party provider via the direct relationship between the Buyer and the third-party-provider. For information only, AWS lists general support instructions and contact information on the third-party provider's details page in AWS Marketplace, which instructs a Buyer how to contact, register and access general support.

However, the Buyer is responsible (and AWS has no responsibility) for contacting the third-party provider through the procurement and licensing arrangements the Buyer has made directly with the third-party provider.

The sole responsibility of AWS for support will be for support of the AWS Marketplace front end and access to the cloud computing infrastructure which AWS provides. Whilst AWS may provide support to third-party providers with regard to their listed offerings, third-party providers are responsible for the pre-configuration and optimisation of products, and the Buyer recognises that AWS may require the third-party provider to provide support where the issue relates to the pre-configuration of their product.

156.7. On-boarding through the AWS Marketplace

In order to deploy a third-party provider's software out of AWS Marketplace as a BYOL Service, a Buyer simply needs to have an active and valid AWS Account. Additional information on how to use AWS Marketplace can be accessed here.

157. AWS Migration Hub

157.1. Service Overview

AWS Migration Hub is the one destination for cloud migration and modernization, giving you the tools, you need to accelerate and simplify your journey with AWS. Perhaps you're making the case for cloud within your organization, or creating a data-driven inventory of existing IT assets. Maybe you're planning, running, and tracking a portfolio of applications migrating to AWS. Or you might be modernizing applications already running on AWS. In all of these cases, Migration Hub can help with your cloud transformation journey.

AWS Migration Hub provides a single place to store IT asset inventory data while tracking migrations to any AWS Region. After migration, use Migration Hub to accelerate the transformation of your applications to native AWS.

157.1.1. Features

- **Import or discover your on-premises server details:** With AWS Migration Hub, you can import information about on-premises servers and applications, or you can perform a deeper discovery using AWS Discovery Agent or AWS Discovery Collector, an agentless approach for VMware environments.
- **Build a migration plan:** AWS Migration Hub network visualization allows you to accelerate migration planning by quickly identifying servers and their dependencies, identifying the role of a server, and grouping servers into applications. To use network visualization, first install Discovery Agents and then start data collection from the Data Collectors page.
- **Strategy recommendations:** AWS Migration Hub Strategy Recommendations helps you easily build a migration and modernization strategy for your applications running on premises or in AWS. Strategy Recommendations analyses your applications to help you determine the optimal strategy and tools to migrate and modernize at scale.
- **Simple and intuitive migration dashboard:** The AWS Migration Hub dashboard shows the latest status and metrics for your rehost and replatform migrations. This allows you to quickly understand the progress of your migrations, as well as identify and troubleshoot any issues that arise. Migration Hub lets you track the status of your migrations into any AWS Region supported by your migration tools. Regardless of which

Regions you migrate into, the migration status will appear in Migration Hub when using an integrated tool.

- **Incremental app refactoring:** AWS Migration Hub Refactor Spaces (Preview) is the starting point for incremental application refactoring to microservices. Refactor Spaces eliminates the undifferentiated work of building and operating AWS infrastructure for incremental refactoring. You can use Refactor Spaces to reduce the risk of evolving applications into microservices or extending existing applications with new features written in microservices. The Refactor Spaces environment bridges networking across AWS accounts to permit old and new services to communicate while maintaining the independence of separate accounts. Refactor Spaces provides an application proxy that models the strangler-fig pattern to let you transparently add new services to an external HTTPS endpoint and incrementally route traffic to the new services. This keeps underlying architecture changes transparent to your app consumers.
- **Multi-region migrations:** AWS Migration Hub lets you track the status of your migrations into any AWS region supported by your migration tools. Regardless of which regions you migrate into, the migration status will appear in Migration Hub when using an integrated tool.

157.1.2. Benefits

- **Centralized tracking:** Migrations involve many components that need to be tracked such as the status of servers or databases being migrated, and these are typically tracked across different tools. AWS Migration Hub helps address this by providing a central location to track the status of all these components, making it easier to view overall migration progress and reducing the time spent determining current status and next steps.
- **Migration flexibility:** AWS Migration Hub provides the flexibility to use the migration tools that work best for your organization. Whether you use AWS Application Migration Service, AWS Database Migration Service, or partner tools such as ATADATA, Migration Hub makes it easy for you to track migrations across multiple migration tools.
- **Discovery, assessment, and planning:** AWS Migration Hub assists in all phases of migration and modernization readiness. It simplifies the discovery of existing applications and infrastructure and their dependencies, assessment of an application's ability to be migrated and modernized, and recommendations for modernization strategies.
- **Fast-track application refactor:** AWS Migration Hub helps you fast-track refactoring applications, simplify development and operations, and manage existing apps and microservices as a single application. Migration Hub eliminates the undifferentiated work of building and operating AWS infrastructure for incremental refactoring. It reduces the business risk of evolving applications into microservices or extending existing applications that can't be modified with new features written in microservices.

157.2. Backup/Restore and Disaster Recovery

This requirement is not applicable for this service. For additional information beyond what is described herein, please refer to http://aws.amazon.com/documentation/.

157.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

157.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/migrationhub/
- **Service quotas:** https://docs.aws.amazon.com/migrationhub/latest/ug/limits.html
- **Service FAQs:** https://aws.amazon.com/migration-hub/faqs/

157.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/migrationhub/ and the following links for comprehensive technical documentation regarding this service.

- Migration Hub User Guide: Walks through how to set up Migration Hub and integrate it with other services and includes the API reference.

158. AWS Network Firewall

158.1. Service Overview

AWS Network Firewall is a managed service that makes it easy to deploy essential network protections for all of your Amazon Virtual Private Clouds (VPCs). The service can be setup with just a few clicks and scales automatically with your network traffic, so you don't have to worry about deploying and managing any infrastructure. AWS Network Firewall's flexible rules engine lets you define firewall rules that give you fine-grained control over network traffic, such as blocking outbound Server Message Block (SMB) requests to prevent the spread of malicious activity. You can also import rules you've already written in common open source rule formats as well as enable integrations with managed intelligence feeds sourced by AWS partners. AWS Network Firewall works together with AWS Firewall Manager so you can build policies based on AWS Network Firewall rules and then centrally apply those policies across your VPCs and accounts.

AWS Network Firewall includes features that provide protections from common network threats. AWS Network Firewall's stateful firewall can incorporate context from traffic flows, like tracking connections and protocol identification, to enforce policies such as preventing your VPCs from accessing domains using an unauthorized protocol. AWS Network Firewall's intrusion prevention system (IPS) provides active traffic flow inspection so you can identify and block vulnerability exploits using signature-based detection. AWS Network Firewall also offers web filtering that can stop traffic to known bad URLs and monitor fully qualified domain names.

158.1.1. Features

- **High availability and automated scaling:** AWS Network Firewall offers built-in redundancies to ensure all traffic is consistently inspected and monitored. AWS Network Firewall offers a Service Level Agreement with an uptime commitment of 99.99%. AWS Network Firewall enables you to automatically scale your firewall capacity up or down based on the traffic load to maintain steady, predictable performance to minimize costs.
- **Stateful firewall:** The stateful firewall takes into account the context of traffic flows for more granular policy enforcement, such as dropping packets based on the source address or protocol type. The match criteria for this stateful firewall is the same as AWS Network Firewall's stateless inspection capabilities, with the addition of a match setting for traffic direction. AWS Network Firewall's flexible rule engine gives you the ability to write thousands of firewall rules based on source/destination IP, source/destination port,

and protocol. AWS Network Firewall will filter common protocols without any port specification, not just TCP/UDP traffic filtering.

- **Web filtering:** AWS Network Firewall supports inbound and outbound web filtering for unencrypted web traffic. For encrypted web traffic, Server Name Indication (SNI) is used for blocking access to specific sites. SNI is an extension to Transport Layer Security (TLS) that remains unencrypted in the traffic flow and indicates the destination hostname a client is attempting to access over HTTPS. In addition, AWS Network Firewall can filter fully qualified domain names (FQDN).
- **Intrusion prevention:** AWS Network Firewall's intrusion prevention system (IPS) provides active traffic flow inspection with real-time network and application layer protections against vulnerability exploits and brute force attacks. Its signature-based detection engine matches network traffic patterns to known threat signatures based on attributes such as byte sequences or packet anomalies.
- **Alert and flow logs:** Alert logs are rule specific and provide additional data regarding the rule that was triggered and the particular session that triggered it. Flow logs provide state information about all traffic flows that pass through the firewall, with one line per direction. AWS Network Firewall flow logs can be natively stored in Amazon S3, Amazon Kinesis, and Amazon CloudWatch.
- **Central management and visibility:** AWS Firewall Manager is a security management service that enables you to centrally deploy and manage security policies across your applications, VPCs, and accounts in AWS Organizations. AWS Firewall Manager can organize AWS Network Firewall rules groups into policies that you can deploy across your infrastructure to help you scale enforcement in a consistent, hierarchical manner. AWS Firewall Manager provides an aggregated view of policy compliance across accounts and automates the remediation process. As new accounts, resources, and network components are created, Firewall Manager makes it easy to bring them into compliance by enforcing a common set of firewall policies.
- **Rule management and customization:** AWS Network Firewall enables customers to run Suricata-compatible rules sourced internally, from in-house custom rule development or externally, from third party vendors or open source platforms.

158.1.2. Benefits

- **Managed infrastructure for high availability:** AWS Network Firewall infrastructure is managed by AWS. AWS Network Firewall automatically scales with your network traffic and can support hundreds of thousands of connections, so you don't have to worry about building and maintaining your own network security infrastructure.
- **Flexible protection through fine-grained controls:** AWS Network Firewall has a highly flexible rules engine that supports thousands of custom rules, so you can define firewall rules to protect your unique workloads. AWS Network Firewall rules can be based on IP, port, protocol, domain, and pattern matching and are written in common open source rule formats.
- **Consistent policy management across VPCs and accounts:** AWS Network Firewall works with AWS Firewall Manager so you can centrally manage security policies across existing accounts and VPCs. With AWS Firewall Manager, you can also ensure mandatory security policies are automatically enforced on newly created accounts and

VPCs. AWS Network Firewall provides real-time firewall activity monitoring through Amazon CloudWatch metrics.

158.2. Backup/Restore and Disaster Recovery

This requirement is not applicable for this service. For additional information beyond what is described herein, please refer to http://aws.amazon.com/documentation/.

158.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

158.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/network-firewall/
- **Service quotas:** https://docs.aws.amazon.com/network-firewall/latest/developerguide/quotas.html
- **Service FAQs:** https://aws.amazon.com/network-firewall/faqs/

158.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/network-firewall/ and the following links for comprehensive technical documentation regarding this service.

- Developer Guide: Describes how to use AWS Network Firewall, the Amazon Virtual Private Cloud firewall service.

159. AWS Nimble Studio

159.1. Service Overview

Amazon Nimble Studio empowers creative studios to produce visual effects, animation, and interactive content entirely in the cloud, from storyboard sketch to final deliverable. Rapidly onboard and collaborate with artists globally and create content faster with access to virtual workstations, high-speed storage, and scalable rendering across AWS's global infrastructure.

159.1.1. Features

- **Virtual Workstations:** Nimble Studio includes on-demand virtual workstations via the NICE DCV remote display protocol. This allows a browser or local client, such as a laptop to meet high-fidelity content creation's workflow demands. You can accommodate the way your artists work, such as using hotkeys or pen tablets, and scale workstation instance types based on the complexity of artistic tasks. Artists can leverage instances such as G4dn.xlarge (4 vCPUs, 16GBmemory and a NVIDIA Tesla T4 GPU with RTX) for lighter tasks and scale to instances with 64 vCPUs and 256GB of memory which can handle large data sets and simulation workflows.
- **High-Speed Storage:** Nimble Studio customers can access high-speed storage, including Amazon FSx, which provides cost-effective, high-performance, and scalable storage for content creation. You can stream both Linux and Windows while connected to the same storage for mixed OS production environments. Additionally, you can bring your own cloud-based storage solution, such as Qumulo or WekaIO.

- **Scalable Render Farm:** Scale your rendering workloads to tens of thousands of cores in minutes. Utilizing the Render Farm Deployment Kit (RFDK) and AWS Thinkbox Deadline for render orchestration, Nimble Studio provides an integrated render farm that leveraging EC2 Spot pricing. You can also scale down just as quickly, providing incredible compute elasticity and cost control.
- **StudioBuilder:** Nimble Studio's StudioBuilder helps you create a virtual studio from scratch. Walk through and set up your studio by creating networking, render farm, and storage resources. The StudioBuilder process creates and deploys new resources into your region of choice, helping you build a studio in just a few hours. You can also re-run StudioBuilder to add additional resources.
- **Launch Profile:** Launch Profiles let you share necessary studio resources with artists to accomplish production tasks. Administrators can define and share the virtual workstations, necessary software, file system access, and the render farm. Launch Profiles uses a simple sharing mechanism to simplify resource assignment to your artists.
- **Artist-Friendly Portal UI:** The Nimble Studio portal user interface (UI) simplifies artist onboarding, so they can focus more on creative output. Artists simply login into the Nimble Studio portal UI and launch the virtual workstation needed to complete tasks. Studio administrators can easily share Launch Profiles with artists based on the characteristics of the projects they are currently working on.
- **API:** Take advantage of a full set of APIs to integrate proprietary tools into your workflow. Use the Nimble Studio API to extend your studio's functionality and leverage additional AWS services.
- **User Management:** Nimble Studio uses AWS Single Sign-On (SSO) to provide secure artist access to web identities in the Nimble Studio portal. Nimble Studio portal includes workstation and file system access control via AWS Managed Microsoft Active Directory (AD), enabling directory-aware workloads for your production security needs. Nimble Studio gives account administrators and project owners the ability to share projects, add or remove artists, and control download access rights to sensitive data for added security on the (streaming) virtual workstations.

159.1.2. Benefits

- **Accelerate your cloud transition:** Get your content production pipeline up and running in hours instead of weeks. Nimble Studio's automation and pre-built Amazon Machine Images (AMIs) make it easy to set up your virtual workstations, storage, and render farm, all with an artist-friendly user interface (UI).
- **Scale with project demand:** Nimble Studio automatically configures AWS services, scaling your studio to suit business needs across single or multiple locations. Add more artists to graphics-intensive projects with virtual workstations, utilize high-speed storage with Amazon FSx, and orchestrate compute resources on an integrated cloud-based render farm using EC2 Spot Instances.
- **Access the global talent you need:** Onboard remote artists in minutes. Utilize the latest software and hardware technology, to provide the best possible performance for your artists and studio. With availability in major content creation markets, you can look for and hire the best talent.
- **Simplified workstation pricing:** Setting up virtual streaming workstations involves multiple cost considerations. Nimble Studio offers a simplified pricing structure including

the instance, Elastic Block Store (EBS), and egress charges to take the guesswork out of the total cost of ownership (TCO).

- **Seamless collaboration:** The Nimble Studio portal simplifies granting user permissions, sharing project data, and adding new team members. Using the NICE DCV remote display protocol, stream pixels instead of data to keep your project data in the cloud and to streamline artist collaboration.
- **Build with the highest standard for data security:** Build on the most secure global infrastructure, knowing you always own your data, including the ability to encrypt it, move it, and manage retention. AWS automatically encrypts all data flowing across the AWS global network at the physical layer before it leaves our secured facilities.

159.2. Backup/Restore and Disaster Recovery

This service allows for backup and recovery. The service can be backed up with FSx or S3. The service can back up studio data, which saves the content within the studio and visual animations etc. Users control this manually. Users choose what they backup and to where.

159.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

159.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/nimble-studio/latest/userguide/service-quotas.html
- **Service quotas:** https://docs.aws.amazon.com/nimble-studio/latest/userguide/service-quotas.html
- **Service FAQs:** https://aws.amazon.com/nimble-studio/faqs/?refid=ps_a134p000006gb2oaau&trkcampaign=acq_paid_search_brand

159.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/nimble-studio/latest/userguide/service-quotas.html and the following links for comprehensive technical documentation regarding this service.

- User Guide: Provides a conceptual overview of Amazon Nimble Studio and provides tutorials for using the various features with the console.
- API Reference: Describes all the Nimble Studio API operations in detail. Also provides sample requests, responses, and errors for the supported web service protocols.

160. AWS OpsWorks for Chef Automate

160.1. Service Overview

AWS OpsWorks for Chef Automate provides a fully managed Chef Automate server and suite of automation tools that give you workflow automation for continuous deployment, automated testing for compliance and security, and a user interface that gives you visibility into your nodes and their status. The Chef Automate platform gives you full stack automation by handling operational tasks such as software and operating system configurations, continuous compliance, package installations, database setups, and more. The Chef server centrally stores

your configuration tasks and provides them to each node in your compute environment at any scale, from a few nodes to thousands of nodes. OpsWorks for Chef Automate is completely compatible with tooling and cookbooks from the Chef community and automatically registers new nodes with your Chef server.

160.1.1. Features

Copy and paste features from service features page

- **Chef Automation:** Chef is an automation patform that helps you automate operational tasks at scale. You can use Chef to manage both Amazon Elastic Compute Cloud (Amazon EC2) instances and on-premises servers running Linux or Windows.With Chef, you use code templates, or cookbooks, to describe the desired configuration of instances or on-premises servers. Cookbooks contain recipes that describe the desired state for a configuration item and the steps needed to reach that state, server settings, information on how to distribute files, and more. You can use cookbooks to automate operational tasks such as configuring hosts and applications, installing packages, shutting down instances, and more. You can author your own cookbooks or use over 3,000 publicly available cookbooks from the Chef community.
- **Premium Chef Features:** Chef Automate is an enterprise platform and analytics tool that allows development, operations and security engineers to collaborate with actionable insights for configuration and compliance and an auditable history of changes to environments. Chef Automate provides operational visibility for today's Coded Enterprise with:
    - Real-Time data across the estate
    - Effortless collaboration among teams
    - Powerful auditing capabilities
    - Intelligent access controls
    - Pre-built, supported compliance assets
- **Managed Chef Server:** AWS OpsWorks for Chef Automate provisions a managed Chef server running on an Amazon EC2 instance in your account. There is no need to provision or install the Chef server. At the same time, you retain control over the underlying resources running your Chef server and you can use Knife to SSH into your Chef server instance at any time.
- **Multiple Interface Options:** You can provision your Chef server using the AWS Management Console, AWS CLI, and SDKs. Once you have provisioned your Chef server, you can interface with it using Chef-native tools such as the ChefDK or Knife command-line tool.
- **Maintenance Windows:** AWS OpsWorks for Chef Automate handles security, operating system, and Chef minor version updates for you, helping you keep your Chef server up-to-date. You can set a weekly maintenance window during which OpsWorks for Chef Automate will automatically install updates. OpsWorks for Chef Automate also monitors the health of your Chef server during update windows and automatically rolls back changes if issues are detected.
- **Automated Backups:** You can configure automatic backups for your Chef server. AWS OpsWorks for Chef Automate lets you set the frequency of backups, when to perform them, and how many backups to keep. You can then restore from backups at any time using the AWS CLI. OpsWorks for Chef Automate stores Chef server backups in secure, durable Amazon S3 buckets in your AWS account.

- **Node Registration:** AWS OpsWorks for Chef Automate makes it easier to register new instances as Chef nodes. You can register new nodes to your Chef server by inserting user-data code snippets provided by OpsWorks for Chef Automate into your Auto Scaling groups.
- **Manage On-Premises Servers:** You can manage on-premises environments from your Chef server by installing the Chef agent on your on-premises servers.

160.1.2. Benefits

- **Fully Managed Chef Automate Server:** AWS OpsWorks for Chef Automate makes it easy to use Chef Automate on AWS. There is no need for Chef Automate server provisioning and installation. OpsWorks for Chef Automate automatically updates the Chef Automate software and creates backups of the Chef server for you. You can use the AWS Management Console, AWS Command Line Interface (CLI), or AWS SDKs to provision a ready-to-use Chef Automate and Chef server.
- **Programmable Infrastructure:** With Chef server, you can define configurations for your servers in a format that you can maintain and version just like your application source code. Chef server ensures your servers are consistently configured and maintained, and handles complex operational tasks when there are interdependencies across your servers.
- **Scaling Made Easy:** AWS OpsWorks for Chef Automate dynamically configures newly provisioned instances by automatically registering new instances in Auto Scaling groups with your Chef server. You can also choose the instance size and type that your Chef server runs on to meet the scale of your server fleet, from a few nodes to thousands of nodes.
- **Support from Active Chef Community:** AWS OpsWorks for Chef Automate supports the latest version of Chef server and Chef Automate. You can use any Chef community-built tools or cookbooks with your Chef server. With OpsWorks for Chef Automate, you can interface with your Chef server using native Chef tools, such as the Knife command-line tool or Chef Development Kit (ChefDK).
- **Secure:** AWS OpsWorks for Chef Automate runs your Chef server on an Amazon EC2 instance in an Amazon Virtual Private Cloud. This means you can control inbound and outbound network access to your Chef server. You can also use AWS Identity and Access Management (IAM) to set fine-grained access controls on which users and resources can access your Chef server instance.
- **Simple to Manage Hybrid Environments:** Chef server lets you seamlessly administer your Linux or Windows servers running on Amazon EC2 and on-premises from a single service. Once you have defined your server configurations, nodes across your hybrid environments will periodically converge to the desired configuration state.

160.2. Backup/Restore and Disaster Recovery

This service allows for backup and recovery. The service can back up server backups. Users control this via a manual or scheduled API call. Users schedule and recover backups through a web interface.

160.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

160.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/opsworks/index.html
- **Service quotas:** https://docs.aws.amazon.com/general/latest/gr/opsworks-service.html
- **Service FAQs:** https://aws.amazon.com/opsworks/chefautomate/faqs/?nc=sn&loc=5

160.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/opsworks/index.html and the following links for comprehensive technical documentation regarding this service.

- User Guide: Provides detailed descriptions of AWS OpsWorks Stacks, AWS OpsWorks for Chef Automate, and AWS OpsWorks for Puppet Enterprise concepts, and provides instructions for using both the console and the command line interface.
- AWS OpsWorks Stacks API Reference: Describes all the API operations for AWS OpsWorks Stacks in detail. In addition, it provides sample requests, responses, and errors for the supported web services protocols.
- AWS OpsWorks CM API Reference: Describes the API operations for AWS OpsWorks CM—including AWS OpsWorks for Chef Automate and AWS OpsWorks for Puppet Enterprise—in detail. In addition, it provides sample requests, responses, and errors for the supported web services protocols.

161. AWS OpsWorks for Puppet Enterprise

161.1. Service Overview

AWS OpsWorks for Puppet Enterprise is a fully managed configuration management service that hosts Puppet Enterprise, a set of automation tools from Puppet for infrastructure and application management. OpsWorks also maintains your Puppet master server by automatically patching, updating, and backing up your server. OpsWorks eliminates the need to operate your own configuration management systems or worry about maintaining its infrastructure. OpsWorks gives you access to all of the Puppet Enterprise features, which you manage through the Puppet console. It also works seamlessly with your existing Puppet code.

161.1.1. Features

- **Managed Puppet Master:** AWS OpsWorks for Puppet Enterprise provisions a managed Puppet master server running on an Amazon EC2 instance in your account. There is no need to provision or install the Puppet master. At the same time, you retain control over the underlying resources running your Puppet master.
- **Puppet Enterprise Ecosystem:** You can provision your Puppet master using the AWS Management Console, AWS CLI, and SDKs. Your Puppet master is preconfigured with CodeManager, which lets you use Git to develop your Puppet Code and deploy it to your master. You can also extend Puppet's capabilities by using open source modules available on Puppet Forge.
- **Maintenance Windows:** AWS OpsWorks for Puppet Enterprise handles security, operating system, and Puppet Enterprise software updates for you, helping you keep your Puppet master up-to-date. You can choose the weekly maintenance window during which OpsWorks for Puppet Enterprise will automatically install updates. OpsWorks for

Puppet Enterprise also monitors the health of your Puppet master during update windows and automatically rolls back changes if issues are detected.

- **Automated Backups:** You can configure automatic backups for your Puppet master. AWS OpsWorks for Puppet Enterprise lets you set the frequency of backups, when to perform them, and how many backups to keep. You can then restore from backups at any time using the AWS CLI. OpsWorks for Puppet Enterprise stores Puppet master backups in secure, durable Amazon S3 buckets in your AWS account.
- **Node Registration:** AWS OpsWorks for Puppet Enterprise makes it easier to register new instances as Puppet nodes. Puppet nodes are instances that run the Puppet agent and are automatically provisioned, configured, and managed by the Puppet master. You can register new nodes to your Puppet master by inserting a user-data script, provided in the OpsWorks for Puppet Enterprise StarterKit, into your Auto Scaling groups.
- **Manage On-Premises Servers:** You can manage on-premises environments from your Puppet master by installing the Puppet agent on your on-premises servers.
- **Security:** Puppet uses SSL and a certification approval process when communicating to ensure that the Puppet master responds only to requests made by trusted users. AWS OpsWorks for Puppet Enterprise is integrated with AWS Identity and Access Management allowing you to control which nodes can be registered with your Puppet master. Your Puppet master instance runs in Amazon Virtual Private Cloud, allowing you to configure network settings for subnets and security groups. You can also disable SSH access to your Puppet master instance for added security, or use Amazon EC2 Systems Manager Run Command as an alternative to SSH. OpsWorks for Puppet Enterprise is also integrated with AWS CloudTrail, allowing you to track and record a history of API calls made to the service.

161.1.2. Benefits

- **Fully Managed Puppet Master:** AWS OpsWorks for Puppet Enterprise will automatically patch, update, and backup your Puppet masters as well as maintain their availability through scheduled system maintenance. A Puppet master is a central server that provisions, configures, and manages Puppet nodes. You can automatically register new nodes through the API, and they are backed by AWS Identity and Access Management (IAM) instance profile permissions.
- **Programmable Infrastructure:** With Puppet Enterprise, you can define configurations for your servers in a format that you can maintain and version just like your application source code. The Puppet master ensures your servers are consistently configured and maintained. You can also configure your nodes dynamically based on the state of other nodes.
- **Scaling Made Easy:** You can set up Auto Scaling to automatically register and provision new nodes. Whenever you scale out, the launch configurations you set in Auto Scaling will associate the new node to your Puppet master. You can also choose the instance size and type that your Puppet master runs on to meet the scale of your server fleet, from a few nodes to thousands of nodes.
- **Support from Puppet Community:** Find modules for Open Source Puppet and Puppet Enterprise IT automation software on the Puppet Forge repository. Modules are the building blocks for Puppet, consisting of self-contained, reusable, and shareable units of Puppet code (Puppet's configuration language). You can use any Puppet community-built modules or manifests with your Puppet master.

- **Secure:** AWS OpsWorks for Puppet Enterprise runs your Puppet master on an Amazon EC2 instance in an Amazon Virtual Private Cloud. This means you can control inbound and outbound network access to your Puppet master. You can also use AWS Identity and Access Management (IAM) to set fine-grained access controls on which users and resources can access your Puppet master instance.
- **Simple to Manage Hybrid Environments:** Puppet master lets you seamlessly administer your Linux or Windows server nodes running on Amazon EC2 and on-premises. Once you have defined your server configurations, nodes across your hybrid environments will periodically request and check their state against a catalog from the Puppet master, and they will correct any resources that are not in the desired state.

161.2. Backup/Restore and Disaster Recovery

This service allows for backup and recovery. The service can back up server backups. Users control this via a manual or scheduled API call. Users schedule and recover backups through a web interface.

161.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

161.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/opsworks/index.html
- **Service quotas:** https://docs.aws.amazon.com/general/latest/gr/opsworks-service.html
- **Service FAQs:** https://aws.amazon.com/opsworks/chefautomate/faqs/?nc=sn&loc=5

161.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/opsworks/index.html and the following links for comprehensive technical documentation regarding this service.

- User Guide: Provides detailed descriptions of AWS OpsWorks Stacks, AWS OpsWorks for Chef Automate, and AWS OpsWorks for Puppet Enterprise concepts, and provides instructions for using both the console and the command line interface.
- AWS OpsWorks Stacks API Reference: Describes all the API operations for AWS OpsWorks Stacks in detail. In addition, it provides sample requests, responses, and errors for the supported web services protocols.
- AWS OpsWorks CM API Reference: Describes the API operations for AWS OpsWorks CM—including AWS OpsWorks for Chef Automate and AWS OpsWorks for Puppet Enterprise—in detail. In addition, it provides sample requests, responses, and errors for the supported web services protocols.

162. AWS Organizations

162.1. Service Overview

AWS Organizations helps you centrally manage and govern your environment as you grow and scale your AWS resources. Using AWS Organizations, you can programmatically create new AWS accounts and allocate resources, group accounts to organize your workflows, apply

policies to accounts or groups for governance, and simplify billing by using a single payment method for all of your accounts.

In addition, AWS Organizations is integrated with other AWS services so you can define central configurations, security mechanisms, audit requirements, and resource sharing across accounts in your organization. AWS Organizations is available to all AWS customers at no additional charge.

162.1.1. Features

- **Manage your AWS accounts:** AWS accounts are natural boundaries for permission, security, costs, and workloads. Using a multi-account environment is a recommended best-practice when scaling your cloud environment. You can simplify account creation by programmatically creating new accounts using the AWS Command Line Interface (CLI), SDKs, or APIs, and centrally provision recommended resources and permissions to those accounts with AWS CloudFormation StackSets.
- **Define and manage your organization:** As you create new accounts, you can group them into organizational units (OUs), or groups of accounts that serve a single application or service. Apply tag polices to classify or track resources in your organization, and provide attribute-based access control for users or applications. In addition, you can delegate responsibility for supported AWS services to accounts so users can manage them on behalf of your organization.
- **Secure and monitor your accounts:** You can centrally provide tools and access for your security team to manage security needs on behalf of the organization. For example, you can provide read-only security access across accounts, detect and mitigate threats with Amazon GuardDuty, review unintended access to resources with IAM Access Analyzer, and secure sensitive data with Amazon Macie.
- **Control access and permissions:** Set up Amazon Single Sign-On (SSO) to provide access to AWS accounts and resources using your active directory, and customize permissions based on separate job roles. You can also apply service control policies (SCPs) to users, accounts, or OUs to control access to AWS resources, services, and Regions within your organization.
- **Share resources across accounts:** You can share AWS resources within your organization using AWS Resource Allocation Management (RAM). For example, you can create your AWS Virtual Private Cloud (VPC) subnets once and share them across your organization. You can also centrally agree to software licenses with AWS License Manager, and share a catalog of IT services and custom products across accounts with AWS Service Catalog.
- **Audit your environment for compliance:** You can activate AWS CloudTrail across accounts, which creates a log of all activity in your cloud environment that cannot be turned off or modified by member accounts. In addition, you can set policies to enforce backups on your specified cadence with AWS Backup, or define recommended configuration settings for resources across accounts and AWS Regions with AWS Config.
- **Centrally manage billing and costs:** Organizations provides you with a single consolidated bill. In addition, you can view usage from resources across accounts and track costs using AWS Cost Explorer, and optimize your usage of compute resources using AWS Compute Optimizer.

162.1.2. Benefits

- **Quickly scale your workloads:** AWS Organizations helps you quickly scale your environment by allowing you to programmatically create new AWS accounts. An AWS account is a container for your resources. Using multiple accounts gives you built-in security boundaries. It also empowers your teams by providing them designated accounts, and you can automatically provision resources and permissions using AWS CloudFormation StackSets.
- **Provide custom environments for different workloads:** You can use Organizations to apply policies that give your teams the freedom to build with the resources they need, while staying within the safe boundaries you set. By organizing accounts into organizational units (OUs), which are groups of accounts that serve an application or service, you can apply service control policies (SCPs) to create targeted governance boundaries for your OUs.
- **Centrally secure and audit your environment across accounts:** Manage auditing at scale using AWS CloudTrail to create an immutable log of all events from accounts. You can enforce and monitor backup requirements with AWS Backup, or centrally define your recommended configuration criteria across resources, AWS Regions, and accounts with AWS Config. You can also use AWS Control Tower to establish cross-account security audits, or manage and view policies applied across accounts. In addition, you can protect your resources by centrally managing security services, such as detecting threats with Amazon GuardDuty, or reviewing unintended access with AWS IAM Access Analyzer.
- **Simplify permission management and access control:** Simplify user-based permission management for everyone in your organization with AWS Single Sign-On (SSO) and your Active Directory. You can apply least-privilege practices by creating custom permissions for job categories. You can also control access to AWS services by applying service control policies (SCPs) to users, accounts, or OUs.
- **Efficiently provision resources across accounts:** You can reduce resource duplication by sharing critical resources within your organization using AWS Resource Access Manager (RAM). Organizations also helps you meet your software license agreements with AWS License Manager, and maintain a catalog of IT services and custom products with AWS Service Catalog.
- **Manage costs and optimize usage:** AWS Organizations enables you to simplify costs and take advantage of quantity discounts with a single bill. In addition, you can optimize usage across your organization with services like AWS Compute Optimizer and AWS Cost Explorer.

162.2. Backup/Restore and Disaster Recovery

This requirement is not applicable for this service. For additional information beyond what is described herein, please refer to http://aws.amazon.com/documentation/.

162.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

162.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/organizations/index.html
- **Service quotas:** https://docs.aws.amazon.com/organizations/latest/userguide/orgs_reference_limits.html
- **Service FAQs:** https://aws.amazon.com/organizations/faqs/

162.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/organizations/index.html and the following links for comprehensive technical documentation regarding this service.

- User Guide: Introduces you to AWS Organizations, helps you set up an organization by creating accounts or inviting other accounts to join. It shows you how to manage your accounts as a group and provide organization-wide access control to your AWS resources by using policy-based controls.
- API Reference: Describes all the API operations for AWS Organizations in detail. Also provides sample requests, responses, and errors for the supported web services protocols.
- AWS Organizations section of AWS CLI Reference: Describes the AWS CLI commands that you can use to administer AWS Organizations. Provides syntax, options, and usage examples for each command.
- AWS Account Management Reference Guide: Introduces you to creating and managing your individual AWS accounts. It includes detailed descriptions of the API operations, including sample requests, responses, and errors for the supported web services protocols.

163. AWS Outposts

163.1. Service Overview

AWS Outposts is a family of fully managed solutions delivering AWS infrastructure and services to virtually any on-premises or edge location for a truly consistent hybrid experience. Outposts solutions allow you to extend and run native AWS services on premises, and is available in a variety of form factors, from 1U and 2U Outposts servers to 42U Outposts racks, and multiple rack deployments.

With AWS Outposts, you can run some AWS services locally and connect to a broad range of services available in the local AWS Region. Run applications and workloads on premises using familiar AWS services, tools, and APIs. Outposts supports workloads and devices requiring low latency access to on-premises systems, local data processing, helps meet data residency requirements, and can facilitate application migration with local system interdependencies.

163.1.1. Features

163.1.1.1. AWS Outposts rack

- **Compute and storage:** You can choose from a range of pre-validated Outposts rack configurations offering a mix of Amazon Elastic Compute Cloud (EC2), Amazon Elastic Block Store (EBS), and Amazon Simple Storage Service (S3) capacity designed to meet a variety of application and data residency needs. You can also contact AWS to create a customized configuration designed for your unique application needs.
- **Networking:**

- **VPC extension:** You can seamlessly extend your existing Amazon VPC to your Outpost in your on-premises location. After installation, you can create a subnet in your regional VPC and associate it with an Outpost just as you associate subnets with an Availability Zone in an AWS Region. Instances in Outpost subnets communicate with other instances in the AWS Region using private IP addresses, all within the same VPC.
- **Local gateway:** Each Outpost provides a local gateway service (LGW) that allows you to connect your Outpost resources with your on-premises networks. LGW helps enable low latency connectivity between the Outpost and any local data sources, end users, local machinery and equipment, or local databases.
- **Load Balancer:** You can provision an Application Load Balancer (ALB) to automatically distribute incoming HTTP(S) traffic across multiple targets on your Outposts rack, such as Amazon EC2 instances, containers, and IP addresses. ALB on Outposts is a managed service, operates in a single subnet, and scales automatically up to the capacity available on the Outposts rack to meet varying levels of application load without manual intervention.
- **Private Connectivity:** With AWS Outposts Private Connectivity, you can establish a service link VPN connection from your Outpost to the AWS Region over AWS Direct Connect. Private Connectivity minimizes public internet exposure and removes the need for special firewall configurations.

- AWS services on Outposts rack: You can run a variety of AWS services locally to build and run your applications on premises.
  - **Containers:** Amazon ECS is a highly scalable, high-performance container orchestration service that supports Docker containers and allows you to easily run and scale containerized applications on Outposts rack. With ECS you can run containerized applications that require low latency connectivity to on premises systems on Outposts or on local networks. Amazon EKS is a managed service that makes it easy for you to run Kubernetes on AWS without needing to install and operate your own Kubernetes control plane. You can use EKS to run containerized applications that require particularly low latencies to on premises systems on Outposts or on local networks. With EKS, you can manage containers on premises on Outposts with the same ease as you manage your containers in AWS Regions.
  - **Databases:** Amazon RDS on Outposts allows you to deploy fully managed database instances in your on-premises environments. You can deploy Amazon RDS on Outposts to set up, operate, and scale Microsoft SQL Server, MySQL and PostgreSQL relational databases on premises, just as you would in AWS Regions. Amazon RDS on Outposts provides cost-efficient and resizable capacity for on-premises databases, while automating time-consuming administration tasks including infrastructure provisioning, database setup, patching, and backups. When you deploy Amazon RDS on Outposts, you can run RDS on premises for low latency workloads that need to be run in close proximity to your on-premises data and applications. Amazon RDS on Outposts also enables automatic backup to your Outpost or to an AWS Region. You can manage RDS databases both in AWS Regions and on premises using the same AWS Management Console, APIs, and CLI. Amazon RDS Multi-AZ on Outposts can enhance availability by deploying a standby instance on a second Outpost

and uses synchronous replication technologies to keep data on your standby database instance up to date with the primary. Support for additional database engines are coming soon. Amazon ElastiCache on Outposts is a fully managed in-memory data store, compatible with Redis or Memcached, optimized for real-time applications with sub-millisecond latency. Amazon ElastiCache on Outposts allows you to seamlessly set up, run, and scale popular open-source compatible in-memory data stores on AWS Outposts rack capacities, as in the AWS Regions.

  - **Data analytics:** Amazon EMR clusters running on AWS Outposts rack in your data center, co-location space, or on-premises facility provide a truly consistent and seamless hybrid cloud analytics experience.
  - **Upgrading services running on Outposts rack:** As new versions of AWS services become available in the AWS Region, AWS services running locally on Outposts rack will be upgraded automatically to the latest version just as in the AWS Region today. Services such as Amazon RDS on Outposts patch both OS and database engines within scheduled maintenance windows with minimum downtime.
  - **Access regional services:** AWS Outposts rack is an extension of the AWS Region. You can seamlessly extend your Amazon Virtual Private Cloud on premises and connect to a broad range of services available in the AWS Region. You can access all regional AWS services in your private VPC environment — for example, through Interface Endpoints, Gateway Endpoints, or their regional public endpoints.

- **AWS tools:** You can access AWS tools running in the AWS Region such as AWS CloudFormation, Amazon CloudWatch, AWS CloudTrail, Elastic BeanStalk, Cloud 9, and others to run and manage applications on Outposts rack the same way as you do in the AWS Region today.
- **Security and compliance:**
  - **Enhanced security with AWS Nitro:** AWS Outposts rack builds on the AWS Nitro System technologies that enables AWS to provide enhanced security that continuously monitors, protects, and verifies your Outpost's instance hardware and firmware.
  - **Security model:** AWS Outposts rack has an updated shared responsibility model underlying security. AWS is responsible for protecting Outposts rack's infrastructure similar to how it secures infrastructure in the AWS Regions today. Customers are responsible for securing their applications running on Outposts rack as they do in the Region today. With Outposts rack, customers are also responsible for the physical security of their Outposts racks, and for ensuring consistent networking to the Outpost.
  - **Securing data:** Data-at-rest: Data is encrypted at rest by default on EBS volumes, and S3 objects on Outposts rack. Data-in-transit: Data is encrypted in transit between Outposts rack and the AWS Region, through the service link. Deleting data: All data is deleted when instances are terminated in the same way as in the AWS Region.
- **High availability:** Outposts rack is designed for high availability with redundant top of rack networking switches, power elements, and built-in, always active, additional

capacity (if requested by the customer) to enable reliable auto recovery workflows the same way as in AWS Regions. Similar to AWS Auto Scaling in the AWS Regions today, we recommend best practices for high availability deployments and auto recovery workflows for easy failover in case of any underlying host issue. Customers can also deploy multiple Outposts at a site, each tied to a different Availability Zone for even higher availability. In addition, customers can use EC2 placement groups on AWS Outposts rack to ensure instances within a group are placed on distinct Outposts racks to reduce the impact of hardware failures.

- **AWS Resource Access Manager:** AWS Outposts rack support for AWS Resource Access Manager (RAM) lets customers share access to Outposts rack resources – EC2 instances, EBS volumes, S3 capacity, subnets, and local gateways (LGWs) – across multiple accounts under the same AWS organization. This new capability allows distributed teams and business units in customer organizations to configure VPCs, launch and run instances, and create EBS volumes on the shared Outpost.

163.1.1.2. AWS Outposts servers

- **Compute:** AWS Outposts servers includes a 1U server that supports Arm-based AWS Graviton2 powered EC2 instances, and a 2U server that supports 3rd generation Intel Xeon Scalable powered EC2 instances.
- **Storage:** Outposts servers have up to 4x 1.9 TB raw NVMe SSD instance storage, supporting local storage used for data access and processing on premises, and for launching EBS-backed AMIs. When launching new instances on an Outposts server, storage is allocated as boot volumes, reducing the remaining storage available for data volumes.
- **Networking:**
    - **VPC extension:** You can seamlessly extend your existing Amazon Virtual Private Cloud (VPC) to your Outposts server in your on-premises location. After installation, you can create a subnet in your regional VPC and associate it with an Outpost just as you associate subnets with an Availability Zone in an AWS Region. Instances in Outpost subnets communicate with other instances in the AWS Region using private IP addresses, all within the same VPC.
    - **Local network interface (LNI):** Outposts servers have an LNI that provides a Layer 2 presence on your local network for AWS service endpoints.
- **AWS services on Outposts servers:**
    - **Locally supported AWS services:** You can run ECS, IoT Greengrass, or Sagemaker Edge Manager locally on Outposts servers, and connect to the AWS Region for a broad range of services available in the AWS Region. Support for Amazon Elastic Kubernetes Service (EKS) is coming soon.
    - **Access regional services:** AWS Outposts servers are an extension of the AWS Region. You can seamlessly extend your Amazon VPC on premises and connect to a broad range of services available in the AWS Region. You can access all regional AWS services in your private VPC environment — for example, through interface endpoints, gateway endpoints, or their regional public endpoints.
    - **AWS tools:** You can access AWS tools running in the Region—for example, AWS CloudFormation, Amazon CloudWatch, AWS CloudTrail, Amazon Elastic

Beanstalk, and AWS Cloud9—to run and manage applications on Outposts servers the same way as you do in the cloud today.

- **Security and compliance:**
    - **Enhanced security with the AWS Nitro System:** AWS Outposts servers are built on the Nitro System, which enables AWS to provide enhanced security that continuously monitors, protects, and verifies your Outpost's instance hardware and firmware. With AWS Nitro, virtualization resources are offloaded to dedicated hardware and software, minimizing the attack surface. The Nitro System's security model is locked down and prohibits administrative access, reducing the possibility of human error and tampering.
    - **Security model:** AWS Outposts servers have an updated shared responsibility model underlying security. AWS is responsible for protecting infrastructure for Outposts servers similar to how it secures infrastructure in the cloud today. You're responsible for securing your applications running on Outposts servers as you do in the Region today. You're also responsible for the physical security of your Outpost servers and ensuring consistent networking to them.
    - **Securing data:** Data-at-rest: Data is encrypted by default on instance store and for AMIs used for instance launch. Data-in-transit: Data is encrypted in transit between Outposts servers and the AWS Region through the Service Link. Deleting data: All data is deleted when instances are terminated in the same way as in the AWS Region.

163.1.1.3. Continuous Improvement

AWS is constantly adding new capabilities so you can leverage the latest technologies to experiment and innovate more quickly. To monitor recent AWS announcements relating to Outposts you can visit the AWS "What's New" feed here.

The latest features of Outposts servers are documented on the AWS website:

- AWS Outposts rack - link
- AWS Outposts servers - link

163.1.2. Benefits

- **Run AWS services on premises:** Extend AWS compute, networking, security, and other services on premises for low latency, local data processing, and to help meet data residency requirements.
- **Fully managed infrastructure:** Reduce the time, resources, operational risk, and maintenance downtime required to manage IT infrastructure with a fully managed experience.
- **Truly consistent hybrid experience:** Use the same hardware infrastructure, APIs, tools, and management controls available in the cloud to provide a truly consistent developer and IT operations experience.

163.2. 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. You're responsible for securing your applications

running on Outposts as you do in the Region today. You're also responsible for the physical security of your 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 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. 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.

163.3. Service Onboarding

Customers can select an AWS Outposts rack or AWS Outposts server configuration from the AWS Management Console and proceed to place an order for an Outpost, a pool of AWS compute and storage capacity deployed at a customer's site. 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.

163.4. 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. 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.

163.5. Backup/Restore and Disaster Recovery

This requirement is not applicable for AWS Outposts. For additional information beyond what is described herein, please refer to http://aws.amazon.com/documentation/.

163.6. Pricing Overview

Please see the AWS UK G Cloud 13 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.

163.6.1. Operating Systems

Prices do not include operating system (OS) charges. OS charges will be applied based on usage per instance.

163.6.2. 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.

163.6.3. 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. Learn more about data transfer here.

163.7. Service Constraints

163.7.1. 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 FAQs to understand what will happen if your external network connectivity is unavailable.

163.7.2. 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.

163.8. Technical Requirements

Please refer to https://docs.aws.amazon.com/outposts/ for comprehensive technical documentation.

164. AWS Panorama

164.1. Service Overview

AWS Panorama lets you add computer vision (CV) to your existing fleet of cameras with the AWS Panorama Appliance, which integrates seamlessly with your local area network. You can make predictions locally with high accuracy and low latency from a single management interface, where you can analyse video feeds in milliseconds; and can process video feeds at the edge, enabling you to control where your data is stored and operate with limited internet bandwidth.

164.1.1. Features

- **Setup in minutes:** Once installed and connected to your network, the AWS Panorama Appliance connects to the AWS Management Console where you can register your AWS Panorama Appliance and add video feeds from onsite cameras, deploy trained ML models, and run applications in minutes.
- **Connect IP cameras:** The AWS Panorama Appliance supports RTSP-enabled IP cameras, and allows you to add IP cameras that support the ONVIF standard.
- **Parallel multi-model multi-stream support:** The AWS Panorama Appliance supports connecting to multiple camera streams at a given time and supports running multiple ML models per stream.
- **Inference at the edge:** The AWS Panorama Appliance allows you to deploy CV applications to the edge, allowing you to run cloud-based machine learning where low-latency, data privacy, and limited internet bandwidth are concerns. AWS Panorama Appliance offers a flexible option for adding CV to automate tasks that traditionally require human inspection and monitoring.
- **GPU compute:** The AWS Panorama Appliance has a built-in Nvidia Xavier GPU for familiar development and fast machine learning computation at the edge.
- **Edge to cloud managed service:** You can easily discover existing fleets of IP cameras at the edge, create and deploy computer vision applications across AWS Panorama-enabled devices, manage the versioning and lifecycle of those applications as they're used for a variety of use cases, and other analytics to drive process improvements across multiple sites - all from a single management interface in the AWS console.
- **Flexible options for deploying CV:** AWS Panorama supports a growing ecosystem of pre-built applications from AWS and third party (3P) developers. Customers without machine learning experience can get started quickly with applications from AWS and 3Ps that enable people counting, vehicle classification, license plate recognition, and many more. The Panorama service also enables developers to train their own ML models with a variety of frameworks in the cloud and then quickly optimize them to run fast and accurately at the edge.
- **Ruggedized IP62 rated appliance:** The AWS Panorama Appliance is a ruggedized IP62 rated edge appliance, which makes it suitable for deployment in many types of environments as it is dust proof and water resistant.

164.1.2. Benefits

- **Application SDK:** Easily grab camera frames and perform machine learning (ML) inference on image data with a Python-based software development kit (SDK).

- **Container support:** Port functionality from your existing computer vision (CV) applications and bring custom libraries written in any language to AWS Panorama.
- **Management APIs:** Automate workflows with application programming interfaces (APIs) for device, camera, and application management.
- **End-to-end encryption:** Panorama makes it easier to secure your application by providing encryption for models, credentials, and assets in transit and at rest. Learn more.
- **Application samples:** Learn how to build for popular use cases with working sample code.
- **Wide ML framework support:** Use models trained in Tensorflow, Pytorch or MxNet in your application.

164.2. Backup/Restore and Disaster Recovery

This requirement is not applicable for this service. For additional information beyond what is described herein, please refer to http://aws.amazon.com/documentation/.

164.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

164.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/panorama/
- **Service FAQs:** https://aws.amazon.com/panorama/faqs/

164.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/panorama/ and the following links for comprehensive technical documentation regarding this service.

- Developer Guide: Provides a conceptual overview of AWS Panorama and detailed instructions for using its features.
- API Reference: Describes all the API operations for AWS Panorama in detail.

165. AWS Personal Health Dashboard

165.1. Service Overview

AWS Personal Health Dashboard provides alerts and guidance for AWS events that might affect your environment. While the Service Health Dashboard shows the general status of AWS services, the Personal Health Dashboard provides proactive and transparent notifications about your specific AWS environment.

All AWS customers can access the Personal Health Dashboard. The Personal Health Dashboard shows recent events to help you manage active events, and shows proactive notifications so that you can plan for scheduled activities. Use these alerts to get notified about changes that can affect your AWS resources, and then follow the guidance to diagnose and resolve issues.

165.1.1. Features

- **Proactive notifications:** Unlike the Service Health Dashboard, you can use the Personal Health Dashboard to create alerts for specific events that might affect your account. You can set up alerts across multiple channels, including email and mobile notifications, to receive timely and relevant information to help plan for scheduled changes. For example, if a maintenance event is scheduled for one of your Amazon EC2 instances, you can receive an alert with information to help you plan for, and proactively address any issues for the upcoming change.
- **Detailed troubleshooting guidance:** When you get an alert, it includes remediation details and specific guidance so that you can take action for events that affect your resources. For example, if a hardware issue affects one of your Amazon Elastic Block Store (EBS) volumes, the alert includes a list of affected resources, and recommendations and help links to restore your volume from a snapshot. This helps you reduce the amount of time to resolve issues.
- **Integration and automation:** You can use Amazon CloudWatch Events to build custom rules and select targets, such as AWS Lambda functions, to define automated remediation actions for specific events. You can use the AWS Health API, the service that powers Personal Health Dashboard, to integrate health data and notifications with your existing in-house or third-party IT management tools. The AWS Health API is part of an AWS Business Support or AWS Enterprise Support plan.
- **Fine-grained access control by using IAM:** The Personal Health Dashboard supports access control so that you can set up permissions based on event metadata. This allows you to grant or deny access to an AWS Identity and Access Management (IAM) user based on attributes, such as event types, specific services, or other role-based attributes. You can restrict access of sensitive events, such as security events, to only the users that need to see them.

165.1.2. Benefits

- **Personalized view of service health:** The Personal Health Dashboard gives you a personalized view of the status of AWS services that power your applications. Use the Personal Health Dashboard to learn about specific operational issues that affect your account. For example, if you receive an event for a lost Amazon Elastic Block Store (EBS) volume associated with one of your Amazon EC2 instances, you can use the event to quickly view the status of your impacted resources, and then troubleshoot and determine remediation steps.
- **Aggregate health events across AWS Organizations:** If you use AWS Organizations, you can use AWS Health to aggregate notifications from all accounts in your organization. This provides you a centralized and real-time view for all AWS Health events posted to individual accounts in your organization, including operational issues, scheduled maintenance, and account notifications.

165.2. Backup/Restore and Disaster Recovery

This requirement is not applicable for this service. For additional information beyond what is described herein, please refer to http://aws.amazon.com/documentation/.

165.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

165.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/health/latest/ug/getting-started-phd.html
- **Service quotas:** https://docs.aws.amazon.com/general/latest/gr/awshealth.html
- **Service FAQs:** https://aws.amazon.com/premiumsupport/faqs/?nc=sn&loc=6

165.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/health/latest/ug/getting-started-phd.html for comprehensive technical documentation regarding this service.

166. AWS PrivateLink

166.1. Service Overview

AWS PrivateLink provides private connectivity between VPCs, AWS services, and your on-premises networks, without exposing your traffic to the public internet. AWS PrivateLink makes it easy to connect services across different accounts and VPCs to significantly simplify your network architecture.

Interface VPC endpoints, powered by AWS PrivateLink, connect you to services hosted by AWS Partners and supported solutions available in AWS Marketplace. By powering Gateway Load Balancer endpoints, AWS PrivateLink brings the same level of security and performance to your virtual network appliances or custom traffic inspection logic.

166.1.1. Features

- **Accessing services over AWS PrivateLink:** To use AWS PrivateLink, create an interface VPC endpoint for a service outside of your VPC. This creates an elastic network interface in your subnet with a private IP address that serves as an entry point for traffic destined to the service. For more information, see VPC Endpoints.
- **Sharing your services over AWS PrivateLink:** You can create your own AWS PrivateLink-powered service (endpoint service) and enable other AWS customers to access your service. For more information, see VPC endpoint services (AWS PrivateLink).
- **Privately connecting to your on-premises applications:** Interface VPC endpoints support private connectivity over AWS Direct Connect, so that applications in your premises will be able to connect to these services via the Amazon private network.
- **Integration with AWS Marketplace:** AWS PrivateLink is integrated with AWS Marketplace through an easy lookup of the services that are available over AWS PrivateLink. To facilitate the identification of which services are attached to your endpoint, services that are available from AWS Marketplace are supported with vanity DNS names.You can access AWS Marketplace through the AWS PrivateLink-dedicated page here.

166.1.2. Benefits

- **Secure Your Traffic:** Connect your VPCs to services in AWS in a secure and scalable manner with AWS PrivateLink. Network traffic that uses AWS PrivateLink doesn't traverse the public internet, reducing exposure to brute force and distributed denial-of-

service attacks, along with other threats. You can use private IP connectivity so that your services function as though they were hosted directly on your private network. You can also associate security groups and attach an endpoint policy to interface endpoints, which allow you to control precisely who has access to a specified service. AWS connections powered by PrivateLink, such as interface VPC endpoints and Gateway Load Balancer endpoints, deliver the same benefits of security, scalability, and performance.

- **Simplify Network Management:** You can connect services across different accounts and Amazon VPCs, with no need for firewall rules, path definitions, or route tables. There is no need to configure an Internet gateway, VPC peering connection, or manage VPC Classless Inter-Domain Routing (CIDRs). Because AWS PrivateLink simplifies your network architecture, it is easier for you to manage your global network.
- **Accelerate Your Cloud Migration:** More easily migrate traditional on-premises applications to SaaS offerings hosted in the cloud with AWS PrivateLink. Since your data does not get exposed to the Internet where it can be compromised, you can migrate and use more cloud services with the confidence that your traffic remains secure. You no longer have to choose between using a service and exposing your critical data to the Internet. You can find the latest controls in place to help customers stay compliant on our AWS Compliance Programs page.

166.2. Backup/Restore and Disaster Recovery

This requirement is not applicable for this service. For additional information beyond what is described herein, please refer to http://aws.amazon.com/documentation/.

166.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

166.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/vpc/index.html
- **Service quotas:** https://docs.aws.amazon.com/vpc/latest/privatelink/vpc-limits-endpoints.html
- Service FAQs: https://aws.amazon.com/privatelink/faqs/ Technical Requirements

Please refer to https://docs.aws.amazon.com/vpc/index.html and the following links for comprehensive technical documentation regarding this service.

- AWS PrivateLink Guide: Configure VPC endpoints and VPC endpoint services.
- AWS PrivateLink API Reference: Get syntax and examples for the API actions.

167. AWS Proton

167.1. Service Overview

AWS Proton is the first fully managed delivery service for container and serverless applications. Platform engineering teams can use AWS Proton to connect and coordinate all the different tools needed for infrastructure provisioning, code deployments, monitoring, and updates.

Maintaining hundreds – or sometimes thousands – of microservices with constantly changing infrastructure resources and continuous integration/continuous delivery (CI/CD) configurations is a nearly impossible task for even the most capable platform teams. AWS Proton solves this by giving platform teams the tools they need to manage this complexity and enforce consistent standards, while making it easy for developers to deploy their code using containers and serverless technologies.

167.1.1. Features

- **Automated deployments:** AWS Proton makes it easy for platform teams to create application stacks. This includes the CI/CD pipeline available to developers, so they can make a request through the application programming interface (API), command-line interface (CLI), or user interface (UI) to deploy an application immediately.
- **Customer-managed environments:** You can bring your existing shared resources into AWS Proton without re-creating the infrastructure. This is transparent for developers, who can deploy to a customer-managed environment in the same way they deploy to a standard environment.
- **Flexible definitions:** Create service templates with or without a pipeline. AWS Proton gives teams greater flexibility in defining, provisioning, and deploying their services. Developers only need to provide the required inputs for their service, and platform teams can leverage AWS Proton's central management capabilities to ensure that all deployments are up-to-date.
- **Multi-account support:** AWS Proton supports multi-account infrastructures, which help platform operators configure their architecture securely across multiple AWS accounts. You can manage all your multi-account environments and services from a single account using AWS Proton.
- **Self-service interface:** Customize your user interface using the familiar AWS Management Console or CLI. The AWS Proton interface guides you through the process of creating and deploying shared resources as environments to which you can deploy services. Proton also gives you end-to-end provisioning support, including the ability to deploy infrastructure such as compute, database, and many other resources in a simple, declarative style through AWS CloudFormation.
- **Streamlined upgrades:** AWS Proton supports versioning of infrastructure templates and provides developers with updates for out-of-date deployments.
- **Tagging capabilities:** Establish tagging and tag-based access control for any AWS Proton resource, including templates, environments, and services. Streamline and ensure consistency in your tagging process by propagating tags applied to a parent resource down to any of its child resources. AWS Proton also tags all provisioned resources automatically with unique identifiers, allowing you to identify all provisioned resources coming from an AWS Proton-specific template or environment.
- **Template management:** Platform teams use AWS Proton to create a stack presented to their developers as a reusable version-controlled template. These stacks are defined using infrastructure as code in a simple, declarative style with everything needed to provision, deploy, and manage a service including compute, networking, code pipeline, security, and monitoring resources. Developers log into the AWS Proton console to use published AWS Proton stacks to automate infrastructure provisioning and quickly deploy their application code. Developers using AWS Proton don't need to separately provision the components of their stack (like shared resources, continuous integration/continuous deployment (CI/CD) pipeline, and observability tools).

167.1.2. Benefits

- **Guardrails:** Set guardrails with managed, approved stacks to help developers build and deploy applications.
- **Productivity:** Empower developer productivity and innovation with infrastructure provisioning and code deployment in a single interface.
- **One-click updates:** Update applications with a single click to easily maintain a consistent architecture across your organization.
- **Streamline management:** Gain deployment visibility and set consistent standards for compute, networking, continuous integration and deployment (CI/CD), security, and monitoring.
- **Empower self-service development:** Build a curated, self-service interface for developers to create and deploy production infrastructure from approved application stacks.
- **Adopt infrastructure as code:** Accelerate developer innovation using infrastructure as code to define application stacks and configure resources.

167.2. Backup/Restore and Disaster Recovery

This requirement is not applicable for this service. For additional information beyond what is described herein, please refer to http://aws.amazon.com/documentation/.

167.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

167.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/proton/index.html
- **Service quotas:** https://docs.aws.amazon.com/proton/latest/userguide/ag-limits.html
- **Service FAQs:** https://aws.amazon.com/proton/faqs/

167.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/proton/index.html and the following links for comprehensive technical documentation regarding this service.

- Administrator Guide: Provides a conceptual view of AWS Proton and describes how platform teams create and manage infrastructure and deployment tooling.
- User Guide: Provides a conceptual view of AWS Proton and describes how developers deploy their services or applications.
- API Reference: Documents the AWS Proton Query API.
- AWS Proton section of AWS CLI Reference: Documents the AWS Proton commands available in the AWS Command Line Interface (AWS CLI).

168. AWS Resource Access Manager (RAM)

168.1. Service Overview

AWS Resource Access Manager (RAM) helps you securely share your resources across AWS accounts, within your organization or organizational units (OUs) in AWS Organizations, and with IAM roles and IAM users for supported resource types. You can use AWS RAM to share transit gateways, subnets, AWS License Manager license configurations, Amazon Route 53 Resolver rules, and more resource types.

Many organizations use multiple accounts to create administrative or billing isolation, and to limit the impact of errors. With AWS RAM, you don't need to create duplicate resources in multiple AWS accounts. This reduces the operational overhead of managing resources in every account that you own. Instead, in your multi-account environment, you can create a resource once, and use AWS RAM to share that resource across accounts by creating a resource share. When you create a resource share, you select the resources to share, choose an AWS RAM managed permission per resource type, and specify whom you want to have access to the resources. AWS RAM is available to you at no additional charge.

168.1.1. Features

- **Reduce Operational Overhead:** Procure AWS resources centrally, and use RAM to share resources such as subnets or License Manager configurations with other accounts. This eliminates the need to provision duplicate resources in every account in a multi-account environment, reducing the operational overhead of managing those resources in every account.

168.1.2. Benefits

- **Improve Security and Visibility:** RAM leverages existing policies and permissions set in AWS Identity and Access Management (IAM) to govern the consumption of shared resources. RAM also provides comprehensive visibility into shared resources to set alarms and visualize logs through integration with Amazon CloudWatch and AWS CloudTrail.
- **Optimize Costs:** Sharing resources such as AWS License Manager configurations across accounts allows you to leverage licenses in multiple parts of your company to increase utilization and optimize costs.

168.2. Backup/Restore and Disaster Recovery

This requirement is not applicable for this service. For additional information beyond what is described herein, please refer to http://aws.amazon.com/documentation/.

168.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

168.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/ARG/index.html
- **Service quotas:** https://docs.aws.amazon.com/general/latest/gr/ram.html
- **Service FAQs:** https://aws.amazon.com/ram/faqs/

168.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/ARG/index.html and the following links for comprehensive technical documentation regarding this service.

- AWS RAM User Guide: Introduces you to AWS RAM and helps you set up a resource share. When you create a resource share, you select the resources to share, choose an AWS RAM managed permission per resource type, and specify whom you want to have access to the resources.
- AWS RAM API Reference: Describes all the API operations for AWS RAM. Also provides sample requests, responses, and errors for the supported web services protocols.
- AWS RAM section of the AWS CLI Reference: Describes the AWS CLI commands that you can use to administer AWS RAM.

169. AWS RoboMaker

169.1. Service Overview

AWS RoboMaker provides the most complete cloud solution for robotics development. It features fully-managed simulation capabilities built on scalable infrastructure, and provides application deployment capabilities, an optional independent development environment (IDE), Robot Operating System (ROS) extensions for robots to use AWS Cloud services, and seamless integration with various AWS management, monitoring, security and storage capabilities to empower customers to innovate and provide best-of-class robotic solutions.

169.1.1. Features

- **Simulation with AWS RoboMaker:** AWS RoboMaker is a fully managed service that enables developers, QAs, and DevOps engineers to easily create simulation worlds and run simulation jobs without provisioning or managing any infrastructure. RoboMaker makes simulation at scale affordable and accessible to all robotics companies by providing tools for developers to test and iterate code in 3D virtual environments. The service supports large-scale and parallel simulations, and automatically scales based on the complexity of the scenarios being tested. With RoboMaker simulation, robotics companies can make robotics application testing and machine learning faster, less expensive, and more robust.
- **Run large-scale, parallel simulations**: With the RoboMaker batch simulation API, developers can easily launch a large batch of simulations with a single API call. Running large-scale simulations through an API makes it easier to access compute power, allowing developers to increase the complexity, scale, and frequency of their tests, which speeds up the development lifecycle and makes testing more robust.
- **Use any simulator**: RoboMaker Simulation can be used to run your simulator of choice, including Unity, Unreal, and Nvidia Isaac-based simulations. RoboMaker enables you to run your simulations in the cloud without provisioning, configuring, or managing any infrastructure.
- **Managed ROS/ Gazebo environment**: RoboMaker Simulation can be used to run the open-source software library known as Robot Operating System (ROS) and ROS 2 applications in simulation using the open-source Gazebo robot simulation engine. RoboMaker enables you to run Gazebo-based simulations in the cloud without provisioning, configuring, or managing any infrastructure. The service supports the

Gazebo graphical client for interacting with a running simulation job, rviz for visualizing sensor data, rqt for running various GUI tools, and command line for interacting with the running robot application.

- **Enhance robot functionality:** AWS has developed cloud extensions to Amazon services that enable developers to enhance the functionality of their robots and collect data from them, without installing additional hardware or developing complex software. You can use cloud extensions to enhance ROS-based robot functionality with Amazon Rekognition for object detection, Amazon Kinesis for video streaming, Amazon Polly for converting text to speech, and Amazon Lex for speech recognition. For operations, you can use cloud extensions to pull performance and operational data from robots using Amazon CloudWatch for metrics, logging, and monitoring. You can also use a ROS extension to upload rosbags and files from robots for storage in Amazon S3. AWS provides each of these cloud service extensions as open-source ROS packages that customers access via cloud APIs. This integrated suite of AWS services makes it easy for customers to monitor and tune the performance of their robotic applications in the field. ROS Cloud Extensions are currently supported for ROS Melodic only.
- **Securely manage and deploy applications:** RoboMaker's application deployment service is integrated with AWS IoT Greengrass to provide robot registry, security, and fault-tolerance. The registry service enables companies to identify, track, and organize their robots into optimal fleets. Developers can use RoboMaker application deployment to securely deploy their application to their robots via AWS' fully-managed over-the-air (OTA) update infrastructure. AWS IoT Greengrass uses X.509 certificates, managed subscriptions, AWS IoT policies, and IAM roles for secure connection to AWS cloud services through encrypted connections. RoboMaker's OTA service supports conditional updates which provides intelligence into the OTA process to lower the risk of interrupted or incomplete software updates.
- **Development Environment:** RoboMaker's development environment is a customized environment in AWS Cloud9 for robotics development. This environment comes with ROS pre-installed and includes sample applications. This environment is also integrated with other RoboMaker capabilities such as simulation so that you can use these capabilities from the interface of the development environment.

169.1.2. Benefits

- **Simulations:** Run large-scale and parallel simulations with a single API call.
- **Scalable and cost-effective:** Cost-effectively scale and automate simulation workloads.
- **Easy to use:** Easily create user-defined, randomized 3D virtual environments.
- **Automated regression testing:** Automate testing within a continuous integration and continuous delivery (CI/CD) pipeline.
- **Train reinforcement learning models:** Train reinforcement learning models with high volumes of iterative trials.
- **Multi-robot testing:** Connect multiple concurrent simulations to your fleet management software for testing.

169.2. Backup/Restore and Disaster Recovery

This requirement is not applicable for this service. For additional information beyond what is described herein, please refer to http://aws.amazon.com/documentation/.

169.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

169.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/robomaker/
- **Service quotas:** https://docs.aws.amazon.com/general/latest/gr/robomaker.html
- **Service FAQs:** https://aws.amazon.com/robomaker/faqs/

169.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/robomaker/ and the following link for comprehensive technical documentation regarding this service.

- **Developer Guide**: Provides a conceptual overview of AWS RoboMaker, instructions for managing robots, simulations and fleets, and includes a complete API reference for developers.

170. AWS Secrets Manager

170.1. Service Overview

AWS Secrets Manager helps you protect secrets needed to access your applications, services, and IT resources. The service enables you to easily rotate, manage, and retrieve database credentials, API keys, and other secrets throughout their lifecycle. Users and applications retrieve secrets with a call to Secrets Manager APIs, eliminating the need to hardcode sensitive information in plain text. Secrets Manager offers secret rotation with built-in integration for Amazon RDS, Amazon Redshift, and Amazon DocumentDB. Also, the service is extensible to other types of secrets, including API keys and OAuth tokens. In addition, Secrets Manager enables you to control access to secrets using fine-grained permissions and audit secret rotation centrally for resources in the AWS Cloud, third-party services, and on-premises.

170.1.1. Features

- **Secure secrets storage:** AWS Secrets Manager encrypts secrets at rest using encryption keys that you own and store in AWS Key Management Service (KMS). When you retrieve a secret, Secrets Manager decrypts the secret and transmits it securely over TLS to your local environment. By default, Secrets Manager does not write or cache the secret to persistent storage. And, you can control access to the secret using fine-grained AWS Identity and Access Management (IAM) policies and resource-based policies. You can also tag secrets individually and apply tag-based access controls. For example, you can tag secrets used in the production environment as “Prod,” and then write an IAM policy to grant access to these secrets only if the requests are coming from within the corporate IT network.
- **Automatic secrets rotation without disrupting applications:** With AWS Secrets Manager, you can rotate secrets on a schedule or on demand by using the Secrets Manager console, AWS SDK, or AWS CLI. For example, to rotate a database password, you provide the database type, rotation frequency, and master database credentials when storing the password in Secrets Manager. Secrets Manager natively supports rotating credentials for databases hosted on Amazon RDS and Amazon DocumentDB

and clusters hosted on Amazon Redshift. You can extend Secrets Manager to rotate other secrets by modifying sample Lambda functions. For example, you can rotate OAuth refresh tokens used to authorize applications or passwords used for MySQL databases hosted on-premises. Users and applications retrieve secrets by replacing hardcoded secrets with a call to Secrets Manager APIs, enabling you to automate secret rotation while ensuring applications run without interruption.

- **Automatic replication of secrets to multiple AWS Regions:** With AWS Secrets Manager, you can automatically replicate your secrets to multiple AWS Regions to meet your unique disaster recovery and cross-regional redundancy requirements. By using the Secrets Manager console, AWS SDK, AWS CLI or AWS CloudFormation, you can specify the AWS Regions where a secret need to be replicated and Secrets Manager will securely create regional read replicas, eliminating the need to maintain a complex solution for this functionality. You can now give your multi-Region applications access to replicated secrets in the required Regions and rely on Secrets Manager to keep the replicas in sync with the primary secret.
- **Programmatic retrieval of secrets:** You can store and retrieve secrets using the AWS Secrets Manager console, AWS SDK, AWS CLI, or AWS CloudFormation. To retrieve secrets, you simply replace plaintext secrets in your applications with code to pull in those secrets programmatically using the Secrets Manager APIs. Secrets Manager provides code samples to call Secrets Manager APIs, also available on the Secrets Manager Resources page. You can configure Amazon Virtual Private Cloud (VPC) endpoints to keep traffic between your VPC and Secrets Manager within the AWS network. You can also use Secrets Manager client-side caching libraries to improve the availability and reduce the latency of using your secrets.
- **Audit and monitor secrets usage:** AWS Secrets Manager enables you to audit and monitor secrets through integration with AWS logging, monitoring, and notification services. For example, after enabling AWS CloudTrail for an AWS region, you can audit when a secret is stored or rotated by viewing AWS CloudTrail logs. Similarly, you can configure Amazon CloudWatch to receive email messages using Amazon Simple Notification Service when secrets remain unused for a period, or you can configure Amazon CloudWatch Events to receive push notifications when Secrets Manager rotates your secrets.
- **Compliance:** You can use AWS Secrets Manager to manage secrets for workloads that are subject to Department of Defense Cloud Computing Security Requirements Guide (DoD CC SRG IL2, DoD CC SRG IL4, and DoD CC SRG IL5), Federal Risk and Authorization Management Program (FedRAMP), U.S. Health Insurance Portability and Accountability Act (HIPAA), Information Security Registered Assessors Program (IRAP), Outsourced Service Provider's Audit Report (OSPAR), ISO/IEC 27001, ISO/IEC 27017, ISO/IEC 27018, ISO 9001, Payment Card Industry Data Security Standard (PCI-DSS), or System and Organization Control (SOC). View details of AWS's compliance program and report in AWS Artifact.

170.1.2. Benefits

- **Rotate secrets safely:** AWS Secrets Manager helps you meet your security and compliance requirements by enabling you to rotate secrets safely without the need for code deployments. For example, Secrets Manager offers built-in integration for Amazon RDS, Amazon Redshift, and Amazon DocumentDB and rotates these database credentials on your behalf automatically. You can customize Lambda functions to extend Secrets Manager rotation to other secret types, such as API keys and OAuth tokens.

Retrieving the secret from Secrets Manager ensures that developers and applications are using the latest version of your secrets.

- **Manage access with fine-grained policies:** With Secrets Manager, you can manage access to secrets using fine-grained AWS Identity and Access Management (IAM) policies and resource-based policies. For example, you can create a policy that enables developers to retrieve certain secrets only when they are used for the development environment. The same policy could enable developers to retrieve passwords used in the production environment only if their requests are coming from within the corporate IT network. For the database administrator, a policy can be built to allow the database administrator to manage all database credentials and permission to read the SSH keys required to perform OS-level changes to the particular instance hosting the database.
- **Secure and audit secrets centrally:** Using Secrets Manager, you can help secure secrets by encrypting them with encryption keys that you manage using AWS Key Management Service (KMS). It also integrates with AWS' logging and monitoring services for centralized auditing. For example, you can audit AWS CloudTrail logs to see when Secrets Manager rotates a secret or configure AWS CloudWatch Events to notify you when an administrator deletes a secret.
- **Pay as you go:** Secrets Manager offers pay as you go pricing. You pay for the number of secrets managed in Secrets Manager and the number of Secrets Manager API calls made. Using Secrets Manager, you can enable a highly available secrets management service without the upfront investment and on-going maintenance costs of operating your own infrastructure.
- **Easily replicate secrets to multiple regions:** AWS Secrets Manager enables you to easily replicate secrets in multiple AWS regions to support your multi-region applications and disaster recovery scenarios. The multi-Region secrets feature abstracts the complexity of replicating and managing secrets across multiple regions, enabling you to simply access and read secrets where you need them.

170.2. Backup/Restore and Disaster Recovery

This requirement is not applicable for this service. For additional information beyond what is described herein, please refer to http://aws.amazon.com/documentation/.

170.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

170.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/secretsmanager/index.html
- **Service quotas:** https://docs.aws.amazon.com/secretsmanager/latest/userguide/reference_limits.html
- **Service FAQs:** https://aws.amazon.com/secrets-manager/faqs/

170.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/secretsmanager/index.html and the following links for comprehensive technical documentation regarding this service.

- User Guide: Walks you through creating and managing your secrets, retrieving them in your application, and automatically rotating them to help keep them secure.
- API Reference: Describes all the API operations for AWS Secrets Manager. It provides sample requests, responses, and errors that can be generated.

171. AWS Security Hub

171.1. Service Overview

AWS Security Hub is a cloud security posture management service that performs automated, continuous security best practice checks against your AWS resources. Security Hub aggregates your security alerts (i.e. findings) from various AWS services and partner products in a standardized format so that you can more easily take action on them. To maintain a complete view of your security posture in AWS, you need to integrate multiple tools and services including threat detections from Amazon GuardDuty, vulnerabilities from Amazon Inspector, sensitive data classifications from Amazon Macie, resource configuration issues from AWS Config, and AWS Partner Network Products. Security Hub simplifies how you understand and improve your security posture with automated security best practice checks powered by AWS Config rules and automated integrations with dozens of AWS services and partner products.

Security Hub enables you to understand your overall security posture via a consolidated security score across all of your AWS accounts, automatically assesses the security of your AWS accounts resources via the AWS Foundational Security Best Practices standard and other compliance frameworks. It also aggregates all of your security findings from dozens of AWS security services and APN products in a single place and format via the AWS Security Finding Format, and reduces your Mean Time To Remediation (MTTR) with automated response and remediation support. Security Hub has out-of-the-box integrations with ticketing, chat, Security Information and Event Management (SIEM), Security Orchestration Automation and Response (SOAR), threat investigation, Governance Risk and Compliance (GRC), and incident management tools to provide your users with a complete security operations workflow.

Getting started with Security Hub requires just a few clicks from the Management Console to begin aggregating findings and conducting security checks using our 30-day free trial. You can integrate Security Hub with AWS Organizations to automatically enable the service in all accounts in your organization.

171.1.1. Features

- **Automated, continuous security best practice checks:** Security Hub provides you with a set of automated security controls called the AWS Foundational Security Best Practices standard. This is a highly curated set of security best practices vetted by our AWS security experts that either run continuously whenever there are changes to the associated resources or on a set periodic schedule. Each control has a specific severity score to help you prioritize your remediation efforts. We recommend that this standard is enabled across all accounts and regions, and we are continuously updating it with new controls and additional service coverage.
- **Consolidated findings across AWS services and partner integrations:** Security Hub automatically collects and consolidates findings from AWS security services enabled in

your environment, such as intrusion detection findings from Amazon GuardDuty, vulnerability scans from Amazon Inspector, Amazon Simple Storage Service (Amazon S3) bucket policy findings from Amazon Macie, publicly accessible and cross-account resources from IAM Access Analyzer, and resources lacking WAF coverage from AWS Firewall Manager. AWS Security Hub also consolidates findings from dozens of integrated AWS Partner Network (APN) security solutions. All findings are stored in Security hub for 90 days after last update date.

- **A single, standardized data format for all of your findings:** Traditionally, when combining security alerts into a single system, you would need to parse and normalize each data source to get it into a common format for search, analytics, and response and remediation actions. Security Hub eliminates these time-consuming and resource-intensive processes by introducing the AWS Security Findings Format (ASFF). With the ASFF, all of Security Hub's integration partners (including both AWS services and external partners) send their findings to Security Hub in a well-typed JSON format consisting of over 1,000 available fields. This means that all of your security findings are normalized before they are ingested into Security Hub, and you don't need to do any parsing and normalization yourself. The findings identify resources, severities, and timestamps in a consistent way, so that you can more easily search and take action on them.
- **Security standards aligned to regulatory and industry compliance frameworks:** In addition to the AWS Foundational Security Best Practices standard, Security Hub also offers additional standards aligned to industry and regulatory frameworks, such as the Payment Card Industry Data Security Standard (PCI DSS) and the Center for Internet Security (CIS) AWS Foundations Benchmark. These standards are also powered by continuous, automated security checks, and you only pay once for the a security check regardless of how many standards it is mapped to.
- **Automated response, remediation, and enrichment actions:** You can create custom automated response, remediation, and enrichment workflows using Security Hub's integration with Amazon EventBridge. All of Security Hub findings are automatically sent to EventBridge, and you can create EventBridge rules that have AWS Lambda functions, AWS Step Function functions, or AWS Systems Manager Automation runbooks as their targets. These functions or runbooks can automatically enrich findings with additional data or take automated response and remediation actions on the findings. Security Hub also supports sending findings to EventBridge on demand via custom actions, so that you can have an analyst decide when to trigger an automated response or remediation action. The Security Hub Automated Response and Remediation (SHARR) solution provides you with prepackaged EventBridge rules for you to deploy via AWS CloudFormation.
- **Multi-account and AWS Organizations support:** You can connect multiple AWS accounts and consolidate findings across those accounts with a few clicks in the AWS Security Hub console. By designating an administrator account, you can enable your security team to see consolidated findings for all accounts, while individual account owners see only findings associated with their account. Integration with AWS Organizations allows you to automatically enable any account in your organization with Security Hub and the AWS Foundational Security Best Practices standard.
- **Cross-Region aggregation of findings:** AWS Security Hub allows you to designate an aggregator Region and link some or all Regions to that aggregator Region to give you a centralized view of all your findings across all your accounts and all your linked Regions. After linking a Region to the aggregator Region, your findings are continuously synced

between the Regions, so that any update made to a finding in one Region is replicated to the other Region. Your Security Hub administrator or delegated administrator account in your aggregator Region can view and manage all of your findings. Individual Security Hub member accounts in the aggregator Region can also view and manage all of their findings across all linked Regions. Your Amazon EventBridge feed in your administrator account and aggregator Region also now includes all your findings across all member accounts and linked Regions, which allows you to simplify integrations with ticketing, chat, incident management, logging, and auto-remediation tools by consolidating those integrations into your aggregator Region.

- **Integrations with ticketing, chat, incident management, investigation, GRC, SOAR, and SIEM tools:** In addition to integrating with dozens of AWS security services and partner products that send Security Hub findings, Security Hub also has integrations with various ticketing, chat, incident management, threat investigation, Governance Risk and Compliance (GRC), Security Orchestration Automation and Response (SOAR), and Security Information and Event Management (SIEM) tools that can automatically receive findings from Security Hub. These integrations include AWS services such as Amazon Detective (threat investigations) and AWS Audit Manager and various partner tools such as Splunk, Slack, PagerDuty, Sumo Logic, ServiceNow ITSM, and Atlassian's Jira Service Management. The integration with ServiceNow and Jira are bi-directional, so that any updates to tickets are synced with the findings in Security Hub.
- **Security scores and summary dashboards:** Security Hub provides a simple 0-100 security score for each standard, for each account across all enabled standards, and a total score for all accounts associated with your administrator account. This score is based on the number of controls that have passed vs. failed for a standard, account, and/or organization. This information is presented along with other key insights, such as which resources have the most failed security checks in summary dashboards to help you monitor your security posture.
- **Filtering, grouping, and saved searches for your findings:** You can filter findings based on fields in the AWS Security Finding Format and use GroupBy statements to aggregate findings into buckets. For example, you can filter findings to show only Critical or High severity findings and then group them by resource IDs to see which resources have the most critical or high findings. Security Hub calls these types searches insights, and Security Hub provides both prepackaged managed insights and lets you define your own custom insights. Each insight includes a time series sparkline to show the trend over time in findings that match the insight.

171.1.2. Benefits

- **Improve Security:** Detect deviations from security best practices with a single click.
- **Security as data:** Automatically aggregate security findings in a standardized data format from AWS and partner services.
- **Automated, fast response:** Accelerate mean time to resolution with automated response and remediation actions.

171.2. Backup/Restore and Disaster Recovery

This requirement is not applicable for this service. For additional information beyond what is described herein, please refer to http://aws.amazon.com/documentation/.

171.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

171.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/securityhub/index.html
- **Service quotas:** https://docs.aws.amazon.com/securityhub/latest/userguide/securityhub_limits.html
- **Service FAQs:** https://aws.amazon.com/security-hub/faqs/

171.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/securityhub/index.html and the following links for comprehensive technical documentation regarding this service.

- AWS Security Hub User Guide: Describes how to set up, configure, and use AWS Security Hub to evaluate the security and compliance state of your AWS environment.
- AWS Security Hub API Reference: Describes all of the API operations for AWS Security Hub.
- AWS Security Hub Partner Integration Guide: Describes how AWS Partner Network (APN) Partners set up an integration with AWS Security Hub.

172. AWS Service Catalog

172.1. Service Overview

AWS Service Catalog allows organizations to create and manage catalogs of IT services that are approved for use on AWS. These IT services can include everything from virtual machine images, servers, software, and databases to complete multi-tier application architectures. AWS Service Catalog allows you to centrally manage deployed IT services and your applications, resources, and metadata. This helps you achieve consistent governance and meet your compliance requirements, while enabling users to quickly deploy only the approved IT services they need. With AWS Service Catalog AppRegistry, organizations can understand the application context of their AWS resources. You can define and manage your applications and their metadata, to keep track of cost, performance, security, compliance and operational status at the application level.

172.1.1. Features

- **Products:** A product is an IT service that you want to make available for deployment on AWS. A product can comprise one or more AWS resources, such as EC2 instances, storage volumes, databases, monitoring configurations, and networking components, or packaged AWS Marketplace products. A product can be a single compute instance running AWS Linux, a fully configured multi-tier web application running in its own environment, or anything in between. You create your products by importing AWS CloudFormation templates. These templates define the AWS resources required for the product, the relationships between resources, and the parameters that the end user can plug in when they launch the product to configure security groups, create key pairs, and perform other customizations.

- **Portfolios:** A portfolio is a collection of products, together with configuration information. Portfolios help manage product configuration, and who can use specific products and how they can use them. With AWS Service Catalog, you can create a customized portfolio for each type of user in your organization and selectively grant access to the appropriate portfolio. When you add a new version of a product to a portfolio, that version is automatically available to all current users of that portfolio. You also can share your portfolios with other AWS accounts and allow the administrator of those accounts to distribute your portfolios with additional constraints. For example, for developers, you can define a portfolio of development environments, such as a LAMP stack with approved versions that users can use for software development and testing. You could also define a portfolio for the marketing organizations that includes campaign websites and market analysis applications.
- **Versioning:** AWS Service Catalog allows you to manage multiple versions of the products in your catalog. This allows you to add new versions of templates and associated resources based on software updates or configuration changes. When you create a new version of a product, the update is automatically distributed to all users who have access to the product, allowing the user to select which version of the product to use. Users can update running instances of the product to the new version quickly and easily.
- **Granular access control:** Granting a user access to a portfolio enables that user to browse the portfolio and launch the products in it. You apply AWS Identity and Access Management (IAM) permissions to control who can view and modify your products and portfolios. IAM permissions can be assigned to IAM users, groups, and roles. When a user launches a product that has an IAM role assigned to it, AWS Service Catalog uses the role to launch the product's cloud resources using AWS CloudFormation. By assigning an IAM role to each product, you can avoid giving users permissions to perform unapproved operations, and enable them to provision resources using the catalog.
- **Constraints:** Constraints restrict the ways that specific AWS resources can be deployed for a product. You can use them to apply limits to products for governance or cost control. There are two types of constraints: template and launch. Template constraints restrict the configuration parameters that are available for the user when launching the product (for example, EC2 instance types or IP ranges). Template constraints allow you to reuse generic AWS CloudFormation templates for products and apply restrictions to the templates on a per-product or per-portfolio basis. Launch constraints allow you to specify a role for a product in a portfolio. This role is used to provision the resources at launch, so you can restrict user permissions without impacting users' ability to provision products from the catalog. For example, for marketing users, you can enable them to create campaign websites, but use constraints to restrict their access to provision the underlying databases.
- **Stack:** Every AWS Service Catalog product is launched as an AWS CloudFormation stack., which is a set of resources provisioned for that instance of the product. AWS CloudFormation stacks make it easier to manage the lifecycle of your product by allowing you to provision, tag, update, and terminate your product instance as a single unit.
- **Service actions:** Using service actions, you can enable end users to perform operational tasks, troubleshoot issues, run approved commands, or request permissions in AWS Service Catalog on your provisioned products, without needing to grant end

users full access to AWS services. You use AWS Systems Manager documents to define service actions. The AWS Systems Manager documents provide access to pre-defined actions that implement AWS best practices, such as Amazon EC2 stop and reboot, and you can define custom actions too.

- **Applications:** Builders can define their applications within Service Catalog AppRegistry by providing a name, description, associations to application metadata, and associations to CloudFormation stacks. The associated attribute groups represent the metadata that your enterprise creates and manages for the application. The associated CloudFormation stacks represent the AWS resources associated to the application. This might be the infrastructure required in a single environment, or it could also include the code repositories and pipelines that support the application across all environments. Either existing or new CloudFormation Stacks can be associated to applications. Stacks can be associated to applications within the template itself, automating the application association during provisioning.
- **Attribute Groups:** Your enterprise creates and manages attributes that capture the application metadata that are important to your enterprise. Application attributes support an open JSON schema, providing you the flexibility you need to capture the complexity of your enterprise metadata taxonomy. Application attributes might include items such as the application security classification, organizational ownership, application type, cost center, and support information. Builders associate the necessary attributes to their applications. When attributes are updated, this is automatically reflected within all associated applications.

172.1.2. Benefits

- **Ensure compliance with corporate standards:** AWS Service Catalog provides a single location where organizations can centrally manage catalogs of IT services. With AWS Service Catalog you can control which IT services and versions are available, what is configured in each of the available service, and who gets permission access by individual, group, department or cost center.
- **Help employees quickly find and deploy approved IT services:** With AWS Service Catalog, you define your own catalog of AWS services and AWS Marketplace software, and make them available for your organization. Then, end users can quickly discover and deploy IT services using a self-service portal.
- **Centrally manage IT service lifecycle:** AWS Service Catalog enables you to add new versions of IT services, and end users are notified so they can keep abreast of the latest updates. With AWS Service Catalog you can control the use of IT services by specifying constraints, such as limiting the AWS regions in which a product can be launched.
- **Connect with ITSM/ITOM software:** The AWS Service Management Connector helps IT Service Management (ITSM) administrators improve governance over provisioned AWS and third-party products. ITSM tools, such as ServiceNow and Jira Service Desk, connect with the AWS Management and Governance services AWS Service Catalog, AWS Config, and AWS Systems Manager. ITSM users to can connect to AWS Service Catalog to request, provision, and manage AWS and third-party services and resources.
- **Manage all of your application information on AWS:** AWS Service Catalog AppRegistry provides a single repository for collecting and managing your application resources on AWS. You define your application metadata, which may include information from your internal systems, other AWS services, and software vendors.

Builders can include a reference to their application within the infrastructure code, and business stakeholders have up-to-date information on application contents and metadata, such as organizational ownership, data sensitivity, and cost center.

172.2. Backup/Restore and Disaster Recovery

This requirement is not applicable for this service. For additional information beyond what is described herein, please refer to http://aws.amazon.com/documentation/.

172.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

172.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/servicecatalog/
- **Service quotas:** https://docs.aws.amazon.com/servicecatalog/latest/adminguide/limits.html
- **Service FAQs:** https://aws.amazon.com/servicecatalog/faqs/

172.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/servicecatalog/ and the following links for comprehensive technical documentation regarding this service.

- Administrator Guide: Provides a conceptual overview of AWS Service Catalog and includes detailed instructions for using the service as an administrator.
- User Guide: Provides a conceptual overview of AWS Service Catalog and includes detailed instructions for using the service as an end user.
- Developer Guide: Provides a conceptual overview of AWS Service Catalog, includes detailed instructions for using the various features, and provides a complete API reference for developers.
- AWS CLI Reference: Describes the AWS CLI commands for AWS Service Catalog.

173. AWS Shield

173.1. Service Overview

AWS Shield is a managed Distributed Denial of Service (DDoS) protection service that safeguards applications running on AWS. AWS Shield provides always-on detection and automatic inline mitigations that minimize application downtime and latency, so there is no need to engage AWS Support to benefit from DDoS protection. There are two tiers of AWS Shield - Standard and Advanced.

All AWS customers benefit from the automatic protections of AWS Shield Standard, at no additional charge. AWS Shield Standard defends against most common, frequently occurring network and transport layer DDoS attacks that target your web site or applications. When you use AWS Shield Standard with Amazon CloudFront and Amazon Route 53, you receive comprehensive availability protection against all known infrastructure (Layer 3 and 4) attacks.

For higher levels of protection against attacks targeting your applications running on Amazon Elastic Compute Cloud (EC2), Elastic Load Balancing (ELB), Amazon CloudFront, AWS Global Accelerator and Amazon Route 53 resources, you can subscribe to AWS Shield Advanced. In addition to the network and transport layer protections that come with Standard, AWS Shield Advanced provides additional detection and mitigation against large and sophisticated DDoS attacks, near real-time visibility into attacks, and integration with AWS WAF, a web application firewall. AWS Shield Advanced also gives you 24x7 access to the AWS Shield Response Team (SRT) and protection against DDoS related spikes in your Amazon Elastic Compute Cloud (EC2), Elastic Load Balancing (ELB), Amazon CloudFront, AWS Global Accelerator and Amazon Route 53 charges.

AWS Shield Advanced is available globally on all Amazon CloudFront, AWS Global Accelerator, and Amazon Route 53 edge locations. You can protect your web applications hosted anywhere in the world by deploying Amazon CloudFront in front of your application. Your origin servers can be Amazon S3, Amazon Elastic Compute Cloud (EC2), Elastic Load Balancing (ELB), or a custom server outside of AWS. You can also enable AWS Shield Advanced directly on an Elastic IP or Elastic Load Balancing (ELB) in the following AWS Regions - Northern Virginia, Ohio, Oregon, Northern California, Montreal, São Paulo, Ireland, Frankfurt, London, Paris, Stockholm, Singapore, Tokyo, Sydney, Seoul, and Mumbai.

173.1.1. Features

- **Static threshold DDoS protection for underlying AWS services:** AWS Shield Standard provides always-on network flow monitoring, which inspects incoming traffic to AWS services and applies a combination of traffic signatures, anomaly algorithms, and other analysis techniques to detect malicious traffic in real time. Shield Standard sets static thresholds for each AWS resource type but doesn't provide custom protections to your applications.
- **Inline attack mitigation:** Automated mitigation techniques are built into AWS Shield Standard, giving underlying AWS services protection against common, frequently occurring infrastructure attacks. Automatic mitigations are applied inline to protect AWS services, so there is no latency impact. Shield Standard uses techniques such as deterministic packet filtering and priority-based traffic shaping to automatically mitigate basic network layer attacks.
- **Tailored detection based on application traffic patterns:** AWS Shield Advanced provides customized detection based on traffic patterns to your protected Elastic IP address, ELB, CloudFront, Global Accelerator, and Route 53 resources. Using additional region- and resource-specific monitoring techniques, Shield Advanced detects and alerts you of smaller DDoS attacks. Shield Advanced also detects application layer attacks such as HTTP floods or DNS query floods by baselining traffic on your application and identifying anomalies.
- **Health-based detection:** AWS Shield Advanced uses the health of your applications to improve responsiveness and accuracy in attack detection and mitigation. You can define a health check in Route 53 and associate it with a resource that is protected by Shield Advanced through the console or API. This allows Shield Advanced to detect attacks impacting the health of your application more quickly and at lower traffic thresholds, improving the DDoS resiliency of your application and preventing false positive notifications. Resource health status is also available to the SRT so they can appropriately prioritize response to unhealthy applications. You can apply health-based detection to all resource types that Shield Advanced supports: Elastic IP, ELB, CloudFront, Global Accelerator, and Route 53.

- **Advanced attack mitigation:** AWS Shield Advanced provides more sophisticated automatic mitigations for attacks targeting your applications running on protected EC2, ELB, CloudFront, Global Accelerator, and Route 53 resources. Using advanced routing techniques, Shield Advanced automatically deploys additional mitigation capacity to protect your application against DDoS attacks. For customers with Business or Enterprise support, the SRT also applies manual mitigations for more complex and sophisticated DDoS attacks that might be unique to your application. For application layer attacks, you can use AWS WAF at no additional charge for resources protected by Shield Advanced to set up proactive rules (such as rate-based blocking to automatically block web requests from attacking source IP addresses) or respond immediately to incidents as they happen. You can also engage directly with the SRT to place custom AWS WAF rules on your behalf in response to an application layer DDoS attack. The SRT will diagnose the attack and, with your permission, apply mitigations on your behalf, reducing the amount of time your applications might be impacted by an ongoing DDoS attack.
- **Automatic application layer DDoS mitigation:** AWS Shield Advanced can automatically protect web applications by mitigating application layer (L7) DDoS events with no manual intervention needed by you or the AWS SRT. Shield Advanced can create WAF rules in your WebACLs to automatically mitigate an attack, or you can activate them in count-only mode. This lets you quickly respond to DDoS events to prevent application downtime due to an application layer DDoS attack.
- **Proactive event response:** AWS Shield Advanced offers proactive engagement from the SRT when a DDoS event is detected. When you activate proactive engagement, the SRT will directly contact you if a Route 53 health check associated with your protected resource becomes unhealthy during a DDoS event. This allows you to engage with experts more quickly when the availability of your application is affected by a suspected attack. You can receive proactive engagement for network layer and transport layer events on Elastic IP addresses and Global Accelerator accelerators, and for application layer attacks on CloudFront distributions and Application Load Balancers.
- **Protection groups:** AWS Shield Advanced allows you to bundle resources into protection groups, giving you a self-service way to customize the scope of detection and mitigation for your application by treating multiple resources as a single unit. Resource grouping improves the accuracy of detection, reduces false positives, eases automatic protection of newly created resources, and accelerates the time to mitigate attacks against multiple resources. For example, if an application consists of four CloudFront distributions, you can add them to one protection group to receive detection and protection for the collection of resources as a whole. Reporting can also be consumed at the protection group level, giving a more holistic view of overall application health.
- **Visibility and attack notification:** AWS Shield Advanced gives you complete visibility into DDoS attacks with near real-time notification through Amazon CloudWatch and detailed diagnostics on the AWS WAF and AWS Shield console or APIs. You can also view a summary of prior attacks from the console.
- **DDoS cost protection:** AWS Shield Advanced comes with DDoS cost protection to safeguard against scaling charges resulting from DDoS-related usage spikes on protected EC2, ELB, CloudFront, Global Accelerator, and Route 53 resources. If any of these protected resources scale up in response to a DDoS attack, you can request Shield Advanced service credits through your regular AWS Support channel.

173.1.2. Benefits

- **Seamless integration and deployment:** Your AWS resources automatically have AWS Shield Standard and are protected from common, most frequently occurring network and transport layer DDoS attacks. You can achieve a higher level of defense by simply enabling AWS Shield Advanced protection for Elastic IP, Elastic Load Balancing (ELB), Amazon CloudFront, AWS Global Accelerator, or Amazon Route 53 resources you want to protect using the management console or APIs. No routing changes are required for enabling these protections.
- **Customizable protection:** With AWS Shield Advanced, you have the flexibility to choose the resources to protect for infrastructure (Layer 3 and 4) protection. You can write customized rules with AWS WAF to mitigate sophisticated application layer attacks. These customizable rules can be deployed instantly, allowing you to quickly mitigate attacks. You can set up rules proactively to automatically block bad traffic, or respond to incidents as they occur. You also have 24x7 access to the AWS Shield Response Team (SRT), who can write rules on your behalf to mitigate application layer DDoS attacks.
- **Managed Protection and Attack Visibility:** With AWS Shield Standard you get always-on heuristics-based network flow monitoring and inline mitigation against common, most frequently occurring network and transport layer DDoS attacks. AWS Shield Advanced provides enhanced resource specific detection and employs advanced mitigation and routing techniques for sophisticated or larger attacks. You also get 24x7 access to the AWS Shield Response Team (SRT) for manual mitigation of edge cases affecting your availability. AWS Shield Advanced also provides visibility and insights into all your DDoS incidents through AWS CloudWatch metrics and attack diagnostics. Finally, you can also see the DDoS threat environment on AWS with the Global threat environment dashboard.
- **Cost Efficient:** With AWS Shield Standard is automatically enabled for all AWS customers at no additional cost. With AWS Shield Advanced, customers get AWS WAF and AWS Firewall Manager at no additional cost for usage on resources protected by AWS Shield Advanced as described on the Shield pricing page. Additionally, you get "DDoS cost protection for scaling", a feature that protects your AWS bill from usage spikes on your AWS Shield Advanced protected EC2, Elastic Load Balancing (ELB), Amazon CloudFront, AWS Global Accelerator, and Amazon Route 53 resources as a result of a DDoS attack.

173.2. Backup/Restore and Disaster Recovery

This requirement is not applicable for this service. For additional information beyond what is described herein, please refer to http://aws.amazon.com/documentation/.

173.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

173.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/shield/
- **Service quotas:** https://docs.aws.amazon.com/general/latest/gr/shield.html

- **Service FAQs:** https://aws.amazon.com/shield/faqs/

173.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/shield/ the following links for comprehensive technical documentation regarding this service.

- Developer Guide: Describes how to get started with AWS Shield Advanced. Explains key concepts, and provides step-by-step instructions that show you how to use the features.
- API Reference: Describes all the API operations for AWS Shield Advanced in detail.

174. AWS Single Sign-On

174.1. Service Overview

AWS Single Sign-On (AWS SSO) is where you create, or connect, your workforce identities in AWS once and manage access centrally across your AWS organization. You can choose to manage access just to your AWS accounts or cloud applications. You can create user identities directly in AWS SSO, or you can bring them from your Microsoft Active Directory or a standards-based identity provider, such as Okta Universal Directory or Azure AD. With AWS SSO, you get a unified administration experience to define, customize, and assign fine-grained access. Your workforce users get a user portal to access all of their assigned AWS accounts, Amazon EC2 Windows instances, or cloud applications. AWS SSO can be flexibly configured to run alongside or replace AWS account access management via AWS IAM.

It's easy to get started with AWS SSO. With just a few clicks in the management console, you can connect AWS SSO to your existing identity source and configure permissions that grant users access to their assigned AWS accounts, cloud applications, and other SAML-based applications that you add to AWS SSO.

174.1.1. Features

Copy and paste features from service features page

- **Integrated with AWS Organizations:** AWS SSO is integrated with AWS Organizations, enabling you to select one or more accounts from your organization and grant users access to these accounts. AWS SSO builds on AWS Identity and Access Management (IAM) roles and policies to help you manage access centrally across all AWS accounts in your AWS organization. No additional configuration is required in the individual accounts. With just a few clicks, you can grant users access to all of the AWS accounts being used for an application or by a team.
- **Manage SSO access for multiple AWS accounts:** Using AWS Single Sign-On (SSO), you can manage SSO access for multiple AWS accounts centrally. When users sign in to their personalized user portals, they will see all of their assigned roles in AWS accounts in one place.
- **Enable SSO access to your Amazon EC2 Windows instances:** By using AWS SSO, you can provide one-click login access to your Amazon EC2 Windows instances from within the AWS Systems Manager Fleet Manager console. This makes it easy for you to access your instance desktops from anywhere without having to enter your credentials multiple times or having to configure remote access client software.
- **Attribute-based access control:** AWS SSO makes it easy for you to create and use fine-grained permissions for your workforce based on user attributes defined in your

AWS SSO identity source. AWS SSO allows you to select multiple attributes, such as cost center, title, or locale, and then use them for attribute-based access control (ABAC) to simplify and centralize your access administration. You can define permissions once for your entire AWS organization, and then grant, revoke, or modify AWS access by simply changing the attributes in the identity source.

- **Create and manage users in AWS SSO:** AWS SSO provides you a directory by default that you can use to create users and organize them in groups within AWS SSO. You can create users in AWS SSO by configuring their email address and name. When you create a user, by default AWS SSO sends an email to the user so that your users can set their own password. Within minutes, you can grant your users and groups permissions to AWS resources in all your AWS accounts as well as many business applications. Your users sign in to a user portal with credentials they configured in AWS SSO to access all of their assigned accounts and applications in a sin
- **Connect and automatically provision users from standards-based identity providers:** You can connect AWS SSO to Okta Universal Directory, Azure AD, or another supported identity provider (IdP) via Security Assertion Markup Language (SAML) 2.0 so your users can sign in with their existing credentials. And, AWS SSO also supports System for Cross-domain Identity Management (SCIM) for automation of user provisioning. You can manage your users in your IdP, get them into AWS quickly, and centrally manage their access to all AWS accounts and business applications. AWS SSO also allows you to select multiple user attributes, such as cost center, title, or locale, from your Okta Universal Directory, and then use them for ABAC to simplify and centralize your access administration.
- **Multi-factor authentication:** With AWS SSO, you can use standards-based strong authentication capabilities for all your users across all your identity sources. If you use a supported SAML 2.0 IdP as your identity source, you can enable multi-factor authentication (MFA) capabilities of your provider. When using Active Directory or AWS SSO as your identity source, AWS SSO supports the Web Authentication specification to help you secure user access to AWS accounts and business applications using with FIDO-enabled security keys, such as YubiKey, and built-in biometric authenticators, such as Touch ID on Apple MacBooks and facial recognition on PCs. You can also enable one-time-passwords (TOTPs) using authenticator apps such as Google Authenticator or Twilio Authy. AWS SSO allows you to enforce MFA for all your users, including the requirement for the users to set up MFA devices during sign-in.
- **User portal:** With AWS SSO, users can find and access all of their assigned accounts and applications in one place. Users can simply sign in to their personalized user portal with their existing corporate credentials and with one click access any of their assigned accounts and applications. The user portal also helps you roll out access to new applications more easily by helping users discover new applications in their user portal.
- **Support for browser, command line, and mobile interfaces:** When users sign in through the AWS Command Line Interface (CLI), they can use their existing corporate credentials and get consistent authentication experience, while getting the benefits of automated short-term credential management. Once signed in, developers can see their AWS SSO assigned accounts and roles, and they can also create profiles that let them switch between roles and accounts in a single command. AWS Mobile Console app also supports AWS SSO so you get a consistent sign-in experience across browser, mobile, and command line interfaces.

- **Built-in SSO integrations to business applications:** AWS SSO offers you built-in SSO integrations to many business applications, including Salesforce, Box, and Microsoft 365. You can easily configure SSO access to these applications by following step by step instructions. AWS SSO guides you through entering the required URLs, certificates, and metadata. For a full list of business applications pre-integrated with AWS SSO, see AWS SSO Cloud Applications.

174.1.2. Benefits

- **Central place to create or connect your identities:** You have the option to create your users' identities and groups in AWS SSO. Or, you can connect to your existing users and groups from Microsoft Active Directory Domain Services, Okta Universal Directory, Azure AD, or another standards-based identity provider. In either case, you manage and authenticate users where you want and AWS SSO authorizes access to the AWS accounts, cloud applications, and other SAML-based applications that you add to AWS SSO.
- **Manage access to multiple AWS accounts from one place:** With AWS Organizations integration, AWS SSO enables you to manage access across multiple accounts with no additional setup within individual accounts. You can assign user permissions based on common job functions, customize them to meet your specific security requirements, and assign fine-grained permissions within the specific accounts where they need access. AWS SSO also allows you to utilize user attributes, such as cost center, title, or locale, for attribute-based access control (ABAC).
- **Manage access to your cloud applications:** With AWS Single Sign-On, you can easily control who has access to your cloud applications. Your users can utilize their directory credentials to sign in to their AWS SSO web user portal and get one-click access to their assigned applications like Amazon SageMaker Studio, AWS Systems Manager Change Manager, and standards-based cloud applications including Salesforce, Box, and Microsoft 365.

174.2. Backup/Restore and Disaster Recovery

This requirement is not applicable for this service. For additional information beyond what is described herein, please refer to http://aws.amazon.com/documentation/.

174.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

174.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/singlesignon/index.html
- **Service quotas:** https://docs.aws.amazon.com/singlesignon/latest/userguide/limits.html
- **Service FAQs:** https://aws.amazon.com/single-sign-on/faqs/

174.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/singlesignon/index.html and the following links for comprehensive technical documentation regarding this service.

- AWS SSO User Guide: Introduces you to AWS SSO, helps you set up your AWS accounts and cloud applications for SSO access, and shows you how to audit and monitor user sign-ins.
- AWS SSO API Reference: Describes the API operations for AWS SSO.
- AWS SSO SCIM Implementation Developer Guide: Documents the implementation of the System for Cross-domain Identity Management (SCIM) v2.0 protocol in AWS SSO.
- AWS SSO Identity Store API Reference: Describes the API operations for the AWS SSO Identity Store service.
- AWS SSO OIDC API Reference: Describes the API operations for the AWS SSO OpenID Connect (OIDC) service.
- AWS SSO Portal API Reference: Describes the API operations for the AWS SSO Portal service.

175. AWS Snowball

175.1. Service Overview

AWS Snowball, a part of the AWS Snow Family, is an edge computing, data migration, and edge storage device that comes in two options. Snowball Edge Storage Optimized devices provide both block storage and Amazon S3-compatible object storage, and 40 vCPUs. They are well suited for local storage and large scale-data transfer. Snowball Edge Compute Optimized devices provide 52 vCPUs, block and object storage, and an optional GPU for use cases like advanced machine learning and full motion video analysis in disconnected environments. You can use these devices for data collection, machine learning and processing, and storage in environments with intermittent connectivity (like manufacturing, industrial, and transportation) or in extremely remote locations (like military or maritime operations) before shipping them back to AWS. These devices may also be rack mounted and clustered together to build larger temporary installations.

Snowball supports specific Amazon EC2 instance types and AWS Lambda functions, so you can develop and test in the AWS Cloud, then deploy applications on devices in remote locations to collect, pre-process, and ship the data to AWS. Common use cases include data migration, data transport, image collation, IoT sensor stream capture, and machine learning.

175.1.1. Features

- **Fast data transfer:** Snowball Edge devices feature high-speed network connections, supporting 10 Gbps to 100 Gbps links with RJ45, SFP+ and QSFP+ copper, and optical interfaces. All encryption is performed on the device itself, helping enable a higher data throughput rate and shorter data transfer times. For device specific networking specifications, please see Snowball Edge documentation.
- **AWS OpsHub for simple management and monitoring:** AWS OpsHub is a graphical user interface that makes it easy to set up and manage AWS Snowball devices enabling you to rapidly deploy edge computing workloads and simplify data migration to the cloud. You can download and install AWS OpsHub on any Windows or Mac client machine, such as a laptop.
- **GPU support:** Snowball Edge Compute Optimized provides an optional NVIDIA Tesla V100 GPU along with Amazon EC2 instances to accelerate an application's performance in disconnected environments. Using the GPU option, you can run

applications such as advanced machine learning and full motion video analysis in environments with little or no connectivity.

- **Clustering:** You can cluster multiple Snowball Edge devices when running edge computing jobs to create a local storage tier with increased durability for your on-premises applications. When creating a new job in the Console, select the option to make a cluster. In the event of a device failure, a replacement device can be ordered easily through the Console. This functionality is available for local storage and compute jobs and is not enabled for data transfer jobs.
- **S3-compatible endpoint for object storage:** Applications can work with Snowball Edge object storage through an S3-compatible endpoint accessed through the S3 SDK or CLI. For specific information, see the API documentation.
- **Block storage:** You can run block storage on both Snowball Edge Compute Optimized and Snowball Edge Storage Optimized devices. You attach block storage volumes to Amazon EC2 instances using a subset of the Amazon EBS API that enable you to configure and manage volumes for EC2 instances on Snowball Edge devices. This makes it easier to develop applications in EC2, and then run them in disconnected and remote locations. Snowball Edge supports both performance optimized and capacity optimized volume types.
- **NFS endpoint:** Applications can work with Snowball Edge as a NFS mount point. NFS v3 and v4.1 are supported so you can easily use Snowball Edge with your existing on-premises servers and file-based applications. Using the NFS interface allows simple file transfer to a Snowball Edge device when the S3 adapter is not feasible. The file system metadata is preserved until the files are converted into objects when they are transferred into your S3 bucket.
- **Encryption:** All data transferred to AWS Snowball is automatically encrypted with 256-bit encryption keys that are managed by the AWS Key Management Service (KMS). The encryption keys are never stored on the device to help ensure your data stays secure during transit.
- **Rugged and portable:** AWS Snowball Edge devices have a ruggedized case designed for durability and portability. A device weighs less than 50 pounds and can be moved by a single person.
- **Tamper evident:** Snowball Edge devices feature a Trusted Platform Module (TPM) that provides a hardware root of trust. The TPM also provides interfaces to the trusted software stack during the measurements and verification of the boot environment integrity after the power is switched on, and before the Snowball Edge device is ready to be used. AWS also uses additional tamper-indicating inspection processes after each device is received back to the AWS Region. This helps to ensure the integrity of the AWS Snowball Edge device, and with the AWS Snowball service's encryption features, it helps preserve the confidentiality of your data.

175.1.2. Benefits

- **Easy data movement:** Snowball moves terabytes of data in about a week. You can use it to move things like databases, backups, archives, healthcare records, analytics datasets, IoT sensor data and media content, especially when network conditions prevent realistic timelines for transferring large amounts of data both into and out of AWS.
- **Simple to use:** Jobs are created in the AWS Management Console. Once a job is created, AWS automatically ships a pre-provisioned Snowball Edge device to your

location. When you receive the device, simply attach it to your local network and connect your applications. Once the device is ready to be returned, the E Ink shipping label automatically updates, and your freight carrier transports it to the correct AWS facility where the upload begins. Job status can be tracked via Amazon SNS-generated text or email messages or directly in the AWS Management Console.

- **Process & analyze data locally:** Run EC2 AMIs and deploy AWS Lambda code on Snowball Edge devices to run local processing or analysis with machine learning or other applications. You can run applications directly on the device as a consistent AWS environment without network connectivity. This capability helps you develop your machine learning and analysis tools and test them in the cloud, but operate them in locations with limited or non-existent network connections before shipping the data back to AWS.
- **Stand-alone storage:** Snowball Edge devices can provide local storage to existing on-premises applications through a file sharing protocol (NFS) or object storage interface (the S3 API). Additionally, you can use on-board block storage volumes for applications running on Amazon EC2 instances on the Snowball. You can also cluster Snowball Edge devices together into a single, larger storage tier with increased durability. If a Snowball Edge device needs to be replaced, it can be removed from the cluster and replaced with a new device.
- **Secure:** Snowball Edge devices use tamper-evident enclosures, 256-bit encryption, and industry-standard Trusted Platform Modules (TPM) designed to ensure both security and full chain-of-custody for your data. Encryption keys are managed with the AWS Key Management Service (KMS) and they are never stored on the device.
- **Scalable:** Snowball can transport multiple terabytes of data and multiple devices can be used in parallel or clustered together to transfer petabytes of data into or out of AWS. Snowball is currently available in select regions and your location will be verified once you create a job in the AWS Management Console.

175.2. Backup/Restore and Disaster Recovery

This requirement is not applicable for this service. For additional information beyond what is described herein, please refer to http://aws.amazon.com/documentation/.

175.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

175.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/snowball/index.html
- **Service quotas:** https://docs.aws.amazon.com/snowball/latest/ug/limits.html
- **Service FAQs:** https://aws.amazon.com/snowball/faqs/

175.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/snowball/index.html and the following links for comprehensive technical documentation regarding this service.

- AWS Snowball User Guide: Provides a conceptual overview of how to use AWS Snowball with a Snowball device, including guidance for importing and exporting data into the AWS Cloud, shipping considerations, and other features of a Snowball device.
- AWS Snow Family API Reference: Describes all the API operations for AWS Snow Family in detail. Also provides sample requests, responses, and errors for the supported web service protocols.

176. AWS Snowcone

176.1. Service Overview

AWS Snowcone is the smallest member of the AWS Snow Family of edge computing, edge storage, and data transfer devices. Weighing in at 4.5 pounds (2.1 kg), AWS Snowcone is equipped with 8 terabytes of usable storage, while AWS Snowcone Solid State Drive (SSD) support 14 terabytes of usable storage. Both referred to as Snowcone, the device is ruggedized, secure, and purpose-built for use outside of a traditional data centre. Its small form factor makes it a perfect fit for tight spaces or where portability is a necessity and network connectivity is unreliable. You can use Snowcone in backpacks on first responders, or for IoT, vehicular, and drone use cases. You can execute compute applications at the edge, and you can ship the device with data to AWS for offline data transfer, or you can transfer data online with AWS DataSync from edge locations.

Like AWS Snowball, Snowcone has multiple layers of security and encryption. You can use either of these services to run edge computing workloads, or to collect, process, and transfer data to AWS. Snowcone is designed for data migration needs up to 8 terabytes per device and from space-constrained environments where AWS Snowball devices will not fit.

176.1.1. Features

Copy and paste features from service features page

- **Small and light:** AWS Snowcone weighs about 4.5 lbs. (2.1 kg) and is roughly the size of a tissue box, approximately 9 inches long, 6 inches wide, and 3 inches tall (8.94” x 5.85” x 3.25” / 227 mm x 148.6 mm x 82.65 mm). In comparison, AWS Snowball weighs about 49.7 lbs., is 28.3 inches long, 10.6 inches wide, and 15.5 inches tall.
- **Ruggedized enclosure:** AWS Snowcone is designed to meet stringent standards for ruggedization, including ISTA-3A, ASTM D4169, and MIL-STD-810G for free-fall shock, operational vibration, and more. It is designed to tolerate falls up to up to 3.8 feet (1.15 meters). It also meets the IP65 International Protection Marking IEC standard, meaning it is both dust-tight – allowing no dust inside the enclosure when sealed – and water resistant, including protection from water jets on all sides. The device has a wide operating temperature range from freezing (0ºC/32ºF) to desert-like conditions (38ºC/100ºF). When in storage or being shipped, Snowcone withstands even harsher temperatures (-32ºC/-25.6ºF to 63ºC/145.4ºF).
- **Encryption:** Like other AWS Snow Family products, all data on AWS Snowcone is always automatically encrypted using 256-bit keys that you manage by using the AWS Key Management Service (KMS). Encryption keys are never stored on the device. This helps ensure that your data stays secure during device transit.
- **Anti-tamper & Tamper-evident:** The AWS Snowcone device has a Trusted Platform Module (TPM) that provides a hardware root of trust. The TPM also provides interfaces to the trusted software stack during the measurements and verification of the boot environment integrity after the power is switched on and before the AWS Snowcone

device is ready to be used. AWS also uses additional tamper-indicating inspection processes after each AWS Snowcone device is received back to the AWS Region. This helps to ensure the integrity of the AWS Snowcone device, and with the Snowcone service's encryption features, it helps preserve the confidentiality of your data.

- **Easy setup and management with AWS OpsHub:** AWS OpsHub is an application for managing the AWS Snow Family devices, including AWS Snowcone. The OpsHub graphical user interface (GUI) makes it easy to setup and manage AWS Snowcone devices so you can rapidly deploy edge computing workloads and migrate data to the cloud – even when you don't have an internet connection. With just a few clicks in OpsHub, you can unlock and configure a Snowcone, drag-and-drop data to it, launch applications, or monitor the device. Download AWS OpsHub here and install it on Windows, Mac, or Linux client machines, such as a laptop. There is no cost to use OpsHub. For more information, refer to the AWS OpsHub documentation.
- **On-board computing:** AWS Snowcone has computing resources to collect and process data with AWS services at the edge. It runs specific Amazon EC2 instances with 2 available CPUs and 4 GB of available memory to support your applications and AWS IoT Greengrass functions. During the ordering process, you can select your EC2 AMIs that you want to install on the device, including AMIs running IoT Greengrass. AWS will load the AMIs onto the device and ship it to you.
- **Offline data transfer:** Like any Snow Family device, AWS Snowcone can efficiently move data by physically shipping the device with your data to your specified AWS Region where data will be put into your specified Amazon S3 bucket. Offline data export from your S3 bucket feature is arriving soon on Snowcone. S3 buckets where data will be imported are setup before you order the device.
- **Online data transfer with AWS DataSync:** AWS DataSync comes pre-installed on AWS Snowcone to enable simple, fast, secure, and cost-effective online data import or export between storage on Snowcone devices at edge locations and Amazon S3, Amazon EFS, or Amazon FSx for Windows File Server. DataSync automatically handles moving data, scheduling transfer jobs, monitoring, encryption, data verification, and network optimization.
- **File-based storage:** AWS Snowcone features 8 TB of usable storage. Applications can work with Snowcone as an NFS v4 mount point. You can easily use Snowcone with your existing on-premises servers and file-based applications to read and write data on the device.
- **Wired and wireless networking:** You can connect AWS Snowcone to wired or Wi-Fi networks. For wired networks, Snowcone provides 2 ports that auto-negotiate for 1 Gb or 10 Gb Ethernet networks.

176.1.2. Benefits

Copy and paste benefits from service landing page

- **Security:** AWS Snowcone uses both hardware and software to provide security that satisfies even the most stringent requirements. It uses hardware-based Trusted Platform Modules (TPM) to store device-specific keys where they are inaccessible to software, helping ensure the integrity of the device. Data stored on the Snowcone is encrypted using two layers of at-rest encryption, which helps to protect the data stored on the device while it is being shipped. Encryption keys are managed with AWS Key Management Service (KMS), and are never stored on the Snowcone device.

- **Portable edge computing:** AWS Snowcone deploys virtually anywhere you need it. It features 2 vCPUs, 4 GB of memory, 8 TB of usable storage, wired access, and USB-C power using a cord or optional battery. You can put it in a messenger bag, run it in an autonomous vehicle or an airplane, or even attach it to a drone.
- **Withstands harsh environments:** AWS Snowcone is built for edge computing and data storage outside of a data center. It is designed to meet stringent standards for ruggedization, including free-fall shock, operational vibration, and more. When sealed, the device is both dust-tight and water-resistant. Snowcone has a wide operating temperature range from freezing to desert-like conditions, and withstands even harsher temperatures in storage.
- **Flexible data transfer & networking:** AWS Snowcone provides options to fit your edge computing, data transfer, and edge storage use case. Using Snowcone's wired 10 GbE networking, you can collect data at the edge, process it with Amazon EC2 instance AMIs, and then move the data to AWS based on your needs. You can ship the device with data to AWS for offline data transfer, or you can transfer data online with AWS DataSync.
- **Works with IoT sensors:** With support for wired connectivity and local compute, you can use AWS Snowcone as an IoT hub, data aggregation point, application monitor, or a lightweight analytics engine. With Snowcone's ruggedization, you can put it nearly anywhere your networked devices need to be.
- **Consistent experience:** develop in the cloud, compute at the edge: With Amazon EC2 instances, you can develop and test functions and applications in the cloud, and then deploy them rapidly at the edge with AWS Snowcone. When you need to update your application, simply ship a new Snowcone on-site in a few days, and swap it in.

176.2. Backup/Restore and Disaster Recovery

This requirement is not applicable for this service. For additional information beyond what is described herein, please refer to http://aws.amazon.com/documentation/.

176.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

176.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/snowball/index.html
- **Service quotas:** https://docs.aws.amazon.com/snowball/latest/snowcone-guide/snowcone-limits.html
- **Service FAQs:** https://aws.amazon.com/snowcone/faqs/

176.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/snowball/index.html and the following links for comprehensive technical documentation regarding this service.

- AWS Snowcone User Guide: Provides a conceptual overview of how to use an AWS Snowcone device. Includes guidance for storage and local compute, importing and exporting data into the AWS Cloud, and other features of a Snowcone device. Use AWS

Snowcone to collect, process, and move data to AWS, either offline, or online when an internet connection is available.

- AWS Snow Family API Reference: Describes all the API operations for AWS Snow Family in detail. Also provides sample requests, responses, and errors for the supported web service protocols.

177. AWS Step Functions

177.1. Service Overview

AWS Step Functions is a low-code, visual workflow service that developers use to build distributed applications, automate IT and business processes, and build data and machine learning pipelines using AWS services. Workflows manage failures, retries, parallelization, service integrations, and observability so developers can focus on higher-value business logic.

177.1.1. Features

- **Workflow configuration:** Using AWS Step Functions, you define your workflows as state machines, which transform complex code into easy to understand statements and diagrams. Building apps, and confirming that they are implementing your desired functionality, is quicker and easier. Read more about **how Step Functions works**.
- **Built-in service primitives:** AWS Step Functions provides ready-made steps for your workflow called states that implement basic service primitives for you, which means you can remove that logic from your application. States can pass data to other states and microservices, handle exceptions, add timeouts, make decisions, execute multiple paths in parallel, and more. Lear more about **states**.
- **AWS service integrations:** Using AWS Step Functions service tasks, you can configure your Step Functions workflow to call other AWS services. This includes compute services (AWS Lambda, Amazon ECS, Amazon EKS, and AWS Fargate), database services (Amazon DynamoDB), messaging services (Amazon SNS and Amazon SQS), data processing and analytics services (Amazon Athena, AWS Batch, AWS Glue, Amazon EMR, and AWS Glue DataBrew), machine learning services (Amazon SageMaker), and APIs created by Amazon API Gateway. Learn more about **service tasks**.
- **Coordination of distributed components:** AWS Step Functions can coordinate any application that can make an HTTPS connection, regardless of where it is hosted—for example, on Amazon EC2 instances, mobile devices, or on-premises servers. Using Step Functions, you can quickly create distributed applications that leverage AWS services as well as your own microservices. Learn more about **activity tasks**.
- **Component reuse:** AWS Step Functions coordinates your existing Lambda functions and microservices into robust applications, and lets you quickly rewire them into new compositions. The tasks in your workflow can run anywhere, including on instances, containers, functions, and mobile devices. Learn how to **reuse existing application components**.
- **Workflow abstraction:** AWS Step Functions keeps the logic of your application strictly separated from the implementation of your application. You can add, move, swap, and reorder steps without having to make changes to your business logic. Through this separation of concerns, your workflows gain modularity, simplified maintenance, scalability, and code reuse.

- **State management:** AWS Step Functions maintains the state of your application during execution, including tracking what step of execution it is in, and storing data that is moving between the steps of your workflow. This means you don't have to manage state yourself with data stores or by building complex state management into all of your tasks.
- **Built-in error handling:** AWS Step Functions automatically handles errors and exceptions with built-in try/catch and retry, whether the task takes seconds or months to complete. You can automatically retry failed or timed-out tasks, respond differently to different types of errors, and recover gracefully by falling back to designated cleanup and recovery code. Learn more about **Step Functions error handling** and how you can **handle error conditions using a state machine**.
- **History of each execution:** AWS Step Functions delivers real-time diagnostics and dashboards, integrates with Amazon CloudWatch and AWS CloudTrail, and logs every execution, including overall state, failed steps, inputs, and outputs. If things go wrong, you can quickly identify not only where, but why, and quickly troubleshoot and remediate failures. Learn more about Step Functions **monitoring and logging**.
- **Visual monitoring:** Launching an application is as simple as pressing a button, then watching the steps execute visually, so you can quickly verify that everything is operating in order – and as expected. The console clearly highlights errors, so you can quickly pinpoint their root-cause, and troubleshoot issues.

177.1.2. Benefits

- **Build and deploy rapidly:** Get started quickly with Workflow Studio, a simple drag-and-drop interface. With Step Functions, you can express complex business logic as low-code, event-driven workflows that connect services, systems or people within minutes.
- **Write less integration code:** Compose AWS resources from over 200 services including Lambda, ECS, Fargate, Batch, DynamoDB, SNS, SQS, SageMaker, EventBridge, or EMR into resilient business workflows, data pipelines, or applications.
- **Build fault-tolerant and stateful workflows:** Step Functions manages state, checkpoints, and restarts for you to make sure that your workflows run in order and as expected. Built-in try/catch, retry, and rollback capabilities deal with errors and exceptions automatically based on your defined business logic.
- **Designed for reliability and scale:** Step Functions offers two workflow types - Standard or Express - that can be used depending on your specific use case. Standard Workflows are used to manage long-running workloads. Express Workflows support high-volume event processing workloads.

177.2. Backup/Restore and Disaster Recovery

This requirement is not applicable for this service. For additional information beyond what is described herein, please refer to http://aws.amazon.com/documentation/.

177.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

177.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/step-functions/index.html

- **Service quotas:** https://docs.aws.amazon.com/step-functions/latest/dg/limits-overview.html
- **Service FAQs:** https://aws.amazon.com/step-functions/faqs/

177.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/step-functions/index.html and the following links for comprehensive technical documentation regarding this service.

- Developer Guide: Describes key concepts of AWS Step Functions and provides instructions for using the features of AWS Step Functions.
- API Reference: Documents the AWS Step Functions API.
- Amazon States Language Specification: Describes the language that is used to define state machines for AWS Step Functions.
- Statelint: A tool to validate your Amazon States Language code.

178. AWS Storage Gateway

178.1. Service Overview

AWS Storage Gateway is a hybrid cloud storage service that gives you on-premises access to virtually unlimited cloud storage. Customers use Storage Gateway to simplify storage management and reduce costs for key hybrid cloud storage use cases. These include moving backups to the cloud, using on-premises file shares backed by cloud storage, and providing low latency access to data in AWS for on-premises applications.

AWS Storage Gateway can help you:

- Deliver low-latency data access to on-premises applications while leveraging the agility, economics and security capabilities of AWS in the cloud.
- Provide on-premises applications access to cloud-backed storage without disruption to your business by maintaining user and application workflows.
- Offer virtually unlimited cloud storage to users and applications without deploying new storage hardware.
- Support your compliance efforts with key capabilities like encryption, audit logging, and write-once, read-many (WORM) storage.

To support these use cases, the service provides four different types of gateways – Tape Gateway, Amazon S3 File Gateway, Amazon FSx File Gateway, and Volume Gateway – that seamlessly connect on-premises applications to cloud storage, caching data locally for low-latency access.

178.1.1. Features

- **Amazon S3 File Gateway**: Amazon S3 File Gateway presents a file interface that enables you to store files as objects in Amazon S3 using the industry-standard NFS and SMB file protocols, and access those files via NFS and SMB from your data center or Amazon EC2, or access those files as objects directly in Amazon S3. POSIX-style metadata, including ownership, permissions, and timestamps are durably stored in Amazon S3 in the user-metadata of the object associated with the file. Once objects are transferred to S3, they can be managed as native S3 objects and bucket policies such as lifecycle management and Cross-Region Replication (CRR), and can be applied directly to objects stored in your bucket. Amazon S3 File Gateway also publishes audit

logs for SMB file share user operations to Amazon CloudWatch. Customers can use Amazon S3 File Gateway to back up on-premises file data as objects in Amazon S3 (including Microsoft SQL Server and Oracle databases and logs), and for hybrid cloud workflows using data generated by on-premises applications for processing by AWS services such as machine learning or big data analytics.

- **Amazon FSx File Gateway**: Amazon FSx File Gateway provides fast, low-latency on-premises access to fully managed, highly reliable, and scalable file shares in the cloud using the industry-standard SMB protocol. Customers can store and access file data in Amazon FSx with Windows-native compatibility including full NTFS support, shadow copies, and Access Control Lists (ACLs). Use Amazon FSx File Gateway for your on-premises file-based business applications and workloads such as user or group file shares, web content management, and media workflows. With Amazon FSx File Gateway, customers can easily migrate and consolidate their on-premises file-based application data stored on Network-Attached-Storage (NAS) arrays or file server VMs into FSx for Windows File Server for scalable shared file access that seamlessly integrates with your existing environment. With the HDD file storage option, Amazon FSx for Windows File Server offers the lowest-cost file storage in the cloud for Windows applications and workloads, or SSD storage for performance-intensive workloads. Customers that use Amazon FSx File Gateway can also benefit from other integrated AWS services for simplified storage management and data protection. You can automatically send logs of SMB user operations to Amazon CloudWatch to perform auditing and analysis, and use AWS Backup for centralized backup and retention.
- **Tape Gateway**; Tape Gateway presents an iSCSI-based virtual tape library (VTL) of virtual tape drives and a virtual media changer to your on-premises backup application. It is compatible with most leading backup applications, so you can continue using your tape-based backup workflows. Tape Gateway stores your virtual tapes in Amazon Simple Storage Service (Amazon S3) and creates new ones automatically, simplifying management and your transition to AWS. Its VTL interface helps you reduce physical tape infrastructure capital expenses, multi-year maintenance contract commitments, and ongoing media costs. You pay only for the capacity you use and scale as your needs grow. For any petabyte-scale tape data migration needs you can use a Snowball Edge Storage Optimized device with Tape Gateway to move your physical tape data to either S3 Glacier Flexible Retrieval or S3 Glacier Deep Archive, further reducing your long-term storage costs. You can access your data stored as virtual tapes in AWS through a Tape Gateway running in AWS or in your data center over the network. With Tape Gateway, you'll no longer need to store media at offsite facilities and do tape media migration from one generation to the next manually.
- **Volume Gateway**: Volume Gateway presents your applications block storage volumes using the iSCSI protocol. Data written to these volumes can be asynchronously backed up as point-in-time snapshots of your volumes, and stored in the cloud as Amazon EBS snapshots. You can back up your on-premises Volume Gateway volumes using the service's native snapshot scheduler or by using the AWS Backup service. In both cases, volume backups are stored as Amazon EBS snapshots in AWS. These snapshots are incremental backups that capture only changed blocks. All snapshot storage is also compressed to minimize your storage charges. Customers often choose Volume Gateway to backup local applications, and use it for disaster recovery based on EBS Snapshots, or Cached Volume Clones. Volume Gateway integration with AWS Backup enables customers to use the AWS Backup service to protect on-premises applications that use Storage Gateway volumes. AWS Backup supports backup and restore of both

cached and stored volumes. Using AWS Backup with Volume Gateway together helps you centralize backup management, reduce your operational burden, and meet compliance requirements.

178.1.2. Benefits

- **Standard Storage Protocols:** Storage Gateway seamlessly connects to your local production or backup applications with NFS, SMB, iSCSI, or iSCSI-VTL, so you can adopt AWS Cloud storage without needing to modify your applications. Its protocol conversion and device emulation enables you to access block data on volumes managed by Storage Gateway on top of Amazon S3, store files as native Amazon S3 objects or in fully managed cloud file shares with Amazon FSx for Windows File Server, and keep virtual tape backups online in a virtual tape library backed by S3 or move the backups to a tape archive tier on Amazon S3 Glacier and Amazon S3 Glacier Deep Archive.
- **Fully Managed Cache:** The local gateway appliance maintains a cache of recently written or read data so your applications can have low-latency access to data that is stored durably in AWS. The gateways use a read-through and write-back cache, committing data locally, acknowledging the write operations, and then asynchronously copying data to AWS, reducing application latency.
- **Optimized and Secure Data Transfer:** Storage Gateway provides secure upload of changed data and secure downloads of requested data, encrypting data in transit between any type of gateway appliance and AWS using SSL. Storage Gateway delivers end-to-end protection of customer data from the Storage Gateway in the enterprise network to the data residing in AWS. The service supports security features, access controls, and supplies compliances and certifications that address enterprise customers' real and perceived security concerns when using AWS Cloud storage via the Storage Gateway Optimizations such as multi-part management, automatic buffering, delta transfers used across all gateway types, and data compression applied for all block and virtual tape data. Storage Gateway offers Federal Information Processing Standard 140-2 (FIPS) compliant endpoints in AWS GovCloud (US-East) and AWS GovCloud (US-West).
- **Easily consume AWS Services:** Storage Gateway enables customers to easily consume AWS services. As a native AWS service, Storage Gateway integrates with other AWS services for storage, backup, and management while still integrating with on-premises environments. The service stores files as native Amazon S3 objects or fully managed file shares in Amazon FSx for Windows File Server, archives virtual tapes in Amazon S3 Glacier and Amazon S3 Glacier Deep Archive, and stores EBS snapshots generated by the Volume Gateway with Amazon EBS. Storage Gateway also integrates with AWS Backup to manage backup and recovery of Volume Gateway volumes, simplifying your backup management, and helping you meet your business and regulatory backup compliance requirements. Storage Gateway publishes health and performance logs and metrics to Amazon CloudWatch and provides monitoring of metrics and alarms in the Storage Gateway console. Storage Gateway integrates with AWS IAM to help manage and secure access to Storage Gateway resources. Your data is encrypted by default at rest using S3-SSE or you can choose to use your own encryption keys through Storage Gateway's integration with AWS KMS.
- **High Availability on VMware:** Storage Gateway provides high availability on VMware through a set of health-checks integrated with VMware vSphere High Availability (VMware HA). With this integration, Storage Gateway deployed in a VMware

environment on-premises, or in VMware Cloud on AWS, will automatically recover from most service interruptions in under 60 seconds. This protects storage workloads against hardware, hypervisor, or network failures, storage errors, or software errors, such as connection timeouts and file share or volume unavailability.

178.2. Backup/Restore and Disaster Recovery

This service allows for backup and recovery. The service can back up block volumes. Users control this via a manual or scheduled API call. Users schedule and recover backups through a web interface.

178.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

178.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/storagegateway/index.html
- **Service quotas:** https://docs.aws.amazon.com/storagegateway/latest/userguide/resource-gateway-limits.html
- **Service FAQs:** https://aws.amazon.com/storagegateway/faqs/?nc=sn&loc=6

178.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/storagegateway/index.html and the following links for comprehensive technical documentation regarding this service.

- User Guide for Amazon S3 File Gateway: Describes Amazon S3 File Gateway concepts and provides instructions on using the various features with both the console and the API.
- User Guide for Amazon FSx File Gateway: Describes Amazon FSx File Gateway, which provides access to in-cloud Amazon FSx for Windows File Server shares from on-premises facilities. Includes instructions on working with the console and the API.
- User Guide for Tape Gateway: Describes Tape Gateway, a durable, cost-effective tape-based solution for archiving data in the AWS Cloud. Provides concepts and instructions on using the various features with both the console and the API.
- User Guide for Volume Gateway: Describes Volume Gateway concepts, including details about cached and stored volume architectures, and provides instructions on using their features with both the console and the API.
- API Reference: Describes the AWS Storage Gateway API and CLI operations. Also provides sample requests, responses, and errors for the supported web services protocols. .

179. AWS Systems Manager

179.1. Service Overview

AWS Systems Manager centralizes operational data from multiple AWS services and automates tasks across your AWS resources. You can create logical groups of resources such as

applications, different layers of an application stack, or production versus development environments. With Systems Manager, you can select a resource group and view its recent API activity, resource configuration changes, related notifications, operational alerts, software inventory, and patch compliance status. It lets you take action on each resource group depending on your operational needs. Systems Manager provides a central place to view and manage your AWS resources, so you can have complete visibility and control over your operations.

179.1.1. Features

- **Explorer:** AWS Systems Manager Explorer is a customizable dashboard, providing key insights and analysis into the operational health and performance of your AWS environment. Explorer aggregates operational data from across AWS accounts and AWS Regions to help you prioritize and identify where action may be required.
- **OpsCenter:** OpsCenter provides a central location where operations engineers and IT professionals can view, investigate, and resolve operational issues related to AWS resources. OpsCenter aggregates and standardizes operational issues, referred to as OpsItems, while providing contextually relevant data that helps with diagnosis and remediation.
- **Incident Manager**: AWS Systems Manager Incident Manager enables faster resolution of critical application availability and performance issues. It helps you prepare for incidents with automated response plans that bring the right people and information together. With Incident Manager, you can automatically take action when a critical issue is detected by an Amazon CloudWatch alarm or Amazon EventBridge event. Incident Manager executes pre-configured response plans to engage responders via SMS and phone calls, links designated chat channels using AWS Chatbot, and executes AWS Systems Manager Automation runbooks. Incident Manager helps you improve service reliability by suggesting post-incident action items, such as automating a runbook step or adding a new alarm, based on Amazon's post-incident analysis template. To learn more, visit the Incident Manager feature page and to get started, visit the Systems Manager console.
- **Application Manager**: AWS Systems Manager Application Manager helps you investigate and remediate issues with your AWS resources in the context of your applications. With Application Manager, you can discover and/or define your application components, view operations data (e.g. deployment status, Amazon CloudWatch alarms, resource configurations, and operational issues) in the context of an application, and perform remedial actions such as patching and running Automation runbooks. This streamlines operational workflows for your applications, avoiding the need to use different consoles to investigate and remediate operational issues. Application Manager will display data and alarms and take action on your existing container clusters in Amazon Elastic Container Service (ECS) and Amazon Elastic Kubernetes Service (EKS) environments. Additionally, you can also manage the full lifecycle of your AWS CloudFormation templates and stacks from within the Application Manager console.
- **AppConfig**: AWS AppConfig helps you deploy application configuration in a managed and a monitored way just like code deployments, but without the need to deploy the code if a configuration value changes. AWS AppConfig scales with your infrastructure so you can deploy configurations to any number of Amazon Elastic Compute Cloud (EC2) instances, containers, AWS Lambda functions, mobile apps, IoT devices, or on-premises instances. AWS AppConfig enables you to update configurations by entering changes through the API or AWS Management Console. AWS AppConfig allows you to validate

those changes semantically and syntactically to ensure configurations are aligned to their respective applications' expectation, thus helping you to prevent potential outages. You can deploy your application configurations with similar best practices as code deployments, including staging roll-outs, monitoring alarms, and rolling back changes should an error occur. To learn more, visit the AWS AppConfig feature page.

- **Parameter Store**: AWS Systems Manager provides a centralized store to manage your configuration data, whether plain-text data such as database strings or secrets such as passwords. This allows you to separate your secrets and configuration data from your code. Parameters can be tagged and organized into hierarchies, helping you manage parameters more easily. For example, you can use the same parameter name, "db-string", with a different hierarchical path, "dev/db-string" or "prod/db-string", to store different values. Systems Manager is integrated with AWS Key Management Service (KMS), allowing you to automatically encrypt the data you store. You can also control user and resource access to parameters using AWS IAM. Parameters can be referenced through other AWS services, such as Amazon ECS, AWS Lambda, and AWS CloudFormation.
- **Change Manager**: AWS Systems Manager Change Manager simplifies the way you request, approve, implement, and report on operational changes to your application configuration and infrastructure on AWS and on premises. With Change Manager, you use pre-approved change workflows to help avoid unintentional results when making operational changes. Change Manager helps you safely implement changes, while detecting schedule conflicts with important business events and automatically notifying impacted approvers. Using Change Manager's change reports, you can monitor progress and audit operational changes across your organization, providing improved visibility and accountability.
- **Automation**: AWS Systems Manager allows you to safely automate common and repetitive IT operations and management tasks. With Systems Manager Automation, you use predefined playbooks, or you can build, run, and share wiki-style automated playbooks to enable AWS resource management across multiple accounts and AWS Regions. You can run Python or PowerShell scripts as part of a playbook in combination with other automation actions such as approvals, AWS API calls, or running commands on your EC2 instances. These playbooks can be scheduled in a maintenance window, triggered based on changes to AWS resources through Amazon CloudWatch Events, or executed directly through the AWS Management Console, CLIs, and SDKs. Automation can track the execution of each step in a playbook, require approvals, incrementally roll out changes, and automatically halt the roll out if errors occur.
- **Maintenance Windows**: AWS Systems Manager lets you schedule windows of time to run administrative and maintenance tasks across your instances. This ensures that you select a convenient and safe time to install patches and updates or make other configuration changes, improving the availability and reliability of your services and applications.
- **Fleet Manager**: AWS Systems Manager Fleet Manager streamlines your remote management process for servers and edge devices. With Fleet Manager, you save time and money by managing and troubleshooting your fleet running in the cloud or on premises, without the need to remotely connect to them. You can drill down to individual nodes (services, devices, or other resources) to perform common system management tasks such as disk and file exploration, log management, Windows Registry operations, and user management, from a console. In break glass scenarios, you can quickly gain

secure shell, CLI, and console-based Remote Desktop Protocol (RDP) access to your instances, from a console, to respond to issues faster.

179.1.2. Benefits

- **Compliance**: AWS Systems Manager automatically aggregates and displays operational data for each resource group through a dashboard. Systems Manager eliminates the need for you to navigate across multiple AWS consoles to view your operational data. With Systems Manager you can view API call logs from AWS CloudTrail, resource configuration changes from AWS Config, software inventory, and patch compliance status by resource group. It also integrates with your AWS CloudWatch Dashboards, AWS Trusted Advisor notifications, and AWS Personal Health Dashboard performance and availability alerts into your Systems Manager dashboard. Systems Manager centralizes all relevant operational data, providing a clear view of your infrastructure compliance and performance.
- **Inventory:** AWS Systems Manager collects information about your instances and the software installed on them, helping you to understand your system configurations and installed applications. You can collect data about applications, files, network configurations, Windows services, registries, server roles, updates, and any other system properties. The gathered data enables you to manage application assets, track licenses, monitor file integrity, discover applications not installed by a traditional installer, and more.
- **Session Manager**: AWS Systems Manager provides a browser-based interactive shell, CLI and browser based remote desktop access for managing instances on your cloud, or on-premises and edge devices, without the need to open inbound ports, manage Secure Shell (SSH) keys, or use bastion hosts. Administrators can grant and revoke access to instances through a central location by using AWS Identity and Access Management (IAM) policies. This allows you to control which users can access each instance, including the option to provide non-root access to specified users. Once access is provided, you can audit which user accessed an instance and log each command to Amazon S3 or Amazon CloudWatch Logs using AWS CloudTrail.
- **Run Command**: AWS Systems Manager provides you safe, secure remote management of your instances at scale without logging into your servers, replacing the need for bastion hosts, SSH, or remote PowerShell. It provides a simple way of automating common administrative tasks across groups of instances such as registry edits, user management, and software and patch installations. Through integration with AWS Identity and Access Management (IAM), you can apply granular permissions to control the actions users can perform on instances. All actions taken with Systems Manager are recorded by AWS CloudTrail, allowing you to audit changes throughout your environment.
- **State Manager**: AWS Systems Manager provides configuration management, which helps you maintain consistent configuration of your Amazon EC2 or on-premises instances. With Systems Manager, you can control configuration details such as server configurations, anti-virus definitions, firewall settings, and more. You can define configuration policies for your servers through the AWS Management Console or use existing scripts, PowerShell modules, or Ansible playbooks directly from GitHub or Amazon S3 buckets. Systems Manager automatically applies your configurations across your instances at a time and frequency that you define. You can query Systems Manager at any time to view the status of your instance configurations, giving you on-demand visibility into your compliance status.

- **Patch Manager**: AWS Systems Manager helps you select and deploy operating system and software patches automatically across large groups of cloud or on-premises instances and edge devices. Through patch baselines, you can set rules to auto-approve select categories of patches to be installed, such as operating system or high severity patches, and specify a list of patches that override these rules and are automatically approved or rejected. You can also schedule maintenance windows for your patches so that they are only applied during preset times. Systems Manager helps ensure that your software is up-to-date and meets your compliance policies.
- **Distributor**: AWS Systems Manager helps you securely distribute and install software packages, such as software agents. Systems Manager Distributor allows you to centrally store and systematically distribute software packages while you maintain control over versioning. You can use Distributor to create and distribute software packages and then install them using Systems Manager Run Command and State Manager. Distributor can also use IAM policies to control who can create or update packages in your account. You can use the existing IAM policy support for Systems Manager Run Command and State Manager to define who can install packages on your hosts.
- **Connect with ITSM / ITOM Software**: IT Service Management (ITSM) tools, such as Jira Service Desk, can connect with AWS Systems Manager to make it easier for ITSM platform users to manage AWS resources. These AWS Service Management Connectors provide Jira Service Desk administrators governance and oversight over AWS products.

179.2. Backup/Restore and Disaster Recovery

This service allows for backup and recovery. The service can back up logs to S3. Users control this via a manual or scheduled API call. Users schedule and recover backups through a web interface.

179.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

179.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/systems-manager/
- **Service quotas:** https://docs.aws.amazon.com/general/latest/gr/ssm.html
- **Service FAQs:** https://aws.amazon.com/systems-manager/faq/

179.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/systems-manager/ and the following links for comprehensive technical documentation regarding this service.

- User Guide: Provides detailed descriptions of AWS Systems Manager concepts and includes instructions for using both the console and the command-line interface. Learn how to configure, manage, and automate tasks on groups of AWS instances and resources.
- Automation runbook reference: Describes the predefined runbooks for Automation, a capability of AWS Systems Manager, in detail.

- API Reference: Describes the API operations for AWS Systems Manager in detail. In addition, the API Reference provides sample requests, responses, and errors for the supported web service protocols.
- AWS Systems Manager in the AWS CLI Reference: Describes the AWS CLI commands that you can use to automate systems management tasks.

180. AWS Transfer Family

180.1. Service Overview

The AWS Transfer Family provides fully managed support for file transfers over SFTP, FTPS, and FTP directly into and out of Amazon S3 and Amazon EFS. AWS helps you seamlessly migrate your file transfer workflows with AWS Transfer for SFTP, AWS Transfer for FTPS, and AWS Transfer for FTP by integrating with existing authentication systems — so nothing changes for your customers, partners, and internal teams, or their applications.

180.1.1. Features

- **Fully managed highly available infrastructure:** AWS transparently operates and manages all of the compute, storage, and other infrastructure necessary to maintain high availability and performance for your endpoint. Your endpoint is designed to be available 24 hours a day, 7 days a week, 365 days a year. You get full redundancy across multiple Availability Zones within an AWS Region.
- **Elastic resources:** The AWS Transfer Family can meet the needs of your dynamic workloads with elastic compute infrastructure. Built-in autoscaling means that you never have to worry about provisioning additional resources if your data loads grow over time. You don't have to worry if workloads spike during certain hours of the day or days of the month.
- **Supports multiple user authentication methods:** The AWS Transfer Family supports common user authentication systems, including Microsoft Active Directory and Lightweight Directory Access Protocol (LDAP). Alternatively, you can also choose to store and manage users' credentials directly within the service. By connecting your existing identity provider to the AWS Transfer Family service, you assure that your external users continue to have the correct, secure level of access to your data resources without disruption.
- **Data stored natively in AWS Storage services:** Store the files you exchange as objects in your Amazon S3 bucket or Amazon EFS file system so you can extract business insights faster. The key piece that makes this exchange possible is that AWS Transfer Family stores your data natively in S3 or EFS, while preserving relevant file metadata. For example, with your files stored in Amazon S3, you can use Amazon Translate to make process documents more tailored for international audiences. You can also use Amazon Comprehend to extract relationships and insights from text files, or you could even use Amazon Athena to query CSV files to analyze historical data. Similarly with files in EFS, you can directly integrate your ERP system to access these files on arrival from your business partners.
- **Simple user experience:** An intuitive user interface and API makes it simple for you to configure your SFTP, FTPS, or FTP endpoint and set up client access. For external users, the service supports commonly used SFTP clients such as WinSCP, FileZilla, and scripts. Users don't have to change their behavior to continue sharing data with you in the cloud.

- **Familiar and comprehensive AWS management services:** With the AWS Transfer Family, you can use AWS Identity and Access Management (IAM) for security and identity management and Amazon CloudWatch for monitoring and event triggers to start post-upload processing. You can use AWS Key Management Service (AWS KMS), Amazon S3 server-side encryption, or Customer Managed Keys with Amazon EFS to control encryption at rest. Additionally, AWS CloudTrail helps you meet compliance requirements with granular auditing of user and API activity.

180.1.2. Benefits

- **Use existing authentication systems to modernize workflows:** Move files seamlessly and modernize your transfer workflows within hours by using your existing authentication systems.
- **Security and compliance made easy:** Meet your security requirements with data encryption, VPC and FIPS endpoints, compliance certifications, and more.
- **Support concurrent users:** Support thousands of concurrent users and quickly scale your business-to-business (B2B) file transfers.
- **Real-time insights:** Store information in Amazon S3 or Amazon EFS to break down data silos, control file access, and gain real-time business insights.

180.2. Backup/Restore and Disaster Recovery

This requirement is not applicable for this service. For additional information beyond what is described herein, please refer to http://aws.amazon.com/documentation/.

180.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

180.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/transfer/
- **Service quotas:** https://docs.aws.amazon.com/general/latest/gr/transfer-service.html
- **Service FAQs:** https://aws.amazon.com/aws-transfer-family/faqs/

180.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/transfer/ and the following links for comprehensive technical documentation regarding this service.

- AWS Transfer Family User Guide: Describes all AWS Transfer Family concepts and provides instructions on using the various features with the console, the command line interface, and the AWS Transfer Family API.

181. AWS Transit Gateway

181.1. Service Overview

AWS Transit Gateway connects VPCs and on-premises networks through a central hub. This simplifies your network and puts an end to complex peering relationships. It acts as a cloud router – each new connection is only made once.

As you expand globally, inter-Region peering connects AWS Transit Gateways together using the AWS global network. Your data is automatically encrypted, and never travels over the public internet. And, because of its central position, AWS Transit Gateway Network Manager has a unique view over your entire network, even connecting to Software-Defined Wide Area Network (SD-WAN) devices.

181.1.1. Features

- **Routing:** AWS Transit Gateways supports dynamic and static layer 3 routing between Amazon Virtual Private Clouds (VPCs) and VPN. Routes determine the next hop depending on the destination IP address of the packet, and can point to an Amazon VPC or to a VPN connection.
- **Edge connectivity:** You can create VPN connections between your AWS Transit Gateway and on-premises gateways using VPN. You can create multiple VPN connections that announce the same prefixes and enable Equal Cost Multipath (ECMP) between these connections. By load-balancing traffic over multiple paths, ECMP can increase the bandwidth.
- **Transit Gateway Connect:** AWS Transit Gateway Connect enables native integration of Software-Defined Wide Area Network (SD-WAN) appliances into AWS. Customers can now seamlessly extend their SD-WAN edge into AWS using standard protocols such as Generic Routing Encapsulation (GRE) and Border Gateway Protocol (BGP). It provides customers with added benefits such as improved bandwidth and supports dynamic routing with increased route limits, thus removing the need to set up multiple IPsec VPNs between the SD-WAN appliances and Transit Gateway.
- **Amazon VPC feature interoperability:** AWS Transit Gateway enables the resolution of public DNS hostnames to private IP addresses when queried from Amazon VPCs that are also attached to the AWS Transit Gateway.An instance in an Amazon VPC can access a NAT gateway, Network Load Balancer, AWS PrivateLink, and Amazon Elastic File System in others Amazon VPCs that are also attached to the AWS Transit Gateway.
- **Monitoring:** AWS Transit Gateway provides statistics and logs that are then used by services such as Amazon CloudWatch and Amazon VPC Flow Logs. You can use Amazon CloudWatch to get bandwidth usage between Amazon VPCs and a VPN connection, packet flow count, and packet drop count. You can also enable Amazon VPC Flow Logs on AWS Transit Gateway so you can capture information on the IP traffic routed through the AWS Transit Gateway.AWS Transit Gateway Network Manager includes events and metrics to monitor the quality of your global network, both in AWS and on premises. Event alerts specify changes in the topology, routing, and connection status. Usage metrics provide information on up/down connection, bytes in/out, packets in/out, and packets dropped.
- **Management:** You can use the command-line interface (CLI), AWS Management Console, or AWS CloudFormation to create and manage your AWS Transit Gateway. AWS Transit Gateway provides Amazon CloudWatch metrics, such as the number of bytes sent and received between Amazon VPCs and VPNs, the packet count, and the drop count. In addition, you can use Amazon VPC Flow Logs with AWS Transit Gateway to capture information about the IP traffic going through the AWS Transit Gateway attachment.
- **Peering:** With transit gateway peering, you can establish peering connections between transit gateways in the same AWS region or across regions. Peering allows customers to directly route traffic between two transit gateways. Inter-region peering provides you

with a simple and cost-effective way to share resources between AWS Regions or replicate data for geographic redundancy. Intra-region peering allows multiple teams within your organization to deploy their own transit gateways and easily interconnect their networks in the same AWS region.

- **Multicast:** With Transit Gateway multicast, you can now easily create and manage multicast groups in the cloud, much easier than deploying and managing legacy hardware on premises. You can scale up and down your multicast solution in the cloud to simultaneously distribute a stream of content to multiple subscribers. With Transit Gateway multicast you have fine-grain control on who can produce and who can consume multicast traffic.
- **Security:** AWS Transit Gateway is integrated with Identity and Access Management (IAM), enabling you to manage access to AWS Transit Gateway securely. Using IAM, you can create and manage AWS users and groups, and use permissions to allow and deny their access to the AWS Transit Gateway.
- **Automated provisioning:** Once you've registered existing AWS Transit Gateways, the Network Manager automatically identifies the Site-to-Site VPN connections and the on-premises resources with which they are associated. The SD-WAN consoles from vendors that have integrated AWS Transit Gateway, such as Cisco, Aruba, Silver Peak, or Aviatrix, automatically provision new AWS Site-to-Site VPN connections in Transit Gateway Network Manager and automate the definition of your on-premises network in Transit Gateway Network Manager. You can also manually define your on-premises network in Transit Gateway Network Manager.

181.1.2. Benefits

- **Easier connectivity:** AWS Transit Gateway acts as a cloud router to simplify your network architecture. As your network grows, the complexity of managing incremental connections doesn't slow you down. When building global applications, you can connect AWS Transit Gateways using inter-Region peering.
- **Better visibility and control:** With AWS Transit Gateway Network Manager, you can easily monitor your Amazon VPCs and edge connections from a central console. Integrated with popular SD-WAN devices, AWS Transit Gateway Network Manager helps you quickly identify issues and react to events on your global network.
- **Improved security:** Traffic between an Amazon VPC and AWS Transit Gateway remains on the AWS global private network and is not exposed to the public internet. AWS Transit Gateway inter-Region peering encrypts all traffic, with no single point of failure or bandwidth bottleneck. This helps protect against distributed denial of service (DDoS) attacks and other common exploits.
- **Flexible multicast:** AWS Transit Gateway multicast support distributes the same content to multiple specific destinations. This eliminates the need for expensive on-premises multicast networks and reduces the bandwidth needed for high-throughput applications such as video conferencing, media, or teleconferencing.

181.2. Backup/Restore and Disaster Recovery

This requirement is not applicable for this service. For additional information beyond what is described herein, please refer to http://aws.amazon.com/documentation/.

181.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

181.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/vpc/index.html
- **Service quotas:** https://docs.aws.amazon.com/vpc/latest/tgw/transit-gateway-quotas.html
- **Service FAQs:** https://aws.amazon.com/transit-gateway/faqs/

181.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/vpc/index.html and the following links for comprehensive technical documentation regarding this service.

- Transit Gateway Guide: Simplify your network architecture.
- Transit Gateway API Reference: Get syntax and examples for the API actions.
- Network Manager User Guide: Monitor the quality of your global network, both in AWS and on premises.
- Network Manager API Reference: Get syntax and examples for the API actions.

182. AWS Trusted Advisor

182.1. Service Overview

AWS Trusted Advisors provides recommendations that help you follow AWS best practices. Trusted Advisor evaluates your account by using checks. These checks identify ways to optimize your AWS infrastructure, improve security and performance, reduce costs, and monitor service quotas. You can then follow the check recommendations to optimize your services and resources.

AWS Basic Support and AWS Developer Support customers can access core security checks and all checks for service quotas. AWS Business Support and AWS Enterprise Support customers can access all checks, including cost optimization, security, fault tolerance, performance, and service quotas. For a complete list of checks and descriptions, see the Trusted Advisor Best Practices.

182.1.1. Features

- **Cost optimization:** Trusted Advisor can help you save cost by recommending you to delete unused or idle resources, or use reserved capacity.
- **Performance:** Trusted Advisor can improve the performance of your services by ensuring you to take advantage of provisioned throughput, and monitoring for overutilized Amazon EC2 instances.

182.1.2. Benefits

- **Security:** Trusted Advisor can improve the security of your application by recommending you to enable AWS security features, and review your permissions.

- **Fault tolerance:** Trusted Advisor can increase the availability of your AWS application by recommending you to take advantage of auto scaling, health checks, multi-AZ Regions, and backup capabilities.
- **Service quotas:** Service quotas, also referred to as Service limits, are the maximum number of service resources or operations that apply to an account or a Region. Trusted Advisor can notify you if you use more than 80% of a service quota. You can then follow recommendations to delete resources or request a quota increase. Check results are based on a snapshot, so your current usage might vary.

182.2. Backup/Restore and Disaster Recovery

This requirement is not applicable for this service. For additional information beyond what is described herein, please refer to http://aws.amazon.com/documentation/.

182.3. Pricing Overview

AWS Basic Support and AWS Developer Support customers can access core security checks and all checks for service quotas. AWS Business Support and AWS Enterprise Support customers can access all checks, including cost optimization, security, fault tolerance, performance, and service quotas.

182.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/awssupport/latest/user/trusted-advisor.html
- **Service FAQs:** https://aws.amazon.com/premiumsupport/faqs/

182.5. Technical Requirements

Please refer to the following links for comprehensive technical documentation regarding this service.

- https://docs.aws.amazon.com/awssupport/latest/user/trusted-advisor.html

183. AWS VPN

183.1. Service Overview

AWS Virtual Private Network solutions establish secure connections between your on-premises networks, remote offices, client devices, and the AWS global network. AWS VPN is comprised of two services: AWS Site-to-Site VPN and AWS Client VPN. Each service provides a highly-available, managed, and elastic cloud VPN solution to protect your network traffic.

AWS Site-to-Site VPN creates encrypted tunnels between your network and your Amazon Virtual Private Clouds or AWS Transit Gateways. For managing remote access, AWS Client VPN connects your users to AWS or on-premises resources using a VPN software client.

183.1.1. Features

- **Accelerated Site-to-Site VPN:** When you connect an on-premises location to the AWS cloud, Accelerated Site-to-Site VPN will route your VPN traffic to the closest AWS edge location. Accelerated VPN improves the performance of your Site-to-Site VPN connections by reducing the distance over which data is being shared on the internet and leveraging instead the reliability and performance of the AWS global fiber network.

Accelerated Site-to-Site VPN is ideal to connect business-critical locations with your global network, both on premises and in AWS. VPN acceleration will incur additional charges from utilizing both AWS Site-to-Site VPN and AWS Global Accelerator.

- **High availability:** With AWS Site-to-Site VPN you can create failover and CloudHub solutions with AWS Direct Connect. CloudHub enables your remote sites to communicate with each other, and not just with the VPC. It operates on a simple hub-and-spoke model that you can use with or without a VPC. This design is suitable for customers with multiple branch offices and existing internet connections who would like to implement a convenient, potentially low-cost hub-and-spoke model for primary or backup connectivity between these remote offices.
- **Customization:** AWS Site-to-Site VPN offers customizable tunnel options including inside tunnel IP address, pre-shared key, and Border Gateway Protocol Autonomous System Number (BGP ASN). In this way, you can set up multiple secure VPN tunnels to increase the bandwidth for your applications or for resiliency in case of a down time. In addition, equal-cost multi-path routing (ECMP) is available with AWS Site-to-Site VPN on AWS Transit Gateway to help increase the traffic bandwidth over multiple paths.
- **Network Address Translation (NAT) Traversal:** AWS Site-to-Site VPN supports NAT Traversal applications so that you can use private IP addresses on private networks behind routers with a single public IP address facing the internet.
- **Authentication:** AWS Client VPN will authenticate using either Active Directory or certificates. Client VPN integrates with AWS Directory Services, which connects to your existing on-premises Active Directory, so it does not require you to replicate data from your existing Active Directory to the cloud. Certificate-based authentication with Client VPN integrates with AWS Certificate Manager to easily provision, manage, and deploy certificates.

183.1.2. Benefits

- **Secure connectivity:** AWS Client VPN uses OpenVPN, which utilizes a TLS encrypted control channel to negotiate the data channel parameters. The data channel is SSL based, but adds additional safeguards (such as HMAC, hashing, and x.509 certificates).
- **Monitoring:** AWS Site-to-Site VPN can send metrics to CloudWatch to provide you with greater visibility and monitoring. CloudWatch also allows you to send your own custom metrics and add data points in any order, and at any rate you choose. You can retrieve statistics about those data points as an ordered set of time-series data.
- **Authorization:** AWS Client VPN provides network-based authorization so you can define access control rules that limit access to specific networks, based on Active Directory groups.
- **Secure connectivity:** AWS Client VPN uses the secure TLS VPN tunnel protocol to encrypt the traffic. A single VPN tunnel terminates at each Client VPN endpoint and provides users access to all AWS and on-premises resources.
- **Connection management:** You can use Amazon CloudWatch Logs to monitor, store, and access your log files from AWS Client VPN connection logs. You can then retrieve the associated log data from CloudWatch Logs. You can easily monitor, conduct forensics analysis, and terminate specific connections, while staying in control of who has access to your network.
- **Compatibility with your employees' devices:** AWS Client VPN is designed to connect devices to your network. It allows you to choose from OpenVPN-based client, giving

employees the option to use the device of their choice, including Windows, Mac, iOS, Android, and Linux-based devices.

183.2. Backup/Restore and Disaster Recovery

This requirement is not applicable for this service. For additional information beyond what is described herein, please refer to http://aws.amazon.com/documentation/.

183.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

183.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/vpn/index.html
- **Service quotas:** https://docs.aws.amazon.com/vpn/latest/s2svpn/vpn-limits.html
- **Service FAQs:** https://aws.amazon.com/vpn/faqs/

183.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/vpn/index.html and the following links for comprehensive technical documentation regarding this service.

- AWS Site-to-Site VPN User Guide: Describes key concepts and provides instructions for using the features of Site-to-Site VPN.
- AWS Site-to-Site VPN Customer Gateway Device Guide: Describes customer gateway devices and helps you configure them. Previously known as the Network Administrator Guide.
- API Reference: Describes the API operations for Site-to-Site VPN.
- EC2 section of the AWS CLI Reference: Describes the AWS CLI commands for administering Site-to-Site VPN.
- AWS Client VPN User Guide: Provides instructions for establishing a Client VPN session using a VPN client application.
- AWS Client VPN Administrator Guide: Describes key concepts and provides instructions for using the features of Client VPN.
- API Reference: Describes the API operations for Client VPN.
- EC2 section of the AWS CLI Reference: Describes the AWS CLI commands for administering Client VPN.

184. AWS WAF

184.1. Service Overview

AWS WAF is a web application firewall that helps protect your web applications or APIs against common web exploits and bots that may affect availability, compromise security, or consume excessive resources. AWS WAF gives you control over how traffic reaches your applications by enabling you to create security rules that control bot traffic and block common attack patterns, such as SQL injection or cross-site scripting. You can also customize rules that filter out specific traffic patterns. You can get started quickly using Managed Rules for AWS WAF, a pre-

configured set of rules managed by AWS or AWS Marketplace Sellers to address issues like the OWASP Top 10 security risks and automated bots that consume excess resources, skew metrics, or can cause downtime. These rules are regularly updated as new issues emerge. AWS WAF includes a full-featured API that you can use to automate the creation, deployment, and maintenance of security rules.

You can deploy AWS WAF on Amazon CloudFront as part of your CDN solution, the Application Load Balancer that fronts your web servers or origin servers running on EC2, Amazon API Gateway for your REST APIs, or AWS AppSync for your GraphQL APIs. With AWS WAF, you pay only for what you use and the pricing is based on how many rules you deploy and how many web requests your application receives.

184.1.1. Features

- **Web traffic filtering:** AWS WAF lets you create rules to filter web traffic based on conditions that include IP addresses, HTTP headers and body, or custom URIs. This gives you an additional layer of protection from web attacks that attempt to exploit vulnerabilities in custom or third party web applications. In addition, AWS WAF makes it easy to create rules that block common web exploits like SQL injection and cross site scripting. AWS WAF allows you to create a centralized set of rules that you can deploy across multiple websites. This means that in an environment with many websites and web applications you can create a single set of rules that you can reuse across applications rather than recreating that rule on every application you want to protect.
- **AWS WAF Bot Control:** AWS WAF Bot Control is a managed rule group that gives you visibility and control over common and pervasive bot traffic that can consume excess resources, skew metrics, cause downtime, or perform other undesired activities. With just a few clicks, you can block, or rate-limit, pervasive bots, such as scrapers, scanners, and crawlers, or you can allow common bots, such as status monitors and search engines. The Bot Control managed rule group can be used alongside other Managed Rules for WAF or your own custom WAF rules to protect your applications.
- **Full feature API:** AWS WAF can be completely administered via APIs. This provides organizations with the ability to create and maintain rules automatically and incorporate them into the development and design process. For example, a developer who has detailed knowledge of the web application could create a security rule as part of the deployment process. This capability to incorporate security into your development process avoids the need for complex handoffs between application and security teams to make sure rules are kept up to date. AWS WAF can also be deployed and provisioned automatically with AWS CloudFormation sample templates that allow you to describe all security rules you would like to deploy for your web applications delivered by Amazon CloudFront.
- **Real-time visibility:** AWS WAF provides real-time metrics and captures raw requests that include details about IP addresses, geo locations, URIs, User-Agent and Referrers. AWS WAF is fully integrated with Amazon CloudWatch, making it easy to setup custom alarms when thresholds are exceeded or particular attacks occur. This information provides valuable intelligence that can be used to create new rules to better protect applications.
- **Integration with AWS Firewall Manager:** You can centrally configure and manage AWS WAF deployments across multiple AWS accounts by using AWS Firewall Manager. As new resources are created, you can ensure that they comply with a common set of security rules. Firewall Manager automatically audits and informs your security team when there is a

policy violation, so they can respond immediately and take action. To learn more about Firewall Manager, visit the product website.

184.1.2. Benefits

- **Agile protection against web attacks:** AWS WAF rule propagation and updates take under a minute, enabling you to quickly update security across your environment when issues arise. WAF supports hundreds of rules that can inspect any part of the web request with minimal latency impact to incoming traffic. AWS WAF protects web applications from attacks by filtering traffic based on rules that you create. For example, you can filter any part of the web request, such as IP addresses, HTTP headers, HTTP body, or URI strings. This allows you to block common attack patterns, such as SQL injection or cross-site scripting.
- **Save time with managed rules:** With Managed Rules for AWS WAF, you can quickly get started and protect your web application or APIs against common threats. You can select from many rule types, such as ones that address issues like the Open Web Application Security Project (OWASP) Top 10 security risks, threats specific to Content Management Systems (CMS), or emerging Common Vulnerabilities and Exposures (CVE). Managed rules are automatically updated as new issues emerge, so that you can spend more time building applications.
- **Improved web traffic visibility:** AWS WAF gives near real-time visibility into your web traffic, which you can use to create new rules or alerts in Amazon CloudWatch. You have granular control over how the metrics are emitted, allowing you to monitor from the rule level to the entire inbound traffic. In addition, AWS WAF offers comprehensive logging by capturing each inspected web request's full header data for use in security automation, analytics, or auditing purposes.
- **Ease of deployment & maintenance:** AWS WAF is easy to deploy and protect applications deployed on either Amazon CloudFront as part of your CDN solution, the Application Load Balancer that fronts all your origin servers, Amazon API Gateway for your REST APIs, or AWS AppSync for your GraphQL APIs. There is no additional software to deploy, DNS configuration, SSL/TLS certificate to manage, or need for a reverse proxy setup. With AWS Firewall Manager integration, you can centrally define and manage your rules, and reuse them across all the web applications that you need to protect.
- **Easily monitor, block, or rate-limit bots:** With AWS WAF Bot Control, you get visibility and control over common and pervasive bot traffic to your applications. Within the AWS WAF console, you can monitor common bots, such as status monitors and search engines, and get detailed, real-time visibility into the category, identity, and other details of bot traffic. You can also block, or rate-limit, traffic from pervasive bots, such as scrapers, scanners, and crawlers. Using AWS Firewall Manager, you can deploy the Bot Control managed rule group across multiple accounts in your AWS Organization.
- **Security integrated with how you develop applications:** Every feature in AWS WAF can be configured using either the AWS WAF API or the AWS Management Console. This allows your DevOps team to define application-specific rules that increase web security as they develop applications. This lets you put web security at multiple points in the development process chain, from the hands of the developer initially writing code, to the DevOps engineer deploying software, to the security administrators enforcing a set of rules across the organization.

184.2. Backup/Restore and Disaster Recovery

This requirement is not applicable for this service. For additional information beyond what is described herein, please refer to http://aws.amazon.com/documentation/.

184.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

184.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/waf/index.html
- **Service quotas:** https://docs.aws.amazon.com/waf/latest/developerguide/limits.html
- **Service FAQs:** https://aws.amazon.com/waf/faqs/

184.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/waf/index.html and the following links for comprehensive technical documentation regarding this service.

- Developer Guide: Describes how to get started with AWS WAF. Explains key concepts and provides step-by-step instructions that show you how to use the features.
- API Reference: Describes all the API operations for AWS WAF in detail.

185. AWS Wavelength

185.1. Service Overview

AWS Wavelength is an infrastructure offering optimized for mobile edge computing applications. Wavelength Zones are AWS infrastructure deployments that embed AWS compute and storage services within communications service providers' (CSP) 5G networks, so application traffic from 5G devices reach application servers running in Wavelength Zones without leaving the telecommunications network. This avoids the latency that would result from application traffic traversing multiple hops across the internet to reach its destination, which allows customers to take full advantage of the latency and bandwidth benefits offered by modern 5G networks.

You can create Amazon Elastic Compute Cloud (EC2) instances, Amazon Elastic Block Store (EBS) volumes, and Amazon Virtual Private Cloud (VPC) subnets and carrier gateways in Wavelength Zones. You can also use services that orchestrate or work with EC2, EBS and VPC such as Amazon EC2 Auto Scaling, Amazon Elastic Kubernetes Service (EKS) clusters, Amazon Elastic Container Service (ECS) clusters, Amazon EC2 Systems Manager, Amazon CloudWatch, AWS CloudTrail, AWS CloudFormation, and AWS Application Load Balancer (ALB). AWS Wavelength services are part of a VPC connected over a reliable, high-bandwidth connection to an AWS Region for easy access to services including Amazon DynamoDB and Amazon Relational Database Service (RDS).

185.1.1. Features

- **Choice of compute instances:** Wavelength Zones support t3.medium, t3.xlarge and r5.2xlarge instances for applications requiring cost-effective general purpose compute. For applications such as game streaming and machine learning (ML) inference at the edge requiring GPU acceleration, Wavelength Zones support g4dn.2xlarge instances.

- **EBS storage:** Wavelength offers Elastic Block Store (EBS) gp2 volumes for persistent block storage. You can use EBS gp2 volumes for boot or data volumes, and attach or detach EBS volumes to EC2 instances. It provides snapshot and restore capabilities and lets you increase volume size without any performance impact. All EBS volumes and snapshots are fully encrypted by default. Any EBS snapshots will be stored using Amazon S3 in the Region associated with the Wavelength Zone.
- **Connectivity to 5G networks:** VPC: Amazon VPCs in an account can be extended to span multiple Availability Zones, including Wavelength Zones. Amazon EC2 instances and related services will appear as part of the user's regional VPC. Carrier Gateway: Wavelength also introduces a new component to the network setup – the Carrier Gateway. The Carrier Gateway enables connectivity from the user's subnet in the Wavelength Zone to the communications service provider's (CSP) network, the internet, or the AWS Region through the CSP's network. The Carrier Gateway serves two purposes. It allows inbound traffic from a CSP network in a specific location, and outbound traffic to the telecommunications network and the internet.
- **Management and monitoring:** You can use familiar AWS tools such as AWS CloudFormation, Amazon CloudWatch, AWS CloudTrail, and others to run and manage workloads in Wavelength Zones as you do for other cloud workloads today. You can use AWS Cost Explorer to monitor the cost of your projects.

185.1.2. Benefits

- **No learning curve:** Build next-generation applications without any learning curve using familiar AWS services, APIs, and tools.
- **Scale applications quickly:** Develop applications once and scale deployments to multiple Wavelength Zones across global 5G networks.
- **Accelerate app development:** Leverage proven AWS infrastructure and services to accelerate innovative 5G edge application development.

185.2. Backup/Restore and Disaster Recovery

This requirement is not applicable for this service. For additional information beyond what is described herein, please refer to http://aws.amazon.com/documentation/.

185.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

185.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/wavelength/index.html
- **Service quotas:** https://docs.aws.amazon.com/wavelength/latest/developerguide/wavelength-quotas.html
- **Service FAQs:** https://aws.amazon.com/wavelength/faqs/

185.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/wavelength/index.html and the following links for comprehensive technical documentation regarding this service.

- Developer Guide: Describes key concepts of AWS Wavelength and provides instructions about how to get started using Wavelength Zones.
- EC2 section of the AWS CLI Reference: Documents the AWS CLI commands for AWS Wavelength.
- Amazon EC2 API Reference: Describes the API operations for AWS Wavelength.

186. AWS Well-Architected Tool

186.1. Service Overview

The AWS Well-Architected Tool is designed to help you review the state of your applications and workloads, and it provides a central place for architectural best practices and guidance. The AWS Well-Architected Tool is based on the AWS Well-Architected Framework, which was developed to help cloud architects build secure, high-performing, resilient, and efficient application infrastructures. The Framework has been used in tens of thousands of workload reviews by AWS solutions architects, and it provides a consistent approach for evaluating your cloud architecture and implementing designs that will scale with your application needs over time.

In addition to the standard guidance provided by the AWS Well-Architected Framework and AWS-developed lenses, the AWS Well-Architected Tool allows you to add specific best practice guidance using custom lenses. By developing your own questions and evaluating your workloads using your organization's best practices, you can perform reviews based on technology or governance requirements specific to your industry.

To use the AWS Well-Architected Tool, just define your workload, apply one of the AWS Well-Architected lenses or your own custom lens, and begin your review. The tool then provides an action plan to help you build for the cloud using the defined best practices. The AWS Well-Architected Tool is available at no charge in the AWS Management Console.

186.1.1. Features

- **Review your workloads consistently:** Use a single tool and a consistent process to review and measure your cloud architectures. The AWS Well-Architected Tool allows you to monitor the status of multiple workloads across your organization and helps you understand potential risks. With the action plan, you can identify next steps for improvement, drive architectural decisions, and build for the cloud with confidence.
- **Customize your review:** You can create custom lenses and share them across your entire organization to measure workloads consistently. Specify rules to help you determine which options can result in a high or medium risk, and provide guidance on resolving those risks.

186.1.2. Benefits

- **Get free architectural guidance:** Benefit from access to knowledge and best practices used by AWS solutions architects whenever you need it. Answer questions about your application or workload, and the AWS Well-Architected Tool delivers an action plan with step-by-step guidance on areas for improvement.
- **Identify and implement improvements:** Support continuous improvement throughout the workload lifecycle. The AWS Well-Architected Tool makes it easy to save point-in-time milestones and track changes to your workload. Initiate new reviews as desired to help ensure that your architecture improves over time.

186.2. Backup/Restore and Disaster Recovery

This requirement is not applicable for this service. For additional information beyond what is described herein, please refer to http://aws.amazon.com/documentation/.

186.3. Pricing Overview

There is no additional charge for the AWS Well-Architected Tool. You pay only for your underlying AWS resources.

186.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/wellarchitected/
- **Service quotas:** https://docs.aws.amazon.com/general/latest/gr/wellarchitected.html
- **Service FAQs:** https://aws.amazon.com/well-architected-tool/faqs/

186.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/wellarchitected/ and the following links for comprehensive technical documentation regarding this service.

User Guide: Provides an overview of AWS Well-Architected Tool and includes instructions for measuring your workloads and tracking improvements to them over time.

API Reference: Describes all the API operations for AWS Well-Architected Tool in detail. Also provides sample requests, responses, and errors for the supported web services protocols.

187. Elastic Load Balancing

187.1. Service Overview

Elastic Load Balancing automatically distributes incoming application traffic across multiple targets, such as Amazon EC2 instances, containers, IP addresses, and AWS Lambda functions. It can handle the varying load of your application traffic in a single Availability Zone or across multiple Availability Zones. Elastic Load Balancing offers three types of load balancers that all feature the high availability, automatic scaling, and robust security necessary to make your applications fault tolerant:

187.1.1. Features

- **Gateway Load Balancer:** Gateway Load Balancer helps you easily deploy, scale, and manage your third-party virtual appliances. It gives you one gateway for distributing traffic across multiple virtual appliances while scaling them up or down, based on demand. This decreases potential points of failure in your network and increases availability.
- **Application Load Balancer:** Application Load Balancer operates at the request level (layer 7), routing traffic to targets (EC2 instances, containers, IP addresses, and Lambda functions) based on the content of the request. Ideal for advanced load balancing of HTTP and HTTPS traffic, Application Load Balancer provides advanced request routing targeted at delivery of modern application architectures, including microservices and container-based applications. Application Load Balancer simplifies and improves the security of your application, by ensuring that the latest SSL/TLS ciphers and protocols are used at all times.

- **Network Load Balancer:** Network Load Balancer operates at the connection level (Layer 4), routing connections to targets (Amazon EC2 instances, microservices, and containers) within Amazon VPC, based on IP protocol data. Ideal for load balancing of both TCP and UDP traffic, Network Load Balancer is capable of handling millions of requests per second while maintaining ultra-low latencies. Network Load Balancer is optimized to handle sudden and volatile traffic patterns while using a single static IP address per Availability Zone. It is integrated with other popular AWS services such as Auto Scaling, Amazon EC2 Container Service (ECS), Amazon CloudFormation, and AWS Certificate Manager (ACM).
- **Classic Load Balancer:** Classic Load Balancer provides basic load balancing across multiple Amazon EC2 instances and operates at both the request level and connection level. Classic Load Balancer is intended for applications that are built within the EC2-Classic network. We recommend Application Load Balancer for Layer 7 traffic and Network Load Balancer for Layer 4 traffic when using Virtual Private Cloud (VPC).

187.1.2. Benefits

- **Security:** When using Amazon Virtual Private Cloud (VPC), you can create and manage security groups associated with Elastic Load Balancing to provide additional networking and security options for Application Load Balancer and Classic Load Balancer. You can configure any of the Load Balancers to be Internet facing or create a load balancer without public IP addresses to serve as an internal (non-internet-facing) load balancer.
- **High availability:** An Elastic Load Balancer is highly available. You can distribute incoming traffic across your Amazon EC2 instances in a single Availability Zone or multiple Availability Zones. An Elastic Load Balancer automatically scales its request handling capacity in response to incoming application traffic. To ensure that your targets are available and healthy, Elastic Load Balancer runs health checks on targets on a configurable cadence.
- **High throughput:** Elastic Load Balancer is designed to handle traffic as it grows and can load balance millions of requests/sec. It can also handle sudden volatile traffic patterns.
- **Health checks:** An Elastic Load Balancer only routes traffic to healthy targets such as EC2 instances, containers, IP addresses, microservices, Lambda functions, and appliances. With Elastic Load Balancing, you get improved insight into the health of your applications in two ways: (1) health check improvements that allow you to configure detailed error codes. The health checks allow you to monitor the health of each of your services behind the the load balancer; and (2) new metrics that give insight into traffic for each of the services running on an EC2 instance.
- **Sticky sessions:** Sticky sessions are a mechanism to route requests from the same client to the same target. Elastic Load Balancers support sticky sessions. Stickiness is defined at a target group level.
- **Operational monitoring & logging:** Amazon CloudWatch reports Application and Classic Load Balancer metrics such as request counts, error counts, error types, request latency, and more. Amazon CloudWatch also tracks Network and Gateway Load Balancer metrics such as Active Flow Count, New Flow Count, Processed Bytes, and more. Elastic Load Balancers are also integrated with AWS CloudTrail which tracks API calls to the ELB.
- **Delete protection:** You can enable deletion protection on an Elastic Load Balancer to prevent it from being accidentally deleted.

187.2. Backup/Restore and Disaster Recovery

This requirement is not applicable for this service. For additional information beyond what is described herein, please refer to http://aws.amazon.com/documentation/.

187.3. Pricing Overview

Please see the AWS UK G Cloud 13 Pricing Document accompanying this service in the Digital Marketplace.

187.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/elasticloadbalancing/index.html
- **Service quotas:** https://docs.aws.amazon.com/elasticloadbalancing/latest/application/load-balancer-limits.html
- **Service FAQs:** https://aws.amazon.com/elasticloadbalancing/faqs/?nc=sn&loc=5

187.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/elasticloadbalancing/index.html and the following links for comprehensive technical documentation regarding this service.

- User Guide: Learn about Elastic Load Balancing.
- User Guide for Application Load Balancers: Use Application Load Balancers for HTTP and HTTPS traffic. The load balancer routes based on the content of the request.
- User Guide for Network Load Balancers: Use Network Load Balancers for TCP, UDP, and TLS traffic where extreme performance is required.
- User Guide for Gateway Load Balancers: Use Gateway Load Balancers to deploy, scale, and manage virtual appliances, such as firewalls.
- AWS CLI Reference: Provides syntax and examples for the commands.
- API Reference: Provides syntax and examples for the API actions.
- User Guide for Classic Load Balancers: Use Classic Load Balancers with applications in the EC2-Classic network.
- API Reference: Provides syntax and examples for the API actions.
- AWS CLI Reference: Provides syntax and examples for the commands.

188. Cross-Service Definitions

The following service definition topics are applicable to all AWS Service Offerings and are detailed once in a cross-service manner below.

188.1. Availability

AWS has the largest global infrastructure footprint of any provider, and this footprint is constantly increasing at a significant rate. The AWS Cloud infrastructure is built around AWS Regions and Availability Zones. A Region is a physical location in the world where we have multiple Availability Zones. Availability Zones consist of one or more discrete data centres, each

with redundant power, networking, and connectivity and housed in separate facilities. These Availability Zones offer customers the ability to operate production applications and databases that are more highly available, fault tolerant, and scalable than would be possible with a single data centre.

AWS currently has 25 Regions and 81 Availability Zones throughout the world—including 14 Local Zones and 17 Wavelength Zones for ultralow latency applications. AWS Regions include: US East (N. Virginia), US East (Ohio), US West (Oregon), US West (N. California), AWS GovCloud (US-West), AWS GovCloud (US-East), Canada (Central), Europe (Ireland), Europe (Frankfurt), Europe (London), Europe (Milan), Europe (Paris), Europe (Stockholm), Middle East (Bahrain), Africa (Cape Town), Asia Pacific (Singapore), Asia Pacific (Tokyo), Asia Pacific (Sydney), Asia Pacific (Seoul), Asia Pacific (Mumbai), Asia Pacific (Osaka), Asia Pacific (Hong Kong), South America (Sao Paulo), China (Beijing), and China (Ningxia).

AWS has also announced plans for nine more AWS Regions in Australia, Canada, India, Indonesia, Israel, New Zealand, Spain, Switzerland, and United Arab Emirates (UAE).

Information about each Region can be found at the AWS Global Infrastructure page.

Figure 1 depicts the current AWS Regions and Availability Zones, along with the nine announced Regions.

Figure 1. AWS Regions and Availability Zones

The AWS products and services that are available in each Region are listed at the Region Table webpage.

Figure 2 illustrates the relationship between AWS Regions and Availability Zones.

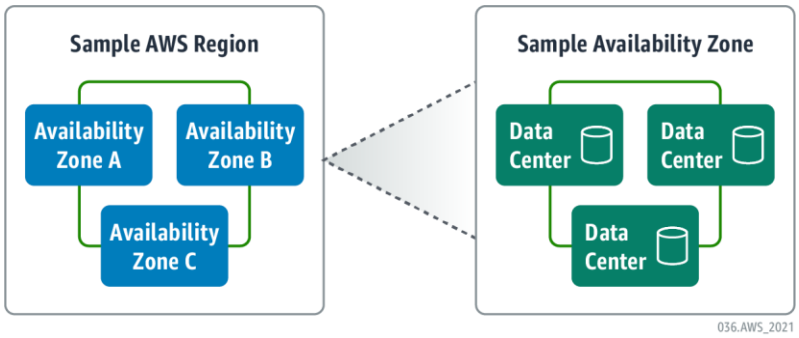

Figure 2. Relationship Between AWS Regions and Availability Zones

188.1.1. Region Availability

Exact service availability depends on a range of factors and choices made by customers when they architect and implement their solution.

The Services Offerings will be delivered from the AWS Region selected by the customer upon opening an AWS account. The customer may specify the AWS Region in which customer content will be stored. It is the customer's responsibility to select the relevant AWS Region in order to comply with its own security and governance requirements. AWS will not access or use customer content except as necessary to maintain or provide the Service Offerings, or as required by law or regulation. Customers acknowledge that AWS does not limit customers to any particular AWS Region. Note that not all AWS Cloud services are available in every AWS

Region; however, we are steadily expanding our service availability across AWS's global regions.

The full list of available AWS services, and their availability by region can be seen on our website at https://aws.amazon.com/about-aws/global-infrastructure/regional-product-services/.

188.1.2. Designing for Availability and Reliability

While AWS goes to great lengths to provide availability and reliability of the cloud, our customers share responsibility for ensuring availability and reliability within the cloud. Some best practices we recommend for building highly resilient systems in the AWS Cloud include designing for failure, automating failover and recovery, testing your recovery procedures, and accessing resources and reference architectures.

188.1.2.1. Design for Failure across Multiple Availability Zones

Although rare, failures can occur that affect the availability of resources that are hosted in the same Availability Zone. If you host all your resources in a single Availability Zone that is affected by such a failure, none of these resources would be available. It is therefore a best practice to architect across multiple Availability Zones in the same Region to achieve extremely high recovery time objectives (RTOs), recovery point objectives (RPOs), and service availability. Availability Zones are connected to each other with fast, private fibre-optic networking, enabling you to easily architect applications that automatically fail over between zones without interruption.

For mission-critical applications, it is a best practice to architect across Regions to handle the rare case of an entire Region failing—perhaps as a result of a major physical attack. You can do so using both private, high-speed networking and public internet connections to provide an additional layer of business continuity or to provide low-latency access across the globe.

188.1.2.2. Automate Failover and Recovery

By monitoring a system for key performance indicators (KPIs), you can trigger automation when a threshold is breached. These KPIs should be a measure of business value and not of the technical aspects of the service operation. This allows for automatic notification and tracking of failures, and automated recovery processes that work around or repair the failure. With sophisticated automation, it is possible to anticipate and remediate failures before they occur.

188.1.2.3. AWS Personal Health Dashboard

AWS Personal Health Dashboard provides alerts and remediation guidance when AWS is experiencing events that may impact you. The dashboard displays relevant and timely information to help you manage events in progress, and provides proactive notification to help you plan for scheduled activities. Alerts are triggered by changes in the health of AWS resources, giving you event visibility and guidance to help quickly diagnose and resolve issues.

The dashboard provides forward-looking notifications, and you can set up alerts across multiple channels, including email and mobile, so you receive timely and relevant information to help plan for scheduled changes that may affect you. In the event of AWS hardware maintenance activities that may impact one of your Amazon Elastic Compute Cloud (Amazon EC2) instances, for example, you would receive an alert with information to help you plan for, and proactively address, any issues associated with the upcoming change.

Personal Health Dashboard can integrate with Amazon CloudWatch Events, enabling you to build custom rules and select targets such as AWS Lambda functions to define automated remediation actions. The AWS Health API, which powers Personal Health Dashboard, allows

you to integrate health data and notifications with your existing in-house or third-party IT management tools.

188.1.3. Test Your Recovery Procedures

One of the benefits of the cloud is that you can test how your system fails, and you can validate your recovery procedures. You can use a test environment to simulate different failures or recreate scenarios that led to failures already. This exposes failure pathways that you can test and fix before a real failure scenario, reducing the risk of components that have not been tested before failing.

188.1.4. Resources and Reference Architecture

The resources described below help customers to understand AWS Cloud services and features and provide architectural guidance on designing and implementing systems that run on the AWS infrastructure:

- The AWS Well-Architected Framework codifies the experiences of thousands of customers, helping them assess and improve their cloud-based architectures and mitigate disruptions.
- The AWS Architecture Center is designed to provide customers with the necessary guidance and application architecture best practices to build highly scalable and reliable applications in the AWS Cloud.
- The AWS Outposts High Availability Design and Architecture Considerations whitepaper discusses architecture considerations and recommended practices to build highly available on-premises application environments with AWS Outposts.
- The AWS advanced continuous delivery best practices video.

188.2. On-Boarding/Off-Boarding Processes and Service Migration

AWS maintains a cadre of Getting Started Guides and schedules regular webinars. These guides and webinars cover a variety of topics, see http://aws.amazon.com/documentation/gettingstarted/ for more details.

AWS allows customers to move data as needed off AWS storage using the public internet or AWS Cloud services such as AWS Direct Connect, AWS Import/Export, and more.

With AWS, you can provision compute power, storage, and other resources, gaining access to a suite of elastic IT infrastructure services as your business demands them. With minimal cost and effort, you can move your application to the AWS Cloud and reduce capital expenses, minimise support and administrative costs, and retain the performance, security, and reliability requirements your business demands. To see a step-by-step migration strategy, refer to the Migrating Your Existing Applications to the AWS Cloud whitepaper.

188.3. Service Management Details

AWS Cloud services are driven by robust APIs that allow for a wide variety of monitoring, management and developer tools to integrate easily with AWS Cloud resources. Common tools from vendors such as Microsoft, VMware, BMC Software, Okta, RightScale, Eucalyptus, CA, Xceedium, Symantec, Racemi, and Dell are supported on AWS. This flexibility allows AWS customers to easily provision, manage, and monitor all of their IT resources through a “single pane of glass” with the tool that best fits their unique needs. This also means that a full

inventory of those resources is only a few clicks away. Below are various AWS-native management options.

188.3.1. AWS Management Console

The AWS Management Console is a single destination for managing all AWS resources, from Amazon Elastic Compute Cloud (Amazon EC2) instances to Amazon DynamoDB tables. Customers can use the AWS Management Console to perform any number of tasks, from deploying new applications to monitoring the health of applications. The AWS Management Console also enables customers to manage all aspects of their AWS account, including accessing monthly spending by service, managing security credentials, or even setting up new AWS Identity and Access Management (AWS IAM) users. The AWS Management Console supports all AWS Regions and lets customers provision resources across multiple AWS Regions.

188.3.2. AWS Developer Tools

AWS offerings are provided with a range of supporting components like management tools, networking services, and application augmentation services, with multiple interfaces to AWS Application Programming Interface (API)-based services, including Software Development Kits (SDKs), Integrated Development Environment (IDE) toolkits, and Command Line Tools. Browse by programming language for tools to develop and manage applications on AWS at our Tools to Build on AWS page.

AWS Developer Tools help you securely store and version control your application's source code and automatically build, test, and deploy your application to AWS or your on-premises environment. They are built to work with AWS, making it easier for your team to get set up and be productive.

AWS Developer Tools are designed to help you build software like Amazon. They facilitate practices such as continuous delivery and infrastructure as code for serverless, containers, and Amazon EC2.

188.3.3. Management and Governance Tools

In the past, organizations have had to choose between innovating faster and maintaining control over cost, compliance, and security. With AWS Management and Governance services, customers don’t have to choose between innovation and control—they can have both. With AWS, customers can enable, provision, and operate their environment for both business agility and governance control.

- **Scale** – AWS Management and Governance services are built to manage highly dynamic cloud resources at massive scale.
- **Simplicity** – AWS reduces complexity, offering a single control plane for customers to manage and govern their resources on AWS and on-premises.
- **Third-party solutions** – AWS offers the broadest partner ecosystem for customers to extend and augment their management and governance system.
- **Cost savings** – Customers can use AWS Management and Governance services to assess their resource utilization and identify ways to reduce costs.

188.4. Service Levels and Service Credits

AWS currently provides SLAs, with a corresponding Service Credit regime, for several products. Due to the rapidly evolving nature of AWS's product offerings, SLAs are best reviewed directly on our website via the links below:

https://aws.amazon.com/api-gateway/sla/

https://aws.amazon.com/appstream2/sla/

https://aws.amazon.com/athena/sla/

https://aws.amazon.com/rds/aurora/sla/

https://aws.amazon.com/braket/sla/

https://aws.amazon.com/chime/sla/

https://aws.amazon.com/cloud-directory/sla/

https://aws.amazon.com/cloudfront/sla/

https://aws.amazon.com/cloudsearch/sla/

https://aws.amazon.com/cloudwatch/sla/

https://aws.amazon.com/cognito/sla/

https://aws.amazon.com/compute/sla/

https://aws.amazon.com/connect/sla/

https://aws.amazon.com/it/detective/sla/

https://aws.amazon.com/devops-guru/sla/

https://aws.amazon.com/documentdb/sla/

https://aws.amazon.com/dynamodb/sla/

https://aws.amazon.com/ecs/anywhere/sla/

https://aws.amazon.com/efs/sla/

https://aws.amazon.com/ecr/sla/

https://aws.amazon.com/eks/sla/

https://aws.amazon.com/elasticloadbalancing/sla/

https://aws.amazon.com/elastictranscoder/sla/

https://aws.amazon.com/emr/sla/

https://aws.amazon.com/eventbridge/sla/

https://aws.amazon.com/finspace/sla/

https://aws.amazon.com/forecast/sla/

https://aws.amazon.com/fraud-detector/sla/

https://aws.amazon.com/fsx/sla/

https://aws.amazon.com/guardduty/sla/

https://aws.amazon.com/healthlake/sla/

https://aws.amazon.com/inspector/sla/
https://aws.amazon.com/ivs/sla/
https://aws.amazon.com/kendra/sla/
https://aws.amazon.com/keyspaces/sla/
https://aws.amazon.com/kinesis/sla/
https://aws.amazon.com/lightsail/sla-lightsail-instances-and-block-storage/
https://aws.amazon.com/lightsail/sla-lightsail-managed-databases/
https://aws.amazon.com/location/sla/
https://aws.amazon.com/lookout-for-equipment/sla/
https://aws.amazon.com/machine-learning/language/sla/
https://aws.amazon.com/macie/sla/
https://aws.amazon.com/managed-blockchain/sla/
https://aws.amazon.com/grafana/sla/
https://aws.amazon.com/msk/sla/
https://aws.amazon.com/managed-workflows-for-apache-airflow/sla/
https://aws.amazon.com/memorydb/sla/
https://aws.amazon.com/messaging/sla/
https://aws.amazon.com/monitron/sla/
https://aws.amazon.com/amazon-mq/sla/
https://aws.amazon.com/neptune/sla/
https://aws.amazon.com/nimble-studio/sla/
https://aws.amazon.com/elasticsearch-service/sla/
https://aws.amazon.com/personalize/sla/
https://aws.amazon.com/qldb/sla/
https://aws.amazon.com/quicksight/sla/
https://aws.amazon.com/rds/proxy/sla/
https://aws.amazon.com/redshift/sla/
https://aws.amazon.com/rekognition/sla/
https://aws.amazon.com/rds/sla/
https://aws.amazon.com/route53/sla/
https://aws.amazon.com/s3/sla-rtc/
https://aws.amazon.com/sagemaker/sla/
https://aws.amazon.com/s3/sla/
https://aws.amazon.com/swf/sla/

https://aws.amazon.com/simpledb/sla/
https://aws.amazon.com/textract/sla/
https://aws.amazon.com/timestream/sla/
https://aws.amazon.com/pinpoint/sla/
https://aws.amazon.com/vpc/ipam-sla/
https://aws.amazon.com/vpc/sla/
https://aws.amazon.com/workdocs/sla/
https://aws.amazon.com/worklink/amazon-worklink-service-level-agreement/
https://aws.amazon.com/workmail/amazon-workmail-service-level-agreement/
https://aws.amazon.com/workspaces/sla/
https://aws.amazon.com/amplify/sla/
https://aws.amazon.com/application-migration-service/sla/
https://aws.amazon.com/appsync/sla/
https://aws.amazon.com/audit-manager/sla/
https://aws.amazon.com/backup/sla/
https://aws.amazon.com/aws-cost-management/aws-budgets/sla/
https://aws.amazon.com/certificate-manager/private-certificate-authority/sla/
https://aws.amazon.com/vpn/client-vpn-sla/
https://aws.amazon.com/cloud-map/sla/
https://aws.amazon.com/cloudhsm/sla/
https://aws.amazon.com/cloudtrail/sla/
https://aws.amazon.com/codeartifact/sla/
https://aws.amazon.com/codebuild/sla/
https://aws.amazon.com/codecommit/sla/
https://aws.amazon.com/codedeploy/sla/
https://aws.amazon.com/codepipeline/sla/
https://aws.amazon.com/compute-optimizer/sla/
https://aws.amazon.com/config/sla/
https://aws.amazon.com/aws-cost-management/aws-cost-explorer/sla/
https://aws.amazon.com/datapipeline/sla/
https://aws.amazon.com/dms/sla/
https://aws.amazon.com/device-farm/sla/
https://aws.amazon.com/directconnect/sla/
https://aws.amazon.com/directoryservice/sla/

https://aws.amazon.com/disaster-recovery/sla/
https://aws.amazon.com/mediaconnect/sla/
https://aws.amazon.com/mediaconvert/sla/
https://aws.amazon.com/medialive/sla/
https://aws.amazon.com/mediapackage/sla/
https://aws.amazon.com/mediastore/sla/
https://aws.amazon.com/mediatailor/sla/
https://aws.amazon.com/firewall-manager/sla/
https://aws.amazon.com/global-accelerator/sla/
https://aws.amazon.com/glue/sla/
https://aws.amazon.com/ground-station/sla/
https://aws.amazon.com/transfer/sla/
https://aws.amazon.com/iot-1-click/sla/
https://aws.amazon.com/iot-analytics/sla/
https://aws.amazon.com/iot-core/sla/
https://aws.amazon.com/iot-device-defender/sla/
https://aws.amazon.com/iot-device-management/sla/
https://aws.amazon.com/iot-events/sla/
https://aws.amazon.com/greengrass/sla/
https://aws.amazon.com/iot-sitewise/sla/
https://aws.amazon.com/iot-things-graph/sla/
https://aws.amazon.com/kms/sla/
https://aws.amazon.com/lambda/sla/
https://aws.amazon.com/migration-hub/sla/refactor-spaces/
https://aws.amazon.com/network-firewall/sla/
https://aws.amazon.com/opsworks/sla/
https://aws.amazon.com/privatelink/sla/
https://aws.amazon.com/resilience-hub/sla/
https://aws.amazon.com/robomaker/sla/
https://aws.amazon.com/secrets-manager/sla/
https://aws.amazon.com/security-hub/sla/
https://aws.amazon.com/servicecatalog/sla/
https://aws.amazon.com/shield/sla/
https://aws.amazon.com/vpn/site-to-site-vpn-sla/

https://aws.amazon.com/step-functions/sla/

https://aws.amazon.com/systems-manager/sla/

https://aws.amazon.com/transit-gateway/sla/

https://aws.amazon.com/waf/sla/

https://aws.amazon.com/xray/sla/

https://aws.amazon.com/elasticache/sla/

See the Supplier Terms document affiliated with this framework catalogue for additional information.

188.5. Trial Service Details

The AWS Free Tier provides customers the ability to explore and try out AWS services free of charge up to specified limits for each service. The Free Tier is comprised of three different types of offerings, a 12-month Free Tier, an Always Free offer, and short term trials. Services with a 12-month Free Tier allow customers to use the product for free up to specified limits for one year from the date the account was created. Services with an Always Free offer allow customers to use the product for free up to specified limits as long as they are an AWS customer. Services with a short-term trial are free to use for a specified period of time or up to a one-time limit depending on the service selected.

Details on the limits and services provided for free are detailed in each card on the Free Tier page. If your application use exceeds the free tier limits, you simply pay standard, pay-as-you-go service rates (see each service page for full pricing details). Restrictions apply; see offer terms for more details.

188.6. Data Restoration/Service Migration

188.6.1. AWS Backup

Backing up your data is an important step towards protecting your applications and ensuring that you meet your business and regulatory backup compliance requirements. Even durable resources are susceptible to threats (e.g., bugs in your application) that could cause accidental deletions or corruption. Building and managing your own backup workflows across all your applications in a compliant and consistent manner can be complex and costly. AWS Backup removes the need for costly, custom solutions or manual processes.

AWS Backup is a fully managed backup service that makes it easy to centralize and automate the back up of data across AWS Cloud services, as well as on premises using the AWS Storage Gateway. With AWS Backup, you can create automated backup policies called backup plans, such as how frequently to back up your data and how long to retain those backups. Together with AWS Organizations, AWS Backup enables you to centrally deploy data protection (backup) policies to configure, manage, and govern your backup activity across your organization's AWS accounts and AWS resources, such as:

- Amazon Elastic Compute Cloud (Amazon EC2) instances
- Amazon Elastic Block Store (Amazon EBS) volumes
- Amazon Relational Database Service (Amazon RDS) databases, including Amazon Aurora clusters
- Amazon DynamoDB tables

- Amazon Elastic File System (Amazon EFS) file systems
- Amazon FSx for Lustre
- Amazon FSx for Windows File Server
- AWS Storage Gateway volumes

You can use AWS Backup to create backup policies that automate backup schedules and retention management. As illustrated in **Figure 3**, AWS Backup provides a fully managed, policy-based solution, simplifying backup management and enabling you to meet business and regulatory compliance requirements.

Figure 3. AWS Backup provides a fully managed, policy-based backup solution

188.6.2. Disaster Recovery

Businesses and public sector organizations of all sizes are using AWS to enable faster and cheaper disaster recovery (DR) of their critical IT systems. Having your own DR site ready and on standby in the cloud, without having to pay for the IT infrastructure, makes the AWS Cloud a perfect solution for DR. With the AWS Cloud, not only can you recover quickly from a disaster and ensure business continuity while keeping costs down, you also can make it easy, secure, and reliable.

With its fault-tolerant architecture and 200+ service offerings, AWS supports many DR architectures—from those built for smaller workloads to enterprise solutions that enable rapid failover at scale. These architectures include “pilot light” environments that are ready to scale up rapidly to “hot standby” environments that fail over at a moment’s notice. As a customer, you have the flexibility to choose the right approach for your DR strategy, depending on your recovery time objective (RTO) and recovery point objective (RPO) goals and budget.

Additionally, AWS offers a business continuity solution called CloudEndure Disaster Recovery that minimizes downtime and data loss by providing fast, reliable, cloud-based DR. The solution continuously replicates applications from physical, virtual, or cloud-based infrastructure to a low-cost staging area that is automatically provisioned in any target AWS Region of your choice. During failover or testing, an up-to-date copy of an application can be spun up on demand and be fully functioning in minutes.

188.6.2.1. High Availability (HA) in the AWS Cloud

A discussion of DR is not complete without addressing high availability (HA). While DR focuses on bringing systems back online once disaster strikes, HA focuses on ensuring that there is no single point of failure in your architecture from the outset. The AWS Cloud global infrastructure,

with its construct of Regions and Availability Zones, is designed for HA. Using AWS Regions (geographically isolated components of the global infrastructure) and Availability Zones (fully fault-tolerant clusters of data centres with redundant power, networking, and connectivity), you can build inherent HA and fault tolerance into your architecture. The AWS global infrastructure allows us to deliver the highest network availability of any cloud provider, with 7x fewer downtime hours than the next largest cloud provider.[1]

Additionally, to better isolate any issues and achieve HA, you can partition applications across multiple Availability Zones in the same AWS Region. AWS control planes and the AWS Management Console are distributed across Regions and include regional application programming interface (API) endpoints. These endpoints are designed to operate securely for at least 24 hours, if isolated from the global control plane functions, without requiring customers to access the Region or its API endpoints via external networks during any isolation.

Both HA and DR rely on some of the same best practices, such as monitoring for failures, deploying to multiple locations, and automatic failover. However, HA focuses on components of the workload, whereas DR focuses on discrete copies of the entire workload. DR has different objectives from HA, measuring time to recovery after the larger scale events that qualify as disasters. You should first ensure your workload meets your availability objectives, as a highly available architecture will enable you to meet customers' needs in the event of availability impacting events. Your DR strategy requires different approaches than those for availability and should focus on deploying discrete systems to multiple locations so that you can failover the entire workload if necessary.

188.6.2.2. Benefits of Using AWS for DR

With AWS, you can eliminate the need for additional physical infrastructure, offsite data replication, and upkeep of spare capacity. Availability Zones are a key AWS feature as they make partitioning applications for high availability easy. This enables customers to protect applications from the failure of a single location, resulting in significant cost savings and increased agility to change and optimize resources during a DR scenario.

AWS provides fine-grained control and building blocks to create the appropriate DR solution in the cloud, given a customer's unique resiliency requirements, recovery objectives (RTO and RPO as shown in **Figure 4**), and budget. With AWS, customers can build highly resilient applications while taking advantage of flexible, cost-effective infrastructure solutions.

[1] Based on downtime hours from 1/1/18 to 12/31/18 pulled directly from the public service health dashboards of the major cloud providers.

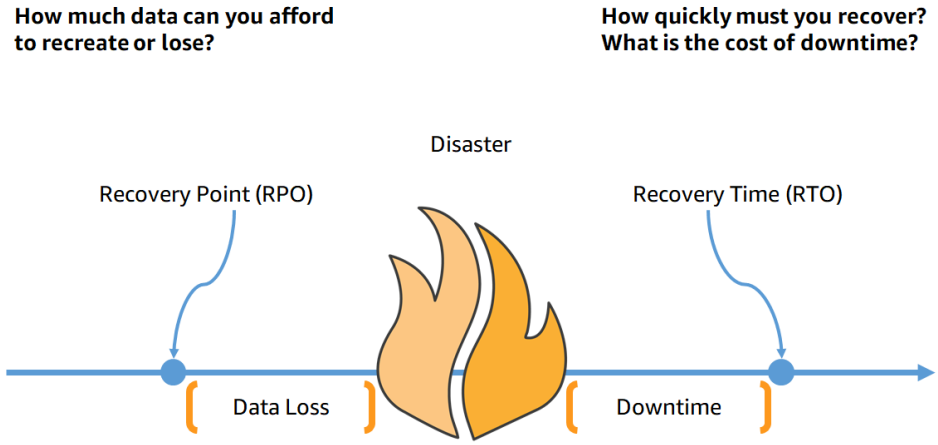

Figure 4. RPO and RTO. RPO is related to data loss in the event of disaster; RTO is related to the amount of time systems are down in the event of disaster.

Figure 5 shows a spectrum of scenarios—multi-site, warm standby, pilot light, and backup and restore—arranged by how quickly a system can be available to users after a DR event. Typically, the shorter the recovery, the higher the cost of the solution.

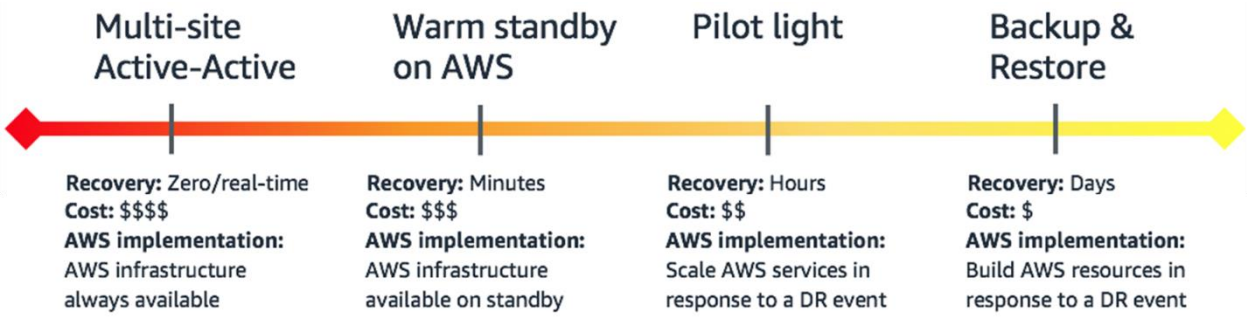

Figure 5. Spectrum of DR Options. Customers can choose a preferred DR option based on preferences for cost and the time it takes to recover systems.

Each DR option is discussed in more detail below.

Backup and Restore

In most traditional environments, data is backed up to tape and sent offsite regularly. Recovery time will be the longest using this method, and lack of automation leads to increased costs. Amazon Simple Storage Service (Amazon S3) is ideal for backup data, as it is designed to provide 99.999999999% durability of objects over a given year. Amazon S3 offers a range of storage classes—including options for one-zone infrequent access and archive—to help you save on costs. Transferring data to and from Amazon S3 is typically done via the network, and it is therefore accessible from any location. To centralize and automate your backup, AWS Backup allows you to configure backup policies and monitor backup activity.

Pilot Light for Simple Recovery into AWS Warm Standby Solution

The idea of the pilot light is an analogy that comes from gas heating. In that scenario, a small flame that's always on can quickly ignite the entire furnace to heat up a house. In this DR approach, you simply replicate part of your IT structure for a limited set of core services so that the AWS cloud environment seamlessly takes over in the event of a disaster. A small part of your infrastructure is always running simultaneously syncing mutable data (as databases or

documents), while other parts of your infrastructure are switched off and used only during testing. Unlike a backup and recovery approach, you must ensure that your most critical core elements are already configured and running in AWS (the pilot light). When the time comes for recovery, you can rapidly provision a full-scale production environment around the critical core.

Warm Standby Solution on AWS

The term “warm standby” is used to describe a DR scenario in which a scaled-down version of a fully functional environment is always running in the cloud. It further decreases recovery time because, in this case, some services are always running. By identifying business-critical systems, you could fully duplicate these systems on AWS and have them always on. These servers can be running a minimum-sized fleet of the smallest Amazon Elastic Compute Cloud (Amazon EC2) instances possible. This solution is not scaled to handle a full-production load, but it is fully functional. It can be used for non-production work, such as testing, quality assurance, and internal use. In a disaster, you can scale out the system quickly to handle the production load by adding more instances. You can automate this process using Amazon EC2 Auto Scaling and Elastic Load Balancing.

Multi-Site Solution Deployed on AWS and Onsite

A multi-site solution runs in AWS as well as on your existing on-site infrastructure, in an active-active configuration. The data replication method that you employ will be determined by the recovery point that you choose. You can use a Domain Name System (DNS) service that supports weighted routing, such as Amazon Route 53, to route production traffic to different sites that deliver the same application or service. A proportion of traffic will go to your infrastructure in AWS, and the remainder will go to your on-site infrastructure.

In an on-site disaster situation, you can adjust the DNS weighting and send all traffic to the AWS servers. The capacity of the AWS Cloud service can be rapidly increased to handle the full production load. You can use Amazon EC2 Auto Scaling to automate this process. You might need some application logic to detect the failure of the primary database services and cut over to the parallel database services running in AWS.

The cost of this scenario is determined by how much production traffic is handled by AWS during normal operation. In the recovery phase, you pay only for what you use for the duration that the DR environment is required at full scale. You can further reduce cost by purchasing Amazon EC2 Reserved Instances for your “always on” AWS servers.

188.6.2.3. AWS Elastic Disaster Recovery

AWS Elastic Disaster Recovery enables organizations to minimize downtime and data loss with fast, reliable recovery of on-premises and cloud-based applications. This service minimizes downtime and data loss with fast, reliable recovery of on-premises and cloud-based applications using affordable storage, minimal compute, and point-in-time recovery.

AWS Elastic Disaster Recovery continuously replicates your source servers to your AWS account without impacting performance. It reduces costs compared to traditional on-premises disaster recovery solutions by removing idle recovery site resources, and instead leverages affordable AWS storage and minimal compute resources to maintain ongoing replication. If you need to recover applications on AWS, you can do so within minutes.

During recovery, you can choose the most up-to-date server state as a recovery point, or choose to recover an operational copy of your applications from an earlier point in time. Point in time recovery is helpful for recovery from data corruption events such as ransomware. After

issues are resolved in your primary environment, you can use AWS Elastic Disaster Recovery to fail back your recovered applications.

Figure 6 illustrates how AWS Elastic Disaster Recovery operates.

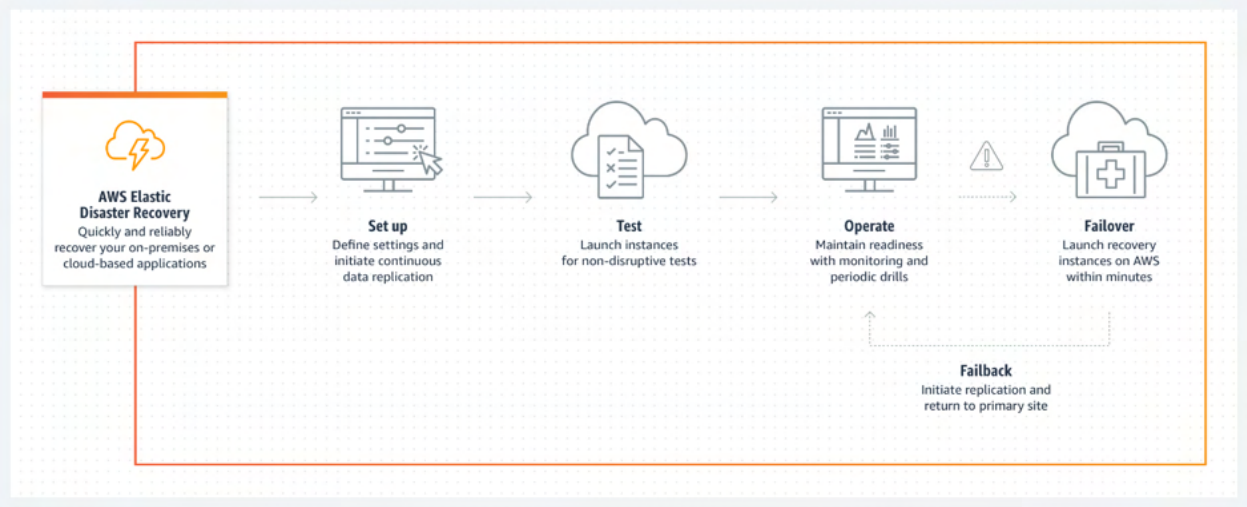

Figure 6. AWS Elastic Disaster Recovery Architecture.This service simplifies recovery of a wide range of applications on AWS and uses a unified process for drills, recovery, and failback.

AWS Elastic Disaster Recovery is the recommended service for disaster recovery to AWS. It provides similar capabilities as CloudEndure Disaster Recovery (see **Section Error! Reference source not found.**) and is operated from the AWS Management Console. This enables seamless integration between AWS DRS and other AWS services, such as AWS CloudTrail, AWS Identity and Access Management (IAM), and Amazon CloudWatch.

188.6.2.4. Complementary AWS Cloud Services for Backup/DR

In addition to CloudEndure Disaster Recovery, many other AWS Cloud services can be used to provide complementary capabilities to enable a successful DR strategy. These services include Amazon EC2, Amazon S3, AWS Storage Gateway, Amazon Relational Database Service (Amazon RDS), Amazon CloudFront, and others.

For example, Amazon RDS is a fully managed relational database service that supports automated backups by default. It can be set up using a Multi-AZ configuration to reduce risk from outages and minimize loss of valuable data. Amazon S3 is designed for 99.999999999% (11 9s) durability and stores data for millions of applications for companies around the world. Amazon EC2 and AWS Storage Gateway enable you to back up key data stores in the AWS Cloud. AWS Backup provides a common way to manage backups across several AWS Cloud services, both in the cloud and on premises. Amazon Elastic File System (Amazon EFS) and Amazon EBS provide HA and durability for file and block storage so that your backups are protected and available when needed.

More details can be found at these links:

- Learn to Build on AWS: Backup and Recovery
- AWS DR Resources
- Backup & Restore Services with AWS

188.6.2.5. DR Resources

There are multiple resources to help organizations start using AWS for a DR solution:

- Read the whitepaper Affordable Enterprise-Grade Disaster Recovery Using AWS.
- Read about CloudEndure Disaster Recovery, which offers fast, cost-effective business continuity for your mission-critical workloads.
- Read the eBook Leverage the Cloud for your Disaster Recovery Strategy.
- Download the IT Disaster Recovery Plan Checklist to ensure you are on track.

188.6.3. Resiliency Planning

An organization's resiliency and continuity plan outlines a range of disaster scenarios and the steps the organization will take to return to a regular state. Plans are written ahead of time, by key staff and multiple departments, with the goal of creating contingencies that minimize potential harm and negative impacts to the organization.

A strong resilience strategy combines both operational and cultural resilience. Operational resilience includes five pillars:

1. Remote workforce enablement
2. Constituent engagement
3. Operational continuity
4. Real-time analytics
5. Process and systems modernization

Figure 7 below illustrates the AWS Organizational Resiliency Framework and displays which core foundational technologies can assist when creating a strong resiliency plan, for each of the five pillars.

Figure 7. AWS Organizational Resiliency Framework. This framework can improve an organization's odds of minimizing the impact of an emergency on employees, customers, and partners.

A resilience plan also addresses cultural resiliency best practices that prepare your team for any disruption. Beyond pandemics or natural disasters, organizations can be adversely impacted by changing market conditions, mergers, or general restructuring, which can also be classified as emergencies. Regardless of the type of disruption, an employee's alignment to organizational objectives is most at stake during times of disruption. Other elements that influence employees during crisis include the existing organizational culture, leadership, and one's work environment—which can include technologies, tools, physical space, and processes. Your resiliency plan should identify employee risks, just as it identifies technology, customer, and financial risks.

Organizations should start by building a baseline resiliency plan. No resiliency plan is bullet-proof; however, having a plan in place drastically improves your organization's odds of minimizing the impact of an emergency on your employees, customers, and partners.

188.6.3.1. Business Continuity Plan (BCP)

Your DR plan should be a subset of your organization's business continuity plan (BCP). There is no point in maintaining aggressive DR targets for restoring a workload if that workload's business objectives cannot be achieved because of the disaster's impact on elements of your business other than your workload.

For organizations serving the public, disasters and emergencies can derail missions and critical emergency response, public safety, and public health services. Many organizations move from crisis to crisis with short-term fixes, without addressing the underlying lack of a long-term strategy for resilience and business continuity. Even with an existing resiliency plan in place in one department, the plan may not cover process re-engineering requirements and dependencies in other departments, like IT, finance, human resources, legal, communications, operators, or others.

Key disruption areas to address with a long-term BCP include the most common and high-impact disruptors such as the following:

- Ensuring IT infrastructure durability, availability, and security
- Supporting employees and tapping expertise
- Access to data for real-time situational awareness
- Lack of financial agility
- Lack of planning

Business Impact Analysis

As part of a BCP, you should carry out a business impact analysis to quantify the business impact of a disruption to your workloads. It should identify the impact on internal and external customers of not being able to use your workloads and the effect that has on your business. The analysis should help to determine how quickly the workload needs to be made available and how much data loss can be tolerated. However, it is important to note that recovery objectives should not be made in isolation; the probability of disruption and cost of recovery are key factors that help to inform the business value of providing DR for a workload.

Business impact may be time dependent. You may want to consider factoring this into your DR planning. For example, disruption to your payroll system is likely to have a very high impact to

the business just before everyone gets paid, but it may have a low impact just after everyone has already been paid.

Risk Assessment

A risk assessment of the type of disaster and geographical impact along with an overview of the technical implementation of your workload will determine the probability of disruption occurring for each type of disaster.

For highly critical workloads, you may consider high availability across multiple Regions with continuous backups in place to minimize business impact. For less critical workloads, a valid strategy may be not to have any DR in place at all. And for some disaster scenarios, it is also valid not to have any DR strategy in place as an informed decision based on a low probability of the disaster occurring. Remember that Availability Zones within an AWS Region are already designed with meaningful distance between them, and careful planning of location, such that most common disasters should only impact one zone and not the others. Therefore a multi-AZ architecture within an AWS Region may already meet your risk mitigation needs.

The cost of the DR options should be evaluated to ensure that the DR strategy provides the correct level of business value considering the business impact and risk.

With all of this information, you can document the threat, risk, impact and cost of different disaster scenarios and the associated recovery options. This information should be used to determine your recovery objectives for each of your workloads.

188.6.4. Migration

Many customers seek to migrate their workloads to the AWS Cloud to benefit from IT costs savings. These workloads include applications, websites, databases, storage, physical or virtual servers, or entire data centers—all residing in the on-premises environment, hosting facility, or other public cloud.

The AWS Cloud allows you to provision elastic compute power, storage, and other resources in the cloud on demand. Moving your applications to the AWS Cloud can help you reduce upfront expenses by using our pay-as-you-go pricing structure. You can also minimize support and administrative costs through our shared responsibility model and use of AWS Support—all while retaining performance, security, and reliability requirements. At AWS, our research and experience have led us to identify the following key drivers that compel businesses to migrate to the cloud.

- Substantial IT costs savings
- Digital transformation
- Improvements in staff productivity and business agility
- Improved security and operational resilience
- Data center consolidation
- Approaching hardware/software end-of-life
- Going global, mergers, and acquisitions
- Exploring new technologies (e.g., artificial intelligence/machine learning and Internet of Things)

Migrating to the AWS Cloud can enable your organization to reduce operational costs by up to 51% and bring products and services to market 18.8% faster. AWS has helped thousands of organizations, including enterprises such as GE, the Coca-Cola Company, BP, Enel, Samsung, NewsCorp, and Twenty-First Century Fox migrate to the cloud, freeing-up resources by lowering IT costs and improving productivity, operational resiliency, and business agility.

To ensure that you have a holistic view of the transformation initiative that is required for an effective move to the cloud and to create a foundation for your cloud strategy that returns ongoing measurable value to your organization, AWS offers the Cloud Adoption Framework. The AWS Cloud Adoption Framework enables you to analyze your environment through different perspectives: Business, People, Governance, Platform, Security, and Operations. This gives you a complete view of which areas to improve before moving forward with a large migration effort.

188.6.4.1. Migration Patterns

A strong migration plan starts with a deeper understanding of the interdependencies between applications, and it evaluates migration strategies to meet your business case objectives. **Figure 8 – Seven Common Migration** depicts the seven most common migration patterns that help you create a modern migration strategy.

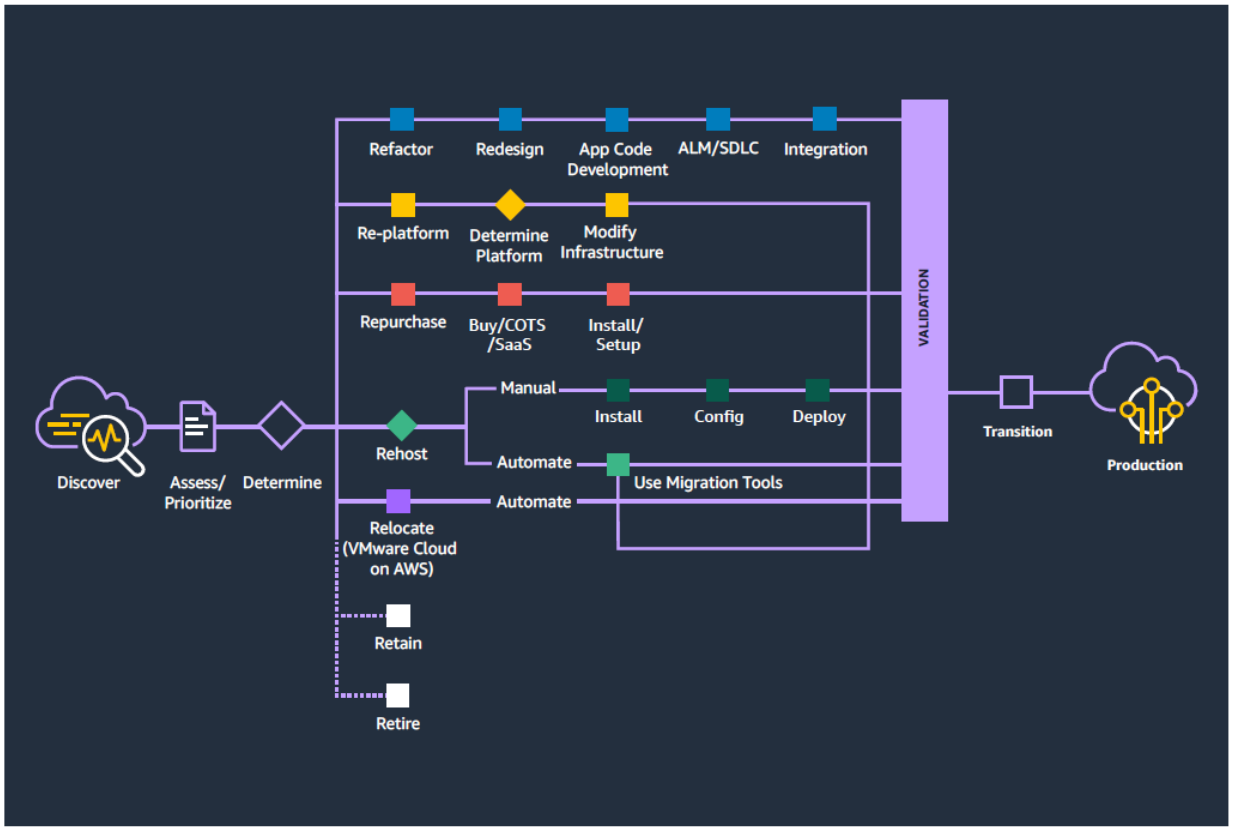

Figure 8 – Seven Common Migration Patterns

The seven migration patterns include:

1. **Rehost** (i.e., “lift and shift”) involves moving applications without changing them.

The AWS Application Migration Service (AWS MGN) enables you to rehost a large number of physical, virtual, or cloud servers to AWS without compatibility issues, performance disruption, or long cutover windows. Features such as automatic server conversion and continuous replication, combined with non-disruptive tests, ensure a smooth cutover for your most critical databases and applications, such as SAP CRM, Oracle E-Business Suite, and Microsoft SharePoint. Using AWS MGN, you can rehost applications from VMware vSphere, Microsoft Hyper-V, Amazon Elastic Compute Cloud (Amazon EC2), Amazon Virtual Private Cloud (Amazon VPC), and other clouds to AWS.

If your preferred AWS Region is not currently supported by AWS MGN, or, if the operating system on which your applications run is not currently supported by AWS MGN, you can automate rehosting using CloudEndure Migration, also designed for enterprise workloads such as SQL Server, Oracle, and SAP.

If you cannot or do not want to install an agent on your servers, you can use the AWS Server Migration Service (AWS SMS). AWS SMS offers agentless capabilities to migrate thousands of on-premises workloads to AWS. AWS SMS utilizes incremental, snapshot-based replication of the existing servers and enables cutover windows measured in hours.

2. **Replatform** (i.e., “lift, tinker, and shift”) involves making a few cloud optimizations to achieve a tangible benefit, but without changing the core architecture of the application. For example, our fully managed Amazon MQ service can easily replace a messaging broker without rewriting your applications or paying for third-party software licenses. The Amazon FSx for Windows File Server can help you migrate a Windows-based application that requires file storage.
3. **Refactor** (i.e., re-architect) enables you to re-imagine how the application is architected and developed by using cloud-native features. Refactoring is driven by a strong business need to add features, scale, or improve performance that would otherwise be difficult to achieve in the application’s existing environment. Many enterprises use the migration effort to modernize their businesses and refactor their legacy technology portfolio.
4. **Relocate** involves moving VMware vSphere®-based applications to AWS without application changes. Common migration patterns usually follow one of the other seven basic patterns, but when you migrate to AWS, you gain this option. For details, see VMware Cloud on AWS.
5. **Repurchase** (i.e., “drop and shop”) involves replacing your current environment with either a newer version of software or an entirely new solution. The AWS Marketplace offers a curated digital catalog of software solutions that support each phase of migration.
6. **Retain** involves retaining portions of your IT portfolio that you are not ready to migrate or believe are best kept on-premises. For applications that remain on-premises, AWS Outposts bring the same hardware, software, services, application programming interface (APIs), management tools, support, and operating model that is used in the AWS Cloud to your data center, co-location space, or on-premises facility.
7. **Retire** involves decommissioning or archiving portions of your IT portfolio that are no longer useful.

For more information, refer to the AWS eBook: 8 Migration Business Drivers and migration-related whitepapers available on our website.

188.7. Customer Responsibilities

Security and compliance responsibilities are shared between AWS and the customer. This shared responsibility model can help relieve customers' operational burdens as AWS operates, manages, and controls the components from the host operating system and virtualization layer down to the physical security of the facilities in which the service operates.

Customers—and in some cases, our AWS Partner Network (APN) Partners who work with those customers—control how they architect and secure their applications and data in the AWS Cloud. AWS provides a wide array of security and compliance services; a customer's responsibilities will vary depending on the services used, the integration of those services into their IT environment, and applicable laws and regulations.

As shown in **Figure 9** below, this differentiation of responsibility is commonly referred to as security "in" the cloud versus security "of" the cloud.

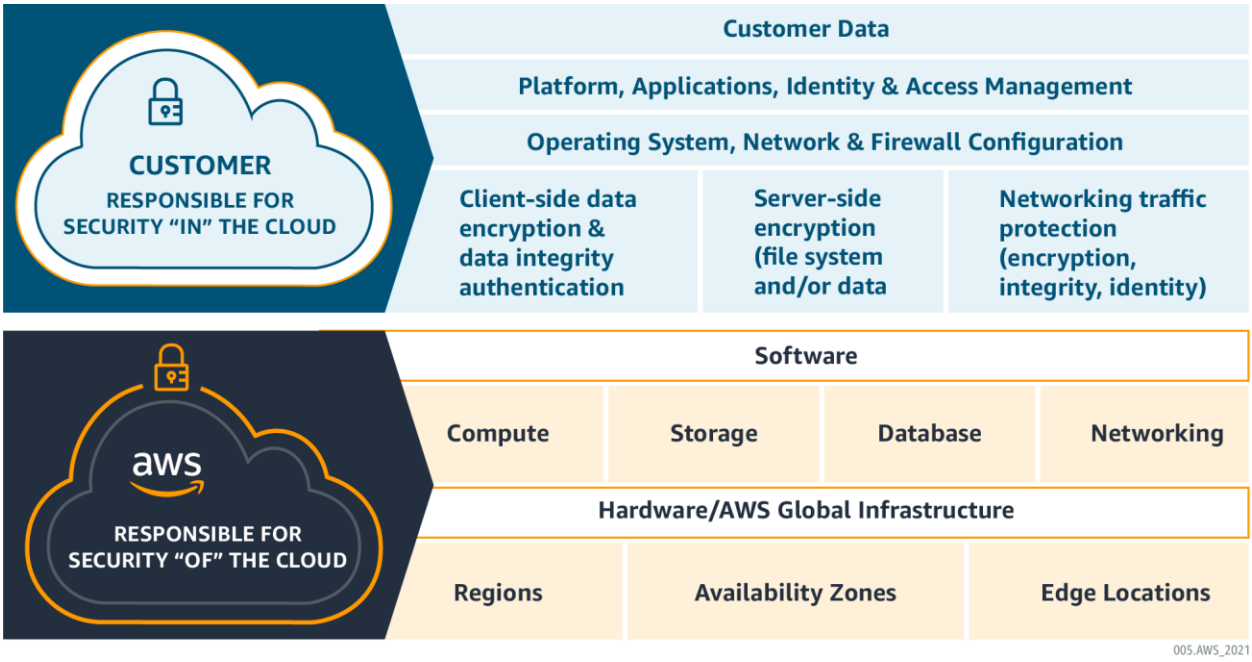

Figure 9. The Shared Responsibility Model

188.7.1. AWS and Customer Responsibilities

As described above, under the shared responsibility model, security and compliance responsibilities are shared between AWS and the customer.

188.7.1.1. AWS Responsibilities (Security of the Cloud)

AWS is responsible for protecting the infrastructure that runs all of the services offered in the AWS Cloud. This infrastructure is composed of the hardware, software, networking, and facilities that run AWS Cloud services. AWS operates, manages, and controls the infrastructure components that customers build upon.

We are vigilant about our customers' security and have implemented sophisticated technical and physical measures against unauthorized access. Customers can validate the security controls in place within the AWS environment through our certifications and reports, including the AWS System & 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.

Our ISO 27018 certification demonstrates that AWS has a system of controls in place that specifically address the privacy protection of customer content. These reports and certifications are produced by independent third-party auditors and attest to the design and operating effectiveness of AWS security controls. More information on AWS compliance certifications, reports, and alignment with best practices and standards can be found at http://aws.amazon.com/compliance/.

188.7.1.2. Customer Responsibilities (Security in the Cloud)

Customer responsibility is determined by the AWS Cloud services that they select. This determines the amount of configuration work the customer must perform as part of their security responsibilities. This determines the amount of configuration work the customer must perform as part of their security responsibilities.

For example, a service such as Amazon Elastic Compute Cloud (Amazon EC2) requires the customer to perform all of the necessary security configuration and management tasks. Customers that deploy an Amazon EC2 instance are responsible for management of the guest operating system (including updates and security patches), any application software or utilities installed on the instance, and the configuration of the AWS-provided firewall (called a *security group*) on each instance.

For abstracted services, such as Amazon Simple Storage Service (Amazon S3) and Amazon DynamoDB, AWS operates the infrastructure layer, the operating system, and platforms, and customers access the endpoints to store and retrieve data. Customers are responsible for managing their data (including encryption options), classifying their assets, and using AWS Identity and Access Management (IAM) tools to apply the appropriate permissions.

Customers are responsible for managing their data (including encryption options), classifying their assets, and using IAM tools to apply the appropriate permissions.

188.7.2. Shared Responsibility and IT Controls

The AWS shared responsibility model also extends to IT controls. Just as the responsibility to operate the IT environment is shared between AWS and its customers, so is the management, operation, and verification of IT controls. AWS can help relieve the customer burden of operating controls by managing those controls associated with the physical infrastructure deployed in the AWS environment (**inherited controls**).

Some controls apply to both the infrastructure layer (AWS responsibility) and customer layers (customer responsibility), but in completely separate contexts or perspectives (**shared controls**). With shared controls, AWS provides the requirements for the infrastructure, and the customer must provide their own control implementation within their use of AWS Cloud 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.

Some controls are the sole responsibility of the customer based on the application(s) they are deploying 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.

188.7.3. Shared Responsibility and Customer Data

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

188.7.3.1. Customer Content

Customers maintain ownership of their content, and they select which AWS Cloud services can process, store, and host it. We do not access or use customer content for any purpose without a customer's consent. We never use customer content or derive information from it for marketing or advertising.

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 Cloud services in connection with that customer's account; and any computational results that a customer or any end user derives from the foregoing, through their use of AWS Cloud services. For example, customer content includes content that a customer or any end user stores in Amazon S3. Customer content does not include account information, which we describe below. The terms of the AWS Customer Agreement and the AWS Service Terms apply to customer content.

There are five important basics concepts regarding customer content in the shared responsibility model:

1. Customers continue to own their content.
2. Customers choose the geographic location(s) in which to store their content—it does not move unless the customer decides to move it.
3. Customers can download or delete their content whenever they like.
4. Customers can "crypto-delete" their content by deleting the master encryption keys that are required to decrypt the data keys, which are, in turn, required to decrypt the data.
5. Customers should consider the sensitivity of their content and decide if and how to use various techniques such as encryption, tokenization, data decomposition, and cyber deception to protect their content.

Legal Requests for Customer Content

We are vigilant about the privacy of our customers. We do not disclose customer content unless we're required to do so to comply with the law, or with a valid and binding order of a governmental or regulatory body. 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, or if there is clear indication of illegal conduct in connection with the use of Amazon products or services, Amazon notifies customers before disclosing customer content so they can seek protection from disclosure. It is also important to point out

that our customers can encrypt their customer content, and we provide customers with the option to manage their own encryption keys.

We know that transparency matters to our customers, so we regularly publish a report about the types and volume of information requests we receive. The report is available on the Law Enforcement Information Requests webpage.

188.7.3.2. 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.

188.7.3.3. 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. Controls in place limit access to systems and data, and provide access to systems or data that is restricted and monitored. Refer to the AWS SOC 1 audit report (available under AWS Nondisclosure Agreement [NDA]) for more information and validation of the control testing related to access permissions and data deletion. Refer to the AWS Payment Card Industry Data Security Standard (PCI DSS) Compliance Package (available under AWS NDA) for testing performed to confirm data deletion. Both the AWS SOC 1 audit report and the AWS PCI Compliance Package can be requested at http://aws.amazon.com/compliance/contact/.

188.7.4. Additional Resources

AWS Website:

- Overview of the shared responsibility model: https://aws.amazon.com/compliance/shared-responsibility-model/
- Detail on customer responsibilities relating to customer data: https://aws.amazon.com/compliance/data-privacy-faq/
- How the shared responsibility model interacts with the GDPR: https://aws.amazon.com/blogs/security/the-aws-shared-responsibility-model-and-gdpr/

AWS Whitepapers:

- Using AWS in the Context of Common Privacy & Data Protection Considerations (Detailed sections on security “of” and “in” the cloud)
- Risk and Compliance (Deep dive on risk and compliance in the context of the shared responsibility model)
- Overview of Security Processes (AWS security measures to deliver security “in” the cloud)
- Data Residency (How customers can manage and protect their data within the shared responsibility model)