G-Cloud 14 Amazon Web Services EMEA SARL, UK Branch (AWS) – Cloud Compute Infrastructure Services Service Definition Catalogue

April 2024

Table of Contents

1 AWS Security Assurance 1
2 Amazon API Gateway 2
3 Amazon AppFabric 4
4 Amazon AppFlow 5
5 Amazon AppStream 2.0 8
6 Amazon Athena 12
7 Amazon Aurora 14
8 Amazon Bedrock 16
9 Amazon Braket 18
10 Amazon Chime (Lot 2) 19
11 Amazon CloudFront 21
12 Amazon CloudWatch 23
13 Amazon CodeCatalyst 26
14 Amazon CodeGuru 28
15 Amazon Cognito 30
16 Amazon Comprehend 32
17 Amazon Comprehend Medical 34
18 Amazon Connect (Lot 2) 36
19 Amazon Data Firehose 40
20 Amazon DataZone 43
21 Amazon Detective 44
22 Amazon DevOps Guru 46
23 Amazon DocumentDB (with MongoDB compatibility) 48
24 Amazon DynamoDB 53
25 Amazon ECS Anywhere 56
26 Amazon EKS Anywhere 58
27 Amazon Elastic Block Store (EBS) 59
28 Amazon Elastic Compute Cloud (EC2) 62
29 Amazon Elastic Container Registry (ECR) 66
30 Amazon Elastic Container Service (ECS) 68
31 Amazon Elastic File System (EFS) 70
32 Amazon Elastic Kubernetes Service (EKS) 73
33 Amazon Elastic MapReduce (EMR) 78
34 Amazon ElastiCache 80
35 Amazon EventBridge 85

36 Amazon File Cache 88
37 Amazon FinSpace 89
38 Amazon Fraud Detector 91
39 Amazon FSx for Lustre 92
40 Amazon FSx for NetApp ONTAP 94
41 Amazon FSx for OpenZFS 96
42 Amazon FSx for Windows File Server 98
43 Amazon GuardDuty 102
44 Amazon Inspector 105
45 Amazon IVS 108
46 Amazon Kendra 111
47 Amazon Keyspaces (for Apache Cassandra) 113
48 Amazon Kinesis Data Streams 116
49 Amazon Lex 118
50 Amazon Lightsail 121
51 Amazon Location Service 123
52 Amazon Lookout for Equipment 126
53 Amazon Lookout for Metrics 127
54 Amazon Lookout for Vision 129
55 Amazon Macie 131
56 Amazon Managed Blockchain 133
57 Amazon Managed Service for Apache Flink 135
58 Amazon Managed Service for Grafana 137
59 Amazon Managed Service for Prometheus 139
60 Amazon Managed Streaming for Apache Kafka 142
61 Amazon Managed Workflows for Apache Airflow 146
62 Amazon MemoryDB for Redis 147
63 Amazon Monitron 150
64 Amazon MQ 152
65 Amazon Neptune 154
66 Amazon OpenSearch Service 157
67 Amazon Personalize 159
68 Amazon Pinpoint 162
69 Amazon Polly 164
70 Amazon Quantum Ledger Database 167
71 Amazon QuickSight 169
72 Amazon Redshift 172

73 Amazon Rekognition 175
74 Amazon Relational Database Service (RDS) 177
75 Amazon Route 53 181
76 Amazon SageMaker 183
77 Amazon Security Lake 187
78 Amazon Simple Email Service (SES) 188
79 Amazon Simple Notification Service (SNS) 191
80 Amazon Simple Queue Service (SQS) 194
81 Amazon Simple Storage Service (S3) 196
82 Amazon Simple Workflow Service (SWF) 199
83 Amazon Textract 201
84 Amazon Timestream 203
85 Amazon Transcribe 205
86 Amazon Transcribe Medical 208
87 Amazon Translate 209
88 Amazon Verified Permissions 211
89 Amazon Virtual Private Cloud (VPC) 213
90 Amazon WorkMail (Lot 2) 215
91 Amazon WorkSpaces (Lot 2) 218
92 AWS Amplify 219
93 AWS App Runner 223
94 AWS Application Discovery Service 225
95 AWS Application Migration Service 226
96 AWS AppSync 228
97 AWS Artifact 230
98 AWS Audit Manager 231
99 AWS Backup 233
100 AWS Batch 236
101 AWS Budgets 239
102 AWS Certificate Manager 240
103 AWS Chatbot 242
104 AWS Cloud9 243
105 AWS Cloud Map 246
106 AWS CloudFormation 247
107 AWS CloudHSM 249
108 AWS CloudTrail 251
109 AWS CodeArtifact 253

110 AWS CodeBuild 254
111 AWS CodePipeline 256
112 AWS Compute Optimizer 258
113 AWS Config 260
114 AWS Control Tower 262
115 AWS Cost Explorer 264
116 AWS Data Exchange (BYOS) 266
117 AWS Database Migration Service 267
118 AWS DataSync 270
119 AWS DeepRacer 273
120 AWS Device Farm 274
121 AWS Digital Investigation and Forensics Storage (created for G-Cloud) 276
122 AWS Direct Connect 278
123 AWS Directory Service 279
124 AWS Elastic Beanstalk 282
125 AWS Elastic Disaster Recovery 284
126 AWS Elemental MediaConnect 286
127 AWS Elemental MediaConvert 287
128 AWS Elemental MediaLive 289
129 AWS Elemental MediaPackage 290
130 AWS Elemental MediaStore 292
131 AWS Elemental MediaTailor 293
132 AWS Fargate 295
133 AWS Fault Injection Service 297
134 AWS Firewall Manager 299
135 AWS Global Accelerator 300
136 AWS Glue 303
137 AWS Ground Station 306
138 AWS Health Dashboard 308
139 AWS HealthImaging 310
140 AWS Healthlake 312
141 AWS HealthOmics 314
142 AWS Identity and Access Management (IAM) 315
143 AWS IoT Analytics 317
144 AWS IoT Core 320
145 AWS IoT Device Defender 322
146 AWS IoT Device Management 324

147 AWS IoT Events....326
148 AWS IoT FleetWise....328
149 AWS IoT Greengrass....330
150 AWS IoT SiteWise....333
151 AWS IoT TwinMaker....335
152 AWS Key Management Service....336
153 AWS Lake Formation....339
154 AWS Lambda....341
155 AWS License Manager....344
156 AWS Mainframe Modernization....347
157 AWS Marketplace – BYOL....350
158 AWS Migration Hub....352
159 AWS Network Firewall....354
160 AWS Organizations....356
161 AWS Outposts....358
162 AWS Panorama....362
163 AWS PrivateLink....363
164 AWS Resilience Hub....365
165 AWS Resource Access Manager (RAM)....366
166 AWS Secrets Manager....368
167 AWS Security Hub....370
168 AWS Service Catalog....373
169 AWS Shield....376
170 AWS SimSpace Weaver....379
171 AWS Snowball Edge....381
172 AWS Step Functions....384
173 AWS Storage Gateway....387
174 AWS Supply Chain....390
175 AWS Systems Manager....392
176 AWS Transfer Family....396
177 AWS Transit Gateway....398
178 AWS Trusted Advisor....400
179 AWS Verified Access....401
180 AWS VPN....403
181 AWS WAF....406
182 AWS Wavelength....407
183 AWS Well-Architected Tool....409

184 AWS X-Ray....411
185 Elastic Load Balancing....413
186 FreeRTOS....414
187 Cross-Service Definitions....417

1 AWS Security Assurance

Moving IT infrastructure to the AWS Cloud 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/.

1.1 Information Assurance

The following subsections provide information relating to information assurance.

1.1.1 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. 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/.

1.1.2 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.

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

2 Amazon API Gateway

2.1 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.

2.1.1 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.
- **Resiliency:** API Gateway helps you manage traffic to your backend systems. You can set throttling rules based on the number of requests per second for each HTTP method in your APIs. API Gateway handles any level of traffic received by an API, so you focus on your business logic and services rather than maintaining infrastructure.
- **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.

Use or disclosure of data contained on this sheet is subject to the restriction on the title page of this document.

- **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.

2.1.2 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.

2.2 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.

2.3 Pricing Overview

Please see the AWS UK G Cloud 14 Pricing Document accompanying this service in the Digital Marketplace.

2.4 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/

2.5 Technical Requirements

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

3 Amazon AppFabric

3.1 Service Overview

AppFabric quickly connects software-as-a-service (SaaS) applications across your organization. IT and security teams can then easily manage and secure applications using a standard schema, and employees can complete everyday tasks faster using generative artificial intelligence (AI).

3.1.1 Features

- **Supported SaaS applications**: AppFabric connects SaaS applications including Asana, the Atlassian Jira suite, Dropbox, Miro, Okta, Slack, ServiceNow, Smartsheet, Webex by Cisco, Zendesk, Zoom, GitHub, Google Workspace, and Microsoft 365, with more coming soon.
- **Supported security tools**: Application data from AppFabric is compatible with security tools like Logz.io, Netskope, NetWitness, Rapid7, and Splunk, or your proprietary security solution.
- **Normalized audit logs**: SaaS application audit logs (for example, event type, workspace name, actor identifier, or device details) are automatically normalized into the Open Cybersecurity Schema Framework (OCSF), or raw data is made available in two data formats—JSON or Apache Parquet.
- **Enriched audit logs**: AppFabric automatically enriches each application's audit log data with a user email address when applicable (when user identity [UID] is matched to the employee email address). By enriching each audit log event with a user identifier (email address), security teams can reduce security incident response time.
- **User access**: Security and IT admin teams can use AppFabric to see who has access to what application(s) by running a simple search using the employee's corporate email address. AppFabric will then verify application access or deprovisioning status across all connected applications.
- **OCSF standard to SaaS applications**: AppFabric closely works with the open-source OCSF community to create new event categories and classes. Together, we have introduced Application Activity as an OCSF event category and introduced activities (like export) specific to SaaS applications in the OCSF.
- **Productivity for application developers**: AWS AppFabric productivity feature reimagines end-user productivity in SaaS applications by generating insights and actions with context from multiple applications including Asana, Atlassian Jira suite, Miro, Slack, Smartsheet, Microsoft 365, and Google Workspace. AppFabric helps application developers provide generative AI-driven insights in the end user's preferred SaaS application user interface (UI), creating new or enhancing existing generative AI assistants across multiple applications.

3.1.2 Benefits

- **Quickly connect SaaS applications**: Connect your applications in minutes and reduce operational costs.
- **Enhance security observability**: Increase observability across SaaS application data to elevate your security posture.
- **Improve employee productivity**: Automatically facilitate tasks across applications with generative AI.

3.2 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/.

3.3 Pricing Overview

Please see the AWS UK G Cloud 14 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/appfabric/
- **Service quotas:** https://docs.aws.amazon.com/appfabric/latest/adminguide/limits.html
- **Service FAQs:** https://aws.amazon.com/appfabric/faqs/

3.5 Technical Requirements

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

- Administration Guide: Describes key concepts of AppFabric and provides instructions for configuring the service.
- API Reference: Describes the API operations for AppFabric. It also provides the basic structure and elements of an API request and response.

4 Amazon AppFlow

4.1 Service Overview

Amazon AppFlow is a fully managed integration service that securely transfers data between Software-as-a-Service (SaaS) applications like Salesforce, Marketo, Slack, and ServiceNow, and AWS services like Amazon S3 and Amazon Redshift.

With AppFlow, you can run data flows at nearly any scale at the frequency you choose - on a schedule, in response to a business event, or on demand. You can configure powerful 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

- **Intuitive user interface**: You can use Amazon AppFlow to set up data flows in minutes with no coding required. Using an intuitive user interface, you can select your data sources and destinations; configure optional transformations, filters, and validations; and run your flow without having to rely on technical teams.
- **Flexible data flow triggers**: Amazon AppFlow helps you run data flows on demand to do bulk transfers or set up a routine schedule to keep data in sync. It can also run flows in response to business events such as the creation of a sales opportunity, the status change of a support ticket, or the completion of a registration form.
- **Native SaaS integrations**: Amazon AppFlow includes native integration with software as a service (SaaS) applications used daily for business operations including Salesforce, SAP, Google Analytics, and Facebook Ads, with additional integrations planned. With Amazon AppFlow, you can transfer data from any supported SaaS application in a few selections.
- **Enterprise grade data transformations**: You can use Amazon AppFlow to map source and destination fields through bulk or individual field mapping, or by using a CSV file. Amazon AppFlow helps you perform data transformations such as merging, masking, filtering, and validation. For example, you can validate your data is in the right numerical format, merge first and last names, or mask credit card details.
- **Partitioning and aggregation**: Amazon AppFlow partitions your data by source object schema when Amazon Simple Storage Service (S3) is selected as a destination. Partitioning enhances the layout of your data to encourage efficient querying by analytic engines such as Amazon Athena and Amazon Redshift Spectrum.
- **Glue Data Catalog integration**: Amazon AppFlow automates the preparation and registration of your data into the AWS Glue Data Catalog in a few selections. Share, discover, and access data from Amazon AppFlow's growing library of SaaS connectors with AWS analytics and machine learning services such as AWS Glue, Amazon Athena, and Amazon SageMaker Data Wrangler.
- **Build your own connector**: The Amazon AppFlow Connector Software Development Kit (SDK) allows you to build your own connectors to securely transfer data between your custom endpoint, application, or other cloud service to and from Amazon AppFlow's library of managed SaaS and AWS connectors. Develop and test connectors with the same open source SDK used by AWS teams internally.
- **Pay as you go**: Amazon AppFlow offers a cost-savings advantage compared to building connectors in-house or through other application integration services. There are no upfront charges or licensing fees to use Amazon AppFlow, and you pay for only the number of flows you run and the volume of data processed.
- **High-scale data transfer**: Amazon AppFlow can run up to 100 GB of data per flow, which helps you more easily transfer millions of Salesforce records, Marketo leads, or Zendesk tickets.
- **Data privacy defaults through PrivateLink**: AWS PrivateLink simplifies the security of data shared with cloud-based applications by decreasing the exposure of data to the internet. For SaaS applications that have PrivateLink enabled, Amazon 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.
- **IAM policy enforcement**: Use your existing AWS Identity and Access Management (IAM) policies to enforce fine-grained permissions and keep access consistent across your organization, without creating new policies. With strictly enforced IAM policies,

application administrators can safely create and manage data flows without depending on technical teams.

- **Built in reliability**: Amazon AppFlow is built with a highly available architecture to prevent single points of failure. Amazon AppFlow takes advantage of AWS scaling, monitoring, auditing, and billing features so you don't need to configure these yourself.

4.1.2 Benefits

- **Fully managed**: Fully managed integration service to transfer data between services such as Salesforce, SAP, Google Analytics, and Amazon Redshift.
- **Speed and agility**: You can integrate applications in a few minutes, without 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 Amazon Appflow, data flow quality is built in, and you can enrich the flow of data through masking, mapping, merging, filtering, and validation as part of the flow itself.
- **Privacy and security**: AppFlow encrypts data at rest and in motion. You can encrypt data with AWS managed keys or bring your own custom keys. It also allows users to restrict data from flowing over the public Internet using Amazon VPC endpoints enabled by AWS PrivateLink. This minimizes the threat from Internet based attacks and the risk of sensitive data leakage.
- **Scalability**: 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. Using Amazon AppFlow, you can easily transfer millions of Salesforce records or Zendesk tickets - all while running a single flow.
- **Reliability**: AppFlow uses a highly available architecture with redundant, isolated resources to prevent any single points of failure while running within the resilient AWS infrastructure.

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 14 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/
- **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/ and the following links for comprehensive technical documentation regarding this service.

- User Guide: Describes key concepts and provides instructions for using Amazon AppFlow.
- API Reference: Describes all the API operations for Amazon AppFlow in detail. Also provides sample requests, responses, and errors for the supported web service protocols.
- Amazon AppFlow in the AWS CLI: Describes the Amazon AppFlow operations that are available in the AWS Command Line Interface (AWS CLI).
- Amazon AppFlow in the AWS CloudFormation User Guide: Describes all Amazon AppFlow resource and property types that are supported by AWS CloudFormation.
- Amazon AppFlow in the AWS SDK for Python (Boto3): Describes all of the Amazon AppFlow operations that are included in the AWS SDK for Python (Boto3).
- Amazon AppFlow in the AWS SDK for Java 2.x: Describes all of the Amazon AppFlow operations that are included the AWS SDK for Java 2.x.
- 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 is an AWS End User Computing (EUC) service that can be configured for SaaS application streaming or delivery of virtual desktops with selective persistence. When AppStream 2.0 is used for virtual desktops, saved files and application settings remain persistent between user sessions, and a fresh virtual desktop is assigned to the user every time they log on. AppStream 2.0 works with Amazon WorkSpaces Thin Client for cost-effective, easy-to-manage, secure access to virtual desktops that improves end-user and IT staff productivity.

AppStream 2.0 offers multi-session capability, allowing you to provision multiple user sessions on a single AppStream 2.0 instance. This gives you the ability to optimize resource utilization and helps lower cost, supporting a wide range of user types without over provisioning.

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. When your testing is complete, you can set application launch configurations, configure default user settings, and publish your image for users to access.
- **Multiple instance types**: Amazon AppStream 2.0 runs your applications on streaming instances, which provide the GPU, CPU, memory, storage, and networking capacity you need. With AppStream 2.0, you can choose from a selection of them to optimize the performance of your applications, and reduce your overall costs. You

can use General Purpose, Compute Optimized, Memory Optimized, Graphics Design, Graphics Pro, and Graphics G4 families based on your use case.

- **GPU optimized instances**: Amazon AppStream 2.0 offers Graphics Design, Graphics Pro and Graphics G4 and Graphics G5 instance families. Graphics Design instances are ideal for delivering applications such as Adobe Premiere Pro, Autodesk Revit, and Siemens NX that rely on hardware acceleration of DirectX, OpenGL, or OpenCL. Powered by AMD FirePro S7150x2 Server GPUs and equipped with AMD Multiuser GPU technology, instances start from 2 vCPU, 7.5 GiB system memory, and 1 GiB graphics memory, to 16 vCPUs, 61 GiB system memory, and 8 GiB graphics memory.
- **Amazon VPC support**: With Amazon AppStream 2.0, your desktop applications can be launched inside an Amazon Virtual Private Cloud (VPC). 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
- Secure access: 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. This allows you to apply your existing AD policies to your streaming instances, and provides your users with single sign on access to Intranet sites, file shares, and network printers from within their applications.
- **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.
- **Fleets**: Amazon AppStream 2.0 offers a variety of fleet options, allowing you to configure the optimal solution for your specific requirements.
- **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. AppStream 2.0 includes applications manager, designed to help you package and mobilize Windows applications across environments, without affecting the underlying base image or operating system. After launching an app block builder instance, install your application(s) and create a self-contained and sharable application image. You can share this application image across multiple AppStream 2.0 elastic fleets provisioned across different AWS Regions
- **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. You can also create your own scaling policies for each of your Always-On and On-Demand fleets. With Elastic fleets, AppStream 2.0 manages the size of the fleet for you without needing to use autoscaling policies.
- **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. You can also track utilization over time, which helps you ensure that your fleet is the correct size. Using CloudWatch, you can set alarms to notify you when there is insufficient

capacity to support your users, and to react to changes in your Always-On and On-Demand fleet.

- **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. You can also access your desktop applications using the AppStream 2.0 Windows Client to use dual monitors and USB peripherals such as 3D mice. The Windows Client also supports keyboard shortcuts, such as Alt + Tab, clipboard shortcuts, and function keys. The Windows Client and Chromium-based web browsers support real-time audio-video (AV) when connecting to AppStream 2.0 sessions.
- **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. Each of these are accessed from the my files tab within an active AppStream 2.0 streaming session, and content can be saved or opened directly from the File menu in most apps.
- **Simple user interface**: Amazon AppStream 2.0 offers an intuitive user interface, making it easy to control your experience. You can use the application-centric view, or use the standard desktop experience. You can run multiple applications and easily switch and interact with them. You can copy and paste between applications and your local machine, quickly upload or download files, and print to your local or network printer. You can also listen to audio output from applications, and control the bandwidth and fidelity for the experience.
- **NICE DCV protocol**: Amazon AppStream 2.0 uses NICE DCV technology to provide secure, high-performance access to your applications. NICE DCV delivers a fluid interactive experience for business, design, engineering, and software development applications. It delivers a secure pixel stream to your end users and always adjusts to network conditions to provide an optimal end-user experience. NICE DCV secures both pixels and end-user inputs using end-to-end AES-256 encryption.
- **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

- **Access applications from anywhere**: Empower remote workers and react quickly to changing conditions with secure access to applications and desktops from anywhere.
- **No refactoring required**: Stream Software-as-a-Service (SaaS) applications or quickly convert desktop applications to SaaS without refactoring.

- **Range of options**: Select from a range of compute, memory, and storage options, while auto scaling helps avoid overprovisioning.
- **Fully managed**: Reduce downtime with fully managed application streaming and virtual desktop services, including a financially-backed 99.9% uptime SLA.

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:

- Amazon AppStream 2.0 Administration Guide: https://docs.aws.amazon.com/appstream2/latest/developerguide/home-folders.html. Users schedule and recover backups through a web interface.
- Best Practices for Deploying Amazon AppStream 2.0: https://docs.aws.amazon.com/whitepapers/latest/best-practices-for-deploying-amazon-appstream-2/disaster-recovery.html

5.3 Pricing Overview

Please see the AWS UK G Cloud 14 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 AppStream 2.0 and includes detailed instructions for using the features available in this service.
- API Reference: Describes all the API operations for AppStream 2.0 in detail. Also provides sample requests, responses, and errors for the supported web services protocols.
- AppStream 2.0 section in the AWS CLI Reference: Describes the AppStream 2.0 commands that are available in the AWS Command Line Interface (AWS CLI).
- Extension SDK Developer Guide: Describes how to get started with the NICE DCV Extension SDK that is supported by AppStream 2.0 and provides an SDK reference.

6 Amazon Athena

6.1 Service Overview

Amazon Athena is a serverless, interactive analytics service built on open-source frameworks, supporting open-table and file formats. Athena provides a simplified, flexible way to analyze petabytes of data where it lives. Analyze data or build applications from an Amazon Simple Storage Service (S3) data lake and 30 data sources, including on-premises data sources or other cloud systems using SQL or Python. Athena is built on open-source Trino and Presto engines and Apache Spark frameworks, with no provisioning or configuration effort required.

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 is based on Trino and Presto, open source, distributed SQL engines optimized for low latency, interactive data analysis. 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.
- **Flexible pricing**. Amazon Athena offers two flexible pricing models. By default, queries are billed based on the data scanned per query in terabytes (TB). This allows you to submit queries without planning ahead for compute. If you prefer to pay based on the compute your queries consume or want to control concurrency and prioritize workloads, use capacity-based pricing available with Provisioned Capacity. For added flexibility, you can use per query billing and capacity-based pricing at the same time in the same account.
- **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.

- **Federated query**. Athena provides built-in connectors to 30 popular AWS, on-premises, and other cloud data stores, including Amazon Redshift, Amazon DynamoDB, Google BigQuery, Google Cloud Storage, Azure Synapse, Azure Data Lake Storage, Redis, Snowflake, and SAP Hana. By using Athena data source connectors, you can generate insights from multiple data sources using the Athena SQL syntax and without the need to move or transform your data.
- **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 simple for anyone with SQL experience to run ML models deployed on Amazon SageMaker.
- **Supports for Apache Spark framework**. Apache Spark is an open-source, distributed processing system used for fast analytics workloads against data of any size that offers a rich system of open-source libraries. With Athena, you can build Spark applications in expressive languages, such as Python, using a simplified notebook experience in the Athena console or through Athena APIs. You can query data from various sources, chain together multiple calculations, and visualize the results of their analyses. For interactive Spark applications, you spend less time waiting and are more productive as Athena starts running applications under a second.

6.1.2 Benefits

- **Start querying instantly:** Athena is serverless. With Athena, you can use SQL or Apache Spark and there is no infrastructure to set up or manage. 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 based on data processed or compute used.
  - **SQL queries**: With per query billing, you can get started quickly and pay only for the data scanned by queries you run. You are charged for the number of bytes scanned per query, rounded up to the nearest megabyte, with a 10 MB minimum per query. There are no charges for Data Definition Language (DDL) statements like CREATE, ALTER, or DROP TABLE statements for managing partitions, or failed queries. Canceled queries are billed for the total amount of data scanned when the query is canceled.
  - **SQL queries with Provisioned Capacity**: Use Provisioned Capacity when you want dedicated compute resources for your SQL queries or prefer a flat-rate, hourly pricing structure. With Provisioned Capacity, you purchase Data Processing Units (DPU), which represent the resources necessary to run queries, and assign them to your workloads. One DPU provides 4 vCPU and 16 GB of memory. You pay only for the capacity you need and the duration it is active in your account, with a 1 hour minimum. There are no charges for data scanned.
- **Open, powerful, standard:** Athena supports ANSI SQL queries. Athena uses Trino, an open-source, in-memory, distributed SQL engine, and can handle complex analysis, including large joins, window functions, and arrays. Athena supports various data formats like CSV, TSV, JSON, or Textfiles and also supports open-source columnar formats, such as ORC 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.

6.1 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/.

6.2 Pricing Overview

Please see the AWS UK G Cloud 14 Pricing Document accompanying this service in the Digital Marketplace.

6.3 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/

6.4 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 Aurora

7.1 Service Overview

Amazon Aurora provides high performance and availability at a global scale with full MySQL and PostgreSQL compatibility, at 1/10th the cost of commercial databases. Aurora has 5x the throughput of MySQL and 3x of PostgreSQL. Aurora has broad compliance standards and best-in-class security capabilities. Aurora offers storage resiliency by making data durable across 3 AZs (customers only pay for 1 copy). Aurora has up to 99.99% availability, and customers can access local read performance with Global Database when deployed across AWS Regions.

With serverless, an on-demand, autoscaling configuration for Amazon Aurora, Aurora can scale-up to hundreds of thousands of transactions in a fraction of a second. Aurora zero-ETL integration with Amazon Redshift delivers analytics on transactional data in near real-time.

7.1.1 Features

- **High performance and scalability**. Aurora offers up to 5x the throughput of MySQL and 3x the throughput of PostgreSQL on similar hardware. Aurora uses a variety of software and hardware techniques to ensure the database engine is able to fully use available compute, memory, and networking. I/O operations use distributed systems techniques, such as quorums to improve performance consistency.

- **Highly secure**. Aurora offers network isolation where it runs in Amazon Virtual Private Cloud (Amazon VPC), and helps you isolate your database in your own virtual network and connect to your on-premises IT infrastructure using industry-standard encrypted IPsec VPNs.
- **Generative AI**: Aurora offers capabilities to enable machine learning (ML) and generative artificial intelligence (AI) models to work with data stored in Aurora in real-time and without moving the data. Aurora machine learning also simplifies adding generative AI model predictions and embeddings to your Aurora database.
- **Zero-ETL integrations**. Aurora zero-ETL integration with Amazon Redshift: Allows near real-time analytics and ML using Amazon Redshift on petabytes of transactional data from Aurora. Transactional data is automatically and continuously replicated within seconds of being written in Aurora and is seamlessly made available in Amazon Redshift. Once data is available in Amazon Redshift, you can start analyzing it immediately and apply advanced features like data sharing, materialized views, and Amazon Redshift ML to get holistic and predictive insights.

7.1.2 Benefits

- **Cost effective**: There is no upfront commitment with Aurora. You pay an hourly charge for each instance that you launch, and when you're finished with an Aurora DB instance, you can delete it. You do not need to overprovision storage as a safety margin, and you only pay for the storage you actually consume.
- **Cost-effective**: Improve productivity and lower total cost of ownership with a fully managed database including innovations like serverless so you can focus on building applications that delight your users.
- **Easy to migrate**: Easily migrate MySQL or PostgreSQL databases to and from Aurora using standard tools, or run legacy SQL Server applications with Babelfish for Aurora PostgreSQL with minimal code change.
- **High availability and durability**. Amazon RDS continuously monitors the health of your Aurora database and underlying Amazon Elastic Compute Cloud (Amazon EC2) instance. In the event of database failure, Amazon RDS will automatically restart the database and associated processes. The backup capability of Aurora enables point-in-time recovery for your instance. This allows you to restore your database to any second during your retention period, up to the last 5 minutes. Your automatic backup retention period can be configured up to 35 days. Automated backups are stored in Amazon Simple Storage Service (Amazon S3), which is designed for very high durability. Aurora backups are automatic, incremental, and continuous and have no impact on database performance.

7.2 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.

7.3 Pricing Overview

Please see the AWS UK G Cloud 14 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/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/

7.5 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.

8 Amazon Bedrock

8.1 Service Overview

Amazon Bedrock is a fully managed service that offers a choice of high-performing foundation models (FMs) from leading AI companies like AI21 Labs, Anthropic, Cohere, Meta, Mistral AI, Stability AI, and Amazon via a single API. Amazon Bedrock offers a broad set of capabilities you need to build generative AI applications with security, privacy, and responsible AI. You can experiment and evaluate top FMs for your use case, privately customize them with your data using techniques such as fine-tuning and Retrieval Augmented Generation (RAG), and build agents that execute tasks using your enterprise systems and data sources. Since Amazon Bedrock is serverless, you don't have to manage any infrastructure, and you can securely integrate and deploy generative AI capabilities into your applications using the AWS services you are already familiar with.

8.1.1 Features

- **Agents**: Agents are fully managed capabilities of Amazon Bedrock. You can securely connect FMs to your company data sources using Agents. Agents can automatically break down tasks and create an orchestration plan–without any manual coding. The agent securely connects to company data through an API, automatically converting data into a machine-readable format, and augmenting the request with relevant information to generate the most accurate response.
- **Developer experience**: Amazon Bedrock makes building with a range of FMs as easy as an API call. You can easily experiment with different FMs using interactive playgrounds for various modalities including text, chat, and image. Amazon Bedrock provides access to leading models including AI21 Labs' Jurassic, Anthropic's Claude, Cohere's Command and Embed, Meta's Llama 2, and Stability AI's Stable Diffusion, as well as our own Amazon Titan models. Amazon Bedrock creates a private, customized copy of the base FM for you, and your data is not used to train the original base models. Your data used to customize models is securely transferred through your Amazon Virtual Private Cloud (VPC).
- **Guardrails**: Guardrails for Amazon Bedrock evaluates user inputs and FM responses based on use case specific policies, and provides an additional layer of safeguards regardless of the underlying FM. Guardrails can be applied across FMs, including Anthropic Claude, Meta Llama 2, Cohere Command, AI21 Labs Jurassic, and Amazon Titan Text, as well as fine-tuned models.
- **Knowledge bases**: Allow you to give FMs and agents contextual information from your company's private data sources for Retrieval Augmented Generation (RAG) to deliver more relevant, accurate, and customized responses. You implement the entire RAG workflow from ingestion to retrieval and prompt augmentation without

having to build custom integrations to data sources and manage data flows. Session context management is built in, so your app can readily support multi-turn conversations.

8.1.2 Benefits

- **Security and Privacy**: With Amazon Bedrock, you have full control over the data you use to customize the foundation models for your generative AI applications. Your data is encrypted in transit and at rest. Additionally, you can create, manage, and control encryption keys using the AWS Key Management Service (AWS KMS). Identity-based policies provide further control over your data, helping you manage what actions users and roles can perform, on which resources, and under what conditions.
- **Comprehensive data protection**: Amazon Bedrock helps ensure that your data stays under your control. When you tune a foundation model, we base it on a private copy of that model. This means your data is not shared with model providers, and is not used to improve the base models. You can use AWS PrivateLink to establish private connectivity from your Amazon Virtual Private Cloud (VPC) to Amazon Bedrock, without having to expose your VPC to internet traffic. Finally, Bedrock is in scope for common compliance standards including ISO, SOC, CSA STAR Level 2, is HIPAA eligible, and customers can use Bedrock in compliance with the GDPR.
- **Responsible AI**: Guardrails for Amazon Bedrock enables you to implement safeguards for your generative AI applications based on your use cases and responsible AI policies. Guardrails helps control the interaction between users and Foundation Models (FMs) by filtering undesirable and harmful content and will soon redact personally identifiable information (PII), enhancing content safety and privacy in generative AI applications. You can create multiple guardrails with different configurations tailored to specific use cases. Additionally, you can continuously monitor and analyze user inputs and FM responses that may violate customer-defined policies in the guardrails.
- **Powerful tools to build at no extra charge**. When using Agents for Amazon Bedrock and Knowledge Bases for Amazon Bedrock, you are only charged for the models and the vector databases you use with these capabilities.

8.2 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/.

8.3 Pricing Overview

Please see the AWS UK G Cloud 14 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/bedrock/
- **Service quotas:** https://docs.aws.amazon.com/bedrock/latest/userguide/quotas.html
- **Service FAQs:** https://aws.amazon.com/bedrock/faqs/

8.5 Technical Requirements

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

- Amazon Bedrock User Guide: Describes all Amazon Bedrock concepts and provides instructions on using the various features with both the console and the API.
- API Reference: Describes the Amazon Bedrock API operations and parameters.

9 Amazon Braket

9.1 Service Overview

Amazon Braket is a fully managed service that helps you get started with quantum computing and accelerate your research. Using Amazon Braket, you can learn how to program quantum computers and explore potential applications. You can design your own quantum algorithms from scratch or choose from a set of pre-built algorithms.

9.1.1 Features

- **Simplified access to quantum computers**: Amazon Braket provides secure access to a variety of quantum computing technologies. Amazon Braket provides access to superconducting, trapped ion, and neutral-atom quantum computers. Following the gate-based quantum computing paradigm, you can access superconducting quantum processors from Oxford Quantum Circuits and Rigetti, and trapped-ion technology from IonQ.
- **Consistent experience with choice of simulators**: 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 on-demand simulators. You can run a circuit on Amazon Braket simulators with a single API call. 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.
- **Hardware-agnostic developer framework**. The Amazon Braket SDKis 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.
- **Braket Direct**: In addition to on-demand access to quantum computers, you can also get dedicated device access using reservations through Braket Direct. Reservations give you exclusive access to a quantum device of your choosing.
- **Management and security**: Amazon Braket is accessible through the AWS Management Console, a centralized web interface 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.

9.1.2 Benefits

- **Accelerate scientific discovery**: Easily work with different types of quantum computers and circuit simulators using a consistent set of development tools.

- **Trusted cloud**: Build quantum projects on a trusted cloud with simple pricing and management controls for both quantum and classical workloads.
- **Priority access**: Run hybrid quantum-classical algorithms faster with priority access to quantum computers and no classical infrastructure to manage.
- **Reserve dedicated capacity**: Reserve dedicated device access and engage directly with quantum computing specialists using Braket Direct.

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 14 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/
- **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/ and the following link for comprehensive technical documentation regarding this service.

- Amazon Braket 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 (Lot 2)

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, and Android devices.

10.1.1 Features

- **Online meetings**: Use your calendar and add meet@chime.aws to your meeting invite. Adding meet@chime.aws will enable Chime to call you so you can ensure your meetings start on time. With Amazon Chime call me, you simply enter your phone number in the web or mobile application and answer the incoming call to join a meeting.
- **Video conferencing**: You can use high-quality wideband audio and high-definition video conferencing for up to 25 people on your desktop or mobile devices. You can create and use personalized meeting URLs for your online meetings.
- **Team Collaboration**: You can use Amazon Chime to chat with colleagues directly, in a group, or in a chat room. You can create Chat Rooms to enable multiple users to collaborate on projects, share files with colleagues, and use @mentions to direct messages to specific participants.
- **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.
- **Security and Administration**: Amazon Chime is an AWS service, which means you benefit from a secure data center and network architecture. Amazon Chime features security capabilities built directly into the service. Messages, voice, video, and content are encrypted using AES 256-bit encryption. Using single sign on, you can integrate Chime with AWS Directory Services for Microsoft Active Directory or Okta SSO to authenticate and manage users.

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.

Use or disclosure of data contained on this sheet is subject to the restriction on the title page of this document.

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 14 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.
- API Reference: Describes the API operations for Amazon Chime.
- 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 content delivery network (CDN) service built for high performance, security, and developer convenience. Amazon CloudFront gives businesses and web application developers an easy and cost effective way to distribute content with low latency and high data transfer speeds.

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.
- **Continuous Deployment**: Continuous deployment with CloudFront gives you a high level of deployment safety. You can now deploy two separate but identical environments—blue and green, and enable simple integration into your continuous integration and delivery (CI/CD) pipelines with the ability to roll out releases gradually without any domain name system (DNS) changes.
- **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 600+ 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.

Use or disclosure of data contained on this sheet is subject to the restriction on the title page of this document.

- **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 14 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 CloudWatch

12.1 Service Overview

Amazon CloudWatch is an AWS monitoring service for cloud resources and the applications that you run on AWS. You can use Amazon CloudWatch to collect and track metrics, collect and monitor log files, and set alarms. Amazon CloudWatch can monitor AWS resources, such as Amazon EC2 instances, Amazon DynamoDB tables, and Amazon RDS DB instances, including custom metrics generated by your applications and services, and any log files that your applications generate, hosted on premises, hybrid, or on other clouds. You can use Amazon CloudWatch to gain system-wide visibility into resource utilization, application performance, and operational health. You can use these insights to react and keep your application running smoothly.

12.1.1 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.
- **Cross-account observability**: Cross-account observability in CloudWatch helps you monitor and troubleshoot applications that span multiple accounts within a Region.
- **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.
- **Internet monitor**: Internet monitor provides visibility into how internet issues impact the performance and availability between your AWS-hosted applications and your end users, reducing the time it takes for you to diagnose these issues from days to minutes.
- **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.
- **Anomaly detection**: Amazon CloudWatch Anomaly Detection applies machine-learning (ML) algorithms to continuously analyze metric data and identify anomalous behavior. It allows you to create alarms that auto-adjust thresholds based on natural metric patterns, such as time of day, day of week, seasonality, or changing trends. You can also visualize metrics with anomaly detection bands on dashboards. This enables you to monitor, isolate, and troubleshoot unexpected changes in your metrics.
- **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).

12.1.2 Benefits

- **Security**: Amazon CloudWatch is integrated with AWS Identity and Access Management (IAM) so you can control which users and resources have permission to access your data and how they can access it. Additionally, Amazon CloudWatch Logs data protection helps you to define data protection policies that can discover and protect sensitive data logged by systems and applications. This feature automatically identifies and masks sensitive information in your logs using machine learning and pattern matching based on the policy that you define.
- **Visualize and analyze your data with end-to-end observability**: Collect, access, and analyze your resource and application data using powerful visualization tools.
- **Operate efficiently with automation**: Improve operational performance using alarms and automated actions set to activate at predetermined thresholds.
- **Obtain an integrated view of your AWS or other resources**: Integrate with more than 70 AWS services for simplified monitoring and scalability.
- **Proactively monitor and get actionable insights to enhance end user experiences**: Troubleshoot operational problems with actionable insights derived from logs and metrics in your CloudWatch dashboards.

12.2 Backup/Restore and Disaster Recovery

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

12.3 Pricing Overview

Please see the AWS UK G Cloud 14 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/cloudwatch/index.html
- **Service quotas:** https://docs.aws.amazon.com/AmazonCloudWatch/latest/monitoring/cloudwatch_limits.html

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

12.5 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.
- CloudWatch CLI Reference: Describes the CloudWatch CLI commands that you can use to monitor AWS. Provides syntax, options, and usage examples for each command.

13 Amazon CodeCatalyst

13.1 Service Overview

Amazon CodeCatalyst is a unified software development service for development teams to quickly build, deliver and scale applications on AWS while adhering to organization-specific best practices. Developers can automate development tasks and innovate faster with generative AI capabilities, and spend less time setting up project tools, managing CI/CD pipelines, provisioning and configuring various development environments or coordinating with team members. IT Leaders can codify organizational best practices at scale via application blueprints to ensure compliance across teams with scale.

13.1.1 Features

- **Enhance developer productivity with generative AI capabilities**: Use Amazon Q feature development in CodeCatalyst to go from an idea in an issue to fully-tested, merge-ready application code with just natural language inputs, in just a few clicks.
- **Accelerate Project setup with built-in, codified best practices:** CodeCatalyst project blueprints automatically set up everything you need to start a new software development project, including CI/CD, deployable code, issue tracking, and AWS services configured according to best practices. With custom blueprints, IT leaders can build their own templates to specify technology to be used by developer teams, control access to project resources and define testing and building methods.
- **Spend more time coding and less time managing local development environments:** CodeCatalyst Dev Environments are available on-demand in the cloud and are automatically created with branch code and consistent project settings, providing faster setup, development, and testing.
- **Collaborate fluidly and scale projects across teams:** Start creating and assigning issues. Use priorities, estimates, labels, and custom fields to help your team prioritize what's next. Exchange feedback on issues and pull requests with your team. Access protected resources with SSO without repeated logins and set project access permissions for several team members at a time.
- **Quickly customize your pipelines and fully-managed build environments:** Simplify the creation and customization of automated workflows by configuring pre-defined actions using the visual editor or directly editing the underlying YAML

configuration. Along with a collection of CodeCatalyst actions, you can also use any GitHub Action in your workflows.

- **Level up your testing strategy with easily added quality gates:** CodeCatalyst has built-in support for code coverage, software composition analysis, and unit tests. You can view your test results at a glance with automated visual test reporting.
- **The easiest way to deploy to AWS:** CodeCatalyst makes it easy to deploy your application to AWS services such as AWS Lambda or Amazon ECS. You can deploy application stacks across accounts or AWS Regions by simply listing them as targets in your pipeline.
- **Track progress across your team:** CodeCatalyst's notifications and personalized activity feed help you stay updated on relevant project activity such as successful deployments or accepted pull requests.

13.1.2 Benefits

- **Speed of set up:** Accelerate project setup with blueprints to codify best practices
- **Collaboration:** Collaborate and scale fluidly across teams with centralized and secure access management
- **Automation:** Automate daily workflows and environment creation and management
- **Productivity Gains:** Improve developer productivity with generative AI capabilities using natural language inputs

13.2 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/.

13.3 Pricing Overview

Please see the AWS UK G Cloud 14 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/codecatalyst/
- **Service quotas:** https://docs.aws.amazon.com/codecatalyst/latest/userguide/quotas.html
- **Service FAQs:** https://codecatalyst.aws/explore/faq

13.5 Technical Requirements

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

- User Guide**:** Learn how to use Amazon CodeCatalyst to track work, develop code, and build, test, and deploy your applications.
- API Reference: Describes API operations for Amazon CodeCatalyst in detail. Also provides details of request and response syntax and errors for the supported web services protocols.
- Administrator Guide: Learn how to set up an Amazon CodeCatalyst space where users can collaborate on projects. Find out how to view your billing information and upgrade your subscription

- Action Developer Guide: Learn how to create custom actions for your workflows using the Amazon CodeCatalystAction Developer Guide.
- Set up AWS Support for Amazon CodeCatalyst: Learn how administrators can set up the support plan and roles required to create and manage support cases for their CodeCatalyst space. Once configured, space and project members can create support cases through AWS Support for Amazon CodeCatalyst.

14 Amazon CodeGuru

14.1 Service Overview

CodeGuru has two components: Amazon CodeGuru Security and Amazon CodeGuru Profiler.

- **CodeGuru Security** is a machine learning (ML) and program analysis-based tool that finds security vulnerabilities in your application code. CodeGuru Security also scans for hardcoded credentials.
- **CodeGuru Profiler** optimizes performance for applications running in production and identifies the most expensive lines of code, reducing operational costs significantly.

14.1.1 Features

Amazon CodeGuru Security

- **Integrations**: CodeGuru Security provides a simplified and flexible API that can be used to integrate and detect vulnerabilities at any stage of the development lifecycle, including integrated development environments (IDEs) and continuous integration and delivery (CI/CD).
- **Precise detection**: CodeGuru Security combines machine learning (ML) and automated reasoning to detect bugs with high precision so that you don't spend time on false positives.
- **Bug tracking**: CodeGuru Security uses powerful algorithms to detect when a code is fixed and closes the finding automatically. This simplifies bug tracking and removes the need to manually track bug fixes.
- **Assisted remediation**: CodeGuru Security uses automated reasoning to provide suggested code fixes for certain vulnerabilities. Use this to quickly apply code fixes and reduce the amount of time it takes to fix an issue.

Amazon CodeGuru Profiler

- **Always-on profiling of applications in production**: You can profile and troubleshoot your application using real customer traffic patterns and discover performance issues. With the profiler data and ML-powered recommendations, you can identify and fix performance issues for your applications in production.
- **Understand application runtime behavior**: CodeGuru Profiler continually 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 more 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.
- **Heap summary**: The heap usage analysis is presented on a heap summary visualization that shows you what objects are allocated on your heap—whether your own domain classes or those owned by libraries or the JDK.

- **Intelligent recommendations**: CodeGuru Profiler automatically identifies performance issues in your application and provides ML-powered recommendations on how to remediate them.
- **Anomaly detection**: CodeGuru Profiler continually analyzes your application profiles in real time and detects anomalies in the behavior of your application and its methods.

14.1.2 Benefits

- **Detect vulnerabilities at any stage of the development workflow**: The CodeGuru Security API-based design provides integrate capabilities to use at any stage of the development workflow. Whether your organization adheres to the “shift left“ or “shift right” ideology, CodeGuru Security plugs into your continuous integration and delivery (CI/CD) tooling to help you identify vulnerabilities in your application code.
- **Reduce false-positive detections**: False positives cost you money by diverting engineering resources to findings that are found not to be vulnerabilities after review. By performing deep semantic analysis, CodeGuru Security detects vulnerabilities with a high precision, significantly reducing the number of false positives. This frees up engineering teams to focus on building applications for your organization.
- **Track bug closure automatically**: The CodeGuru Security bug-tracking feature automatically detects when a bug is closed. The bug-tracking algorithm ensures that you have up-to-date information on your organization's security posture without additional effort. You focus on what matters to you, and CodeGuru Security takes care of the rest.
- **Start immediately without VM provisioning**: There is no need to provision virtual machines (VMs) to run CodeGuru Security. Just integrate CodeGuru Security with your tooling, and it will scale up and down with your workload.
- **Significantly decrease compute costs**: Amazon CodeGuru Profiler helps developers find an application’s most expensive lines of code by helping them understand the runtime behavior of their applications, identify and remove code inefficiencies, improve performance, and help to bring down compute costs.

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 14 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/security-ug/quotas.html for CodeGuru Security
  - https://docs.aws.amazon.com/codeguru/latest/profiler-ug/quotas.html for CodeGuru Profiler
  - 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.

CodeGuru Security

- User Guide: Provides a conceptual overview of Amazon CodeGuru Security, instructions for getting started, product and service integrations, and creating and understanding code scans.
- API Reference: Describes all the Amazon CodeGuru Security API operations in detail. Also provides sample requests, responses, and errors for the supported web services protocols.
- CodeGuru Security section of the AWS CLI Reference: Describes the AWS CLI commands that you can use to work with CodeGuru Security code scans.
- Detector Library: Lists all detectors that Amazon CodeGuru Security uses to analyze code. Provides detailed information about the detectors, their impact on applications, and ways to mitigate risk.

CodeGuru Profiler

- 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.

CodeGuru Reviewer

- 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.

15 Amazon Cognito

15.1 Service Overview

Amazon Cognito processes more than 100 billion authentications per month. The service helps you implement customer identity and access management (CIAM) into your web and mobile applications. You can quickly add user authentication and access control to your applications in minutes.

15.1.1 Features

- **Identity Management**: Amazon Cognito provides both a customizable, pre-packaged, hosted user interface to rapidly get to market and a robust set of APIs to

build a fully custom self-registration solution. Amazon Cognito provides a secure identity store (user pools) that scales to millions of users. Users can migrate into Amazon Cognito using either a batch import or just-in-time (JIT) migration. Amazon Cognito enables B2B interactions with multi-tenant support.

- **User Authentication**: You can use Android, iOS, and JavaScript SDKs for Amazon Cognito to add user sign-up and sign-in pages to your apps. You can add an additional layer of security for your customers by enabling multi-factor authentication (MFA) in an Amazon Cognito user pool. Users can verify their identities using SMS or a Time-based One-time Password (TOTP) generator, such as Google Authenticator.
- **Access Control**: The credential broker for Amazon Cognito, also known as Amazon Cognito identity pools, provides single sign-on access to AWS resources such as Amazon DynamoDB, Amazon S3 buckets, Lambda serverless components, and other Amazon services. Users can be dynamically mapped to different roles to support least privilege access to a service. Using the OAuth Client Credential Flow, Amazon Cognito also provides machine-to-machine authentication, ensuring a secure experience between application components.
- **Customer Experience**: Amazon Cognito provides a robust set of hooks and extensions to fully customize the authentication, registration, and user migration flows. For example, the self-registration flow can be augmented with custom identity proofing and account verification checks and the login process can be extended to create custom authentication flows or modify a token before it is generated.
- **Advanced Security**: When Amazon Cognito detects unusual sign-in activity, such as 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.
- **Auditing and Compliance**: Amazon Cognito is Payment Card Industry Data Security Standard (PCI DSS), System and Organization Controls (SOC), and ISO/IEC 27001, ISO/IEC 27017, ISO/IEC 27018, and ISO 9001 compliant.

15.1.2 Benefits

- **Customizable customer identity and access management**. Deliver frictionless customer identity and access management with a developer-centric, cost-effective, and customizable service.
- **Advanced security features for sign-up and sign-in**: Add advanced security features such as risk-based adaptive authentication, compromised credential monitoring, and security metrics to support your compliance and data residency requirements.
- **High-performing, scalable user directory**: Scale to millions of users with a fully managed, high-performance, and reliable user directory.
- **Standards-based, federated sign-in capabilities**: Federate sign-in by using OpenID Connect or SAML 2.0 and connect to a broad group of AWS services and products.

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 14 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/getting-started/

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 find meaning and insights in text. NLP is a way for computers to analyze, understand, and derive meaning from textual information in a smart and useful way. By utilizing NLP, you can extract important phrases, sentiment, syntax, key entities such as brand, date, location, person, etc., and the language of the text.

16.1.1 Features

- **Custom Entity Recognition**: 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.
- **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.
- **Entity Recognition**: The Entity Recognition API returns the named entities ("People," "Places," "Locations," etc.) that are automatically categorized based on the provided text.
- **Sentiment Analysis**: The Sentiment Analysis API returns the overall sentiment of a text (Positive, Negative, Neutral, or Mixed).
- **Targeted Sentiment**: Targeted Sentiment provides more granular sentiment insights by identifying the sentiment (positive, negative, neutral, or mixed) towards entities within text.
- **PII Identification and Redaction**: Use Amazon Comprehend ML capabilities to detect and redact personally identifiable information (PII) in customer emails, support tickets, product reviews, social media, and more.

- **Toxicity detection**: Comprehend toxicity detection provides a simple, NLP-based solution for toxic content detection in text-based documents. The capability is available out-of-the-box to moderate peer-to-peer conversation in online platforms and generative AI inputs and outputs.
- **Prompt Safety Classification**: Comprehend provides a pre-trained binary classifier that can classify the input prompt as harmful or not. This can be integrated to allow LLMs to only respond to harmless content.
- **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.
- **Events Detection**: Comprehend Events lets you extract the event structure from a document, distilling pages of text down to easily processed data for consumption by your AI applications or graph visualization tools. This API allows you to answer who-what-when-where questions over large document sets, at scale and without prior NLP experience. Use Comprehend Events to extract granular details about real-world events and associated entities expressed in unstructured text.
- **Language Detection**: The Language Detection API automatically identifies text written in over 100 languages and returns the dominant language with a confidence score to support that a language is dominant.
- **Syntax Analysis**: The Amazon Comprehend Syntax API enables customers to analyze text using tokenization and Parts of Speech (PoS), and identify word boundaries and labels like nouns and adjectives within the text.
- **Topic Modeling**: Topic Modeling identifies relevant terms or topics from a collection of documents stored in Amazon S3. It identifies common topics in the collection and organizes them in groups and maps the documents to topics.
- **Multiple language support**: Amazon Comprehend can perform text analysis on German, English, Spanish, Italian, Portuguese, French, Japanese, Korean, Hindi, Arabic, Chinese (simplified), Chinese (traditional) text. To build applications in other languages, customers can use Amazon Translate to convert the text into a language supported by Comprehend and then use Comprehend to perform text analysis.

16.1.2 Benefits

- **Extract insights:** 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 14 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 natural language processing service that makes it easy to use machine learning to extract relevant medical information from unstructured text. Using Amazon Comprehend Medical, you can quickly and accurately gather information, such as medical condition, medication, dosage, strength, and frequency from a variety of sources, like doctor's notes, clinical trial reports, and patient health records.

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.
- **Medical Ontology Linking**: The Medical Ontology Linking APIs identifies medical information and links them to codes and concepts in standard medical ontologies. The Medical Ontology Linking APIs also detects contextual information as entity traits (e.g. negation). You can use ontology linking batch analysis to analyze either a collection of documents or a single document.
- **Detect protected health information**: Healthcare and life science customers often have data that contains protected health information (PHI) that need to be moved/tagged before further processing, such as de-identifying PHI to comply with HIPAA before conducting research. Amazon Comprehend Medical can detect PHI and uses the Safe Harbour guidelines to search for PHI.

17.1.2 Benefits

- **Quickly and accurately extract medical information**: 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.
- **Identify relationships among the extracted 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, RxNorm or SNOMED CT to help you build applications for use cases like revenue cycle management (medical coding), claim validation and processing, and electronic health record creation.

- **Fully managed**: 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.
- **Data privacy**: Amazon Comprehend Medical does not use content processed by the service for any reason other than to provide and maintain the service. Content processed by the service is not used to develop or improve the quality of Amazon Comprehend Medical or other Amazon machine-learning/artificial-intelligence technologies, and will not be stored in any AWS region other than the region in which you are using the service. You always retain ownership of your content and we will only use your content with your consent.
- **Security**: Requests to the Amazon Comprehend Medical APIs and console are made over a secure (SSL) connection. You can use AWS Identity and Access Management (AWS IAM) to control which IAM users have access to specific Amazon Comprehend Medical actions and resources.

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 14 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/
- **Service quotas:** https://docs.aws.amazon.com/comprehend-medical/latest/dev/comprehendmedical-quotas.html
- **Service FAQs:** https://aws.amazon.com/comprehend/medical/faqs/

17.5 Technical Requirements

Please refer to https://docs.aws.amazon.com/comprehend-medical/ 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 the Amazon Comprehend API operations and parameters in detail.

18 Amazon Connect (Lot 2)

18.1 Service Overview

Amazon Connect empowers contact center agents to deliver superior customer experiences from day one. With Amazon Q, a generative AI-powered assistant available in Amazon Connect, customer issues are automatically detected, and agents are provided contextual customer information and suggested responses and actions for faster resolution, all within the unified workspace. In addition, you can guide your agents to resolve customer issues quickly and accurately by automatically recommending appropriate actions to take with step-by-step guides.

18.1.1 Features

Omnichannel and self-service

- **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. The telephony service includes direct inward dial (DID) and toll-free phone numbers for over 20 countries worldwide. There are also over 200 available outbound calling destinations.
- **In-app, web, and video calling**: Amazon Connect's in-app, web, and video calling capabilities make it easy to deliver personalized voice and video experiences in your websites and mobile applications. Using the fully managed communication widget or SDK, you can implement Amazon Connect's in-app, web, and video calling capabilities with as little as a single line of code.
- **High-quality audio**: 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/mobile chat and Short Message Service (SMS)**: Amazon Connect Chat supports asynchronous messaging, so your customers and agents have the ability to message without being available at the same time. Amazon Connect also offers native integration with Apple Business Chat, helping your agents to support customers through the same popular iOS Messaging application that they use every day. Amazon Connect's two-way SMS messaging capabilities help contact centers text with customers using the same routing, configuration, analytics, and management tools they already use for calls and chats. Amazon Connect SMS uses Amazon Lex to detect the intent of the customer message and automate responses without agent involvement.
- **Natural, conversational IVR and chatbots**: With Amazon Lex natively integrated within Amazon Connect, no coding is required to add chatbots that have natural language understanding (NLU). Amazon Connect enables intelligent automation at all stages of the customer journey with fully-integrated, self-service capabilities including chatbots, task routing, and interactive voice response.
- **Drag-and-drop workflow desginer**: Amazon Connect flows is a single drag-and-drop workflow designer you can use to create, personalize, and automate end-to-end customer and agent experiences across channels. With flows, you can design your interactive voice response (IVR) or chatbot experiences to help your customers self-serve, build step-by-step guides for your agents to resolve issues faster and accurately, and create and manage how tasks are automated for your agents.
- **Omnichannel outbound campaigns**: Amazon Connect outbound campaigns helps you communicate across voice, SMS, and email to serve your customers quickly and improve agent productivity, while supporting compliance with local regulations.

- **Predictive dialer with answering machine detection**: Amazon Connect outbound campaigns includes a predictive dialer that is designed to automatically call customers on a list, while throttling outreach based on predicted agent availability. The dialer also uses a machine learning (ML) model to distinguish between a live customer, voicemail greeting, or busy signal, connecting agents only with live customers.
- **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 center agents don't have to learn and work across multiple tools.
- **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.
- **Task management**: Amazon Connect Tasks makes it easier for you to prioritize, assign, and track agent tasks to completion, including work in external applications, to ensure customer issues are quickly resolved.
- **Rules engine**: Rules enable quality managers to automatically check conversations for compliance with company policies or regulatory requirements by specifying criteria (such as mentions of words or phrases in the conversation, customer sentiment, etc.) using Amazon Connect Contact Lens conversational analytics. Rules also empowers supervisors with real-time alerts on customer experience and agent performance, so they can assist agents with handling contacts.
- **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. Voice ID analyzes callers' 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. Voice ID also screens for fraudulent actors in real time, based on your contact center's custom watchlist, reducing potential losses from fraudulent attacks.

Agent productivity

- **Agent workspace**: Amazon Connect agent workspace is a single, intuitive application that provides your agents with all the tools and step-by-step guidance needed to onboard quickly, resolve issues efficiently, and improve customer experiences. From one application your agents can view detailed customer information, work on tasks, view workforce schedules, get generative AI-powered agent assist, and track and manage customer issues that require multiple interactions.
- **Step-by-step agent guides**: Create tailored guides, using the no-code editor that walk agents through the optimal steps to resolve a customer issue accurately the first time. Guides can be used for various types of customer interactions and are presented to the agent in the agent workspace based on context like call queue, customer information, or customer self-service responses.
- **Contact Control Panel**: The Amazon Connect Contact Control Panel (CCP) provides a single, intuitive communication interface for agents to receive calls and chats with contacts. Agents have the ability to transfer contacts to other agents, put contacts on hold, and perform other tasks.
- **Unified customer profiles**: Amazon Connect Customer Profiles provides customer service agents automated experiences with real-time access to up-to-date customer info to personalize each customer interaction. With a complete view of relevant customer information in a single place, using generative AI to combine contact

history information (e.g., transcripts, customer sentiment) with customer information (e.g., current product orders, mobile app interactions), your agents can provide more personalized customer service, addressing issues faster and improving CSAT.

- **Case management**: Amazon Connect Cases allows agents to efficiently manage customer issues that require multiple interactions, track follow-up tasks, and access subject matter experts across a business. Agents can document customer issues in a single, unified view with relevant case details, such as date/time opened, issue summary, customer information, and status.

Analytics, insights, and optimization

- **Contact forecasting**: Predict your customer service contact volumes with high accuracy using Amazon Connect forecasting. With at least 6 months of historical contact data, forecasting uses ML to identify patterns and offer accurate short-term (15- or 30-minute intervals) and long-term (daily, weekly, and monthly) forecasts to help optimize operational accuracy and efficiency. Forecasting can be updated and published to improve capacity planning and agent scheduling.
- **Capacity planning**: Optimize your staffing levels and better manage labor costs with Amazon Connect capacity planning. Estimate full-time equivalent (FTE) and staffing budgets and conduct what-if analysis to fine-tune service-level targets. A primary output of capacity planning is understanding how many FTEs are required to meet the service-level target for a certain period of time. This FTE requirement can be shared with other stakeholders, such as HR, finance, and training departments, to facilitate the staff hiring and training processes.
- **Scheduling**: Contact center schedulers can adjust and view the impact of these changes on agent occupancy and staffing. Contact center schedules can be published for agents to view in the Amazon Connect agent workspace. You can also give your agents the flexibility to choose when they want to work overtime or take time off within predetermined parameters, without the need for manual manager approvals.
- **Real-time conversational analytics with sentiment analysis**: Amazon Connect Contact Lens allows you to better understand the sentiment and trends of conversations and ensure your agents remain in compliance when helping customers. This helps supervisors train agents, replicate successful interactions, and identify crucial company feedback.
- **Key Highlights**: Automatically identify key parts of the customer conversation, assign tags (such as issue, outcome or action items) and display highlights of the customer interaction that can be expanded to view the full transcript of the contact.
- **Post-contact summarization**: Streamline and automate post contact work with generative AI-powered summaries of customer conversations that are structured, concise, and easy to read. These summaries capture important information from customer conversations, so you can quickly review them, understand context, and follow up with customers or take other necessary actions.
- **Data redaction**: Data redaction, or data masking, automatically detects sensitive data such as name, address, and social security number and removes or hides sensitive data in call recordings and transcripts.
- **Performance evaluation and contact scoring**: Define and create a set of agent performance evaluation forms and complete the evaluations side by side with call recordings, transcripts, and conversational analytics outputs such as contact categories, sentiment scores, and issues detected.
- **Screen recording**: Amazon Connect Contact Lens screen recording enables companies to easily record agents' screens along with their audio in Amazon Connect, providing managers the ability to not only listen to, but also watch agents'

actions while handling a customer contact (voice call, chat or tasks). By using screen recording with ML-powered conversational analytics and performance evaluation, managers can monitor, evaluate, and improve contact quality and agent performance.

- **Real-time and historical analytics**: Understanding your contact center 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 analyze them with other data such as conversion rates or customer satisfaction.
- **Call recording**: Amazon Connect comes with integrated call recording for agent performance assessment to help monitor and improve customer experiences.
- **Partner integrations**: Amazon Connect has a broad set of technology partners who provide integrations that can be quickly deployed in a few steps. There are partner integrations for workforce management and organization (WFM/WFO), mapping and location services, sales and service (CRM), and more.

18.1.2 Benefits

- **Increase agent productivity with generative AI**: Amazon Connect empowers contact center agents to deliver superior customer experiences from day one. With Amazon Q, a generative AI-powered assistant available in Amazon Connect, customer issues are automatically detected, and agents are provided contextual customer information and suggested responses and actions for faster resolution, all within the unified workspace. In addition, you can guide your agents to resolve customer issues quickly and accurately by automatically recommending appropriate actions to take with step-by-step guides.
- **Measure, track, and improve performance**: Using AI-powered analytics and optimization capabilities, supervisors and managers can proactively detect and address issues with their customer experience (CX), agent performance, and contact center operations. Supervisors and managers gain instant access to real-time insights from customer interactions, easily identify agent coaching needs, forecast contact volume to optimize staffing, and make the best use of their agents to continuously improve customer satisfaction, reduce costs, and optimize contact center operations.
- **Create a seamless omnichannel experience**: Amazon Connect empowers you to deliver personalized, efficient, and proactive experiences across the channels customers prefer. With AI-powered chatbots, you can deliver natural and intuitive self-service experiences in multiple languages that save customers time and effort. Amazon Connect also helps you proactively engage customers at scale with relevant information, like appointment reminders, in their preferred channel.

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 14 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 center.
- Amazon Connect API Reference: Describes APIs to programmatically set up and manage Amazon Connect.

19 Amazon Data Firehose

19.1 Service Overview

Amazon Data Firehose provides the easiest way to acquire, transform, and deliver data streams within seconds to data lakes, data warehouses, and analytics services. To use Amazon Data Firehose, you set up a stream with a source, destination, and required transformations. Amazon Data Firehose continuously processes the stream, automatically scales based on the amount of data available, and delivers it within seconds.

Data Firehose can capture, transform, and load streaming data into Amazon S3, Amazon Redshift, Amazon OpenSearch Service, and Splunk, enabling near real-time analytics with existing business intelligence tools and dashboards you're already using today. It is a fully managed service that automatically scales to match the throughput of your data and requires no ongoing administration. It can also batch, compress, and encrypt the data before loading it, minimizing the amount of storage used at the destination and increasing security.

19.1.1 Features

- **Easy launch and configuration**: You can launch Amazon Data Firehose and create a Firehose stream to deliver data into Amazon S3, Amazon Redshift, Amazon OpenSearch Service, Snowflake, Splunk, HTTP endpoints, Datadog, New Relic, or MongoDB with just a few clicks in the AWS Management Console. You can send data to the Firehose stream by calling the Firehose API, or running the Linux agent we provide on the data source. Firehose then continuously delivers 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 Firehose 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 Firehose 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.

- **Integration with Amazon MSK and Firehose Streams**: Firehose can read data easily from an existing Amazon MSK cluster or an existing Firehose stream and load it to multiple destinations.
- **Apache Parquet or ORC format conversion**: 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. 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”. 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 Data Firehose to prepare your streaming data before it is loaded to data stores. Simply select an AWS Lambda function from the Amazon Data Firehose stream configuration tab in the AWS Management console. Amazon Data Firehose will automatically apply that function to every input data record and load the transformed data to destinations. Amazon 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 Data Firehose to automatically retry failed jobs and back up the raw streaming data.
- **Support for multiple data destinations**: Amazon 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 Data Firehose provides you the option to have your data automatically encrypted after it is uploaded to the destination. As part of the Firehose stream configuration, you can specify an AWS Key Management System (KMS) encryption key.
- **Metrics for monitoring performance**: Amazon 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 Firehose stream limits, throttled records number and upload success rate.
- **Pay-as-you-go pricing**: With Amazon 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.

19.1.2 Benefits

- **Fully managed**: Amazon 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, Kinesis Data Analytics, generic HTTP endpoints, and service providers like Datadog, New Relic, MongoDB, and Splunk, enabling near real-time analytics and insights.

- **Easy to use**: 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 provision and scale compute, memory, and network resources without ongoing administration.
- **No need to create processing pipelines**: Transform raw streaming data into formats like Apache Parquet, and dynamically partition streaming data without building your own processing pipelines.
- **Stream into data lakes and warehouses**: Stream data into Amazon S3 and convert data into required formats for analysis without building processing pipelines.
- **Boost security**: Monitor network security in real time and create alerts when potential threats arise using supported Security Information and Event Management (SIEM) tools.
- **Build ML streaming applications**: Enrich your data streams with machine learning (ML) models to analyze data and predict inference endpoints as streams move to their destination.

19.2 Backup/Restore and Disaster Recovery

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 Data Firehose delivery stream if you made one of the following choices:

- For Amazon S3 destinations, streaming data is delivered to your S3 bucket. If data transformation is enabled, you can optionally back up source data to another Amazon S3 bucket.

For Amazon Redshift destinations, streaming data is delivered to your S3 bucket first. Amazon Data Firehose then issues an Amazon Redshift COPY command to load data from your S3 bucket to your Amazon Redshift cluster. If data transformation is enabled, you can optionally back up source data to another Amazon S3 bucket.

19.3 Pricing Overview

Please see the AWS UK G Cloud 14 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/firehose/
- **Service quotas:** https://docs.aws.amazon.com/firehose/latest/dev/limits.html
- **Service FAQs:** https://aws.amazon.com/kinesis/data-firehose/faqs/

19.5 Technical Requirements

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

- Amazon Data Firehose section of the AWS CLI Reference: Describes the command line interface for Firehose. Provides basic syntax, options, and usage examples for each command.

20 Amazon DataZone

20.1 Service Overview

Amazon DataZone is a data management service that makes it faster and easier for customers to catalog, discover, share, and govern data stored across AWS, on premises, and third-party sources. With Amazon DataZone, administrators and data stewards who oversee an organization's data assets can manage and govern access to data using fine-grained controls. These controls are designed to ensure access with the right level of privileges and context. Amazon DataZone makes it easier for engineers, data scientists, product managers, analysts, and business users to access data throughout an organization so that they can discover, use, and collaborate to derive data-driven insights.

20.1.1 Features

- **Amazon DataZone business data catalog:** Search for published data, request access, and start working with your data in days instead of weeks.
- **Amazon DataZone projects:** Collaborate with teams through data assets, and manage and monitor data assets across projects.
- **Amazon DataZone portal:** Access analytics with a personalized view for data assets through a web-based application or API.
- **Amazon DataZone governed data sharing:** Ensure that the right data is accessed by the right user for the right purpose with a governed workflow.

20.1.2 Benefits

- **Govern Data:** Govern data across organizational boundaries.
- **Connect people and data:** Connect people and data through shared tools to drive business insight.
- **Automate data:** Automate data discovery and cataloguing with machine learning.

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 14 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/datazone/latest/userguide/what-is-datazone.html
- **Service quotas:** https://docs.aws.amazon.com/datazone/latest/userguide/datazone-limits.html
- **Service FAQs:** https://aws.amazon.com/datazone/faqs/

20.5 Technical Requirements

Please refer to https://docs.aws.amazon.com/datazone/latest/userguide/what-is-datazone.html and the following links for comprehensive technical documentation regarding this service.

- Amazon DataZone User Guide: Provides a conceptual overview of Amazon DataZone and instructions on using various features with both the AWS Management Console and the Amazon DataZone data portal.
- API Reference: Describes all the API operations for Amazon DataZone in detail.

Please refer to and the following links for comprehensive technical documentation regarding this service.

21 Amazon Detective

21.1 Service Overview

Amazon Detective makes it easier to analyze, 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.

21.1.1 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 analyzes events from data sources, such as AWS CloudTrail logs, Amazon VPC Flow Logs, Amazon EKS audit logs, Amazon GuardDuty findings, AWS Security Hub findings, other integrated AWS security services, and maintains up to a year of aggregated data for analysis.
- **Consolidates disparate events into a graph model**: Amazon Detective can analyze trillions of events from various data types, including IP traffic, AWS management operations, and potentially malicious or unauthorized activities. Detective constructs a graph model using machine learning, statistical analysis, and graph theory to build a linked set of data for security investigations. The pre-built graph model contains security-related relationships and offers contextual, and behavioral insights that enable you to quickly validate, compare, and correlate the data to reach conclusions.
- **Interactive visualizations for efficient investigation**: Amazon Detective provides interactive visualizations and insights using generative AI, making it easier to investigate issues faster and more thoroughly with less effort. With a unified view that enables you to visualize all the context and natural language summaries in one place, it becomes more easier to identify patterns that can validate or refute a security issue and understand all of the impacted resources within a security finding.
- **Seamless integration for investigating a security finding**: Amazon Detective is integrated with AWS security services such as Amazon GuardDuty, AWS Security Hub, Amazon Inspector, Amazon Security Lake as well as AWS Partner security products to help quickly investigate security findings identified in these services. Using a single-step from these integrated services you can go to Amazon Detective and immediately see events related to findings, drill down into relevant historical activities and investigate the issue.

- **Security investigation support for Amazon GuardDuty Runtime Monitoring**: Amazon Detective supports security investigations for GuardDuty ECS and EKS Runtime Monitoring, providing enhanced visualizations and additional context for new threat detections. You can use the runtime threat detections from GuardDuty and the investigative capabilities from Detective to improve your detection and response for potential threats to your container workloads.
- **Simple deployment with no upfront data source integration or complex configurations to maintain**: With few steps 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.

21.1.2 Benefits

- **Determine potential security issues.** Through a unified view of user and resource interactions, Amazon Detective simplifies the investigative process and helps security teams conduct faster and more effective investigations.
- **Save time and effort with graph models**. Amazon Detective's prebuilt data aggregations, summaries, and context help you to quickly analyze and determine the nature and extent of possible security issues.
- **Investigate and respond to security findings with streamlined visualizations**. Amazon Detective maintains up to a year of aggregated data and makes it easily available through a set of visualizations that shows changes in the type and volume of activity over a selected time window, and links those changes to security findings.
- **Accelerate security investigations with generative AI**: Finding groups are collections of security findings and resources that are associated with a single potential security incident you should investigate together. You can start your investigation with finding groups, which offer a more complete understanding of the incident. It also offers interactive visualizations that allows you to explore specific findings and insights using generative AI to describe the chain of events in natural language.

21.2 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/.

21.3 Pricing Overview

Please see the AWS UK G Cloud 14 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/detective/
- **Service quotas:** https://docs.aws.amazon.com/general/latest/gr/detective.html
- **Service FAQs:** https://aws.amazon.com/detective/faqs/

21.5 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.

22 Amazon DevOps Guru

22.1 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 behaviors 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 help identify anomalous application behavior (for example, increased latency, error rates, resource constraints, and others) and 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, and context for when and where the issue occurred.

22.1.1 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 and logs. 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.

22.1.2 Benefits

- **Detect abnormal application behavior using machine learning (ML) models**: Informed by years of Amazon.com and AWS operational excellence, Amazon DevOps Guru operational insights aggregate the information needed to investigate and remediate an operational issue directly in the DevOps Guru console.
- **Receive insights and contextual information about anomalous behavior**: An insight is composed of three main sections. It highlights the anomalous metrics and logs related to the operational issue, with graphs to easily visualize abnormal system and application behavior. The insight also includes contextual information such as relevant events and log snippets so you can easily understand the scope and issue timeline. Operational Insights also include recommendations on actions you can take to remediate the issue.
- **Automatically analyze application metrics, logs, and events**: DevOps Guru is designed to continuously analyze the operational data for your AWS resources based on your selection and produces insights whenever it detects ongoing or emergent operational issues. You can choose your analysis coverage boundary to be your entire AWS account or, you can prescribe the specific AWS CloudFormation stacks that you want DevOps Guru to analyze, or use AWS tags to create the resource grouping you want DevOps Guru to analyze.
- **Use ML models to limit alarm noise and focus on remediation and responses**: DevOps Guru uses pre-trained ML models to identify deviations from the established baseline of normal operating bounds. When it identifies anomalous application behavior like increased latency, error rates, or resource constraints that could cause

potential outages or service disruptions, it alerts operators with issue details like the resources involved, the issue timeline, and other related events to help them quickly understand the potential impact and likely causes of the issue.

22.2 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/.

22.3 Pricing Overview

Please see the AWS UK G Cloud 14 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/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/

22.5 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.

23 Amazon DocumentDB (with MongoDB compatibility)

23.1 Service Overview

Amazon DocumentDB (with MongoDB compatibility) is a fast, scalable, highly available, and fully managed enterprise document database service that supports native JSON workloads. As a document database, Amazon DocumentDB makes it easy to store, query, and index JSON data. Developers can use the same MongoDB application code, drivers, and tools as they do today to run, manage, and scale workloads on Amazon DocumentDB. Enjoy improved performance, scalability, and availability without worrying about managing the underlying infrastructure.

23.1.1 Features

Performance at scale

- **Amazon DocumentDB Elastic Clusters**: Amazon DocumentDB Elastic Clusters enables customers to handle millions of writes and reads per second, allowing customers to scale their document databases in minutes with little to no downtime or impact to performance. Customers can also store petabytes of data and only pay for

the capacity they consume with zero management of underlying infrastructure. With Amazon DocumentDB Elastic Clusters customers can now meet the scaling needs of virtually any application.

- **High throughput, low latency for document queries**: 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.
- **Easy scaling of database compute resources**: With a few clicks in the AWS Management Console, customers 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.
- **Automatic storage scaling**: Amazon DocumentDB will automatically grow the size of the storage volume as your cluster storage needs grow. The storage volume will grow in increments of 10 GB up to a maximum of 4 PiB. Customers don't need to provision excess storage for your document 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 DocumentDB replicas share the same underlying storage as the source instance, lowering costs and avoiding the need to perform writes at the replica nodes. This frees up more processing power to serve read requests and reduces the replica lag time–often down to single digit milliseconds. Amazon DocumentDB also provides a single endpoint for read queries, so the application can connect without the need to keep track of replicas as they are added and removed.
- **MongoDB-compatible**: Amazon DocumentDB is compatible with MongoDB 3.6, 4.0, and 5.0 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, 4.0, and 5.0 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. Learn more about supported MongoDB APIs.
- **Geospatial query capabilities**: The launch of Geospatial query capabilities enables customers to use Amazon DocumentDB to support storing, querying and indexing Geospatial data. Customers can create 2dsphere indexes and use popular MongoDB geospatial APIs such as $nearSphere, $geoNear, $minDistance, $maxDistance to perform queries stored on data stored DocumentDB.
- **ACID transactions**: ACID (atomicity, consistency, isolation, durability) is a set of properties of database transactions intended to guarantee data validity despite errors, power failures, and other mishaps. With the launch of support for MongoDB 4.0 compatibility, Amazon DocumentDB supports the ability to perform ACID transactions across multiple documents, statements, collections, and databases.
- **Migration support**: Customers can easily migrate their MongoDB databases on-premises or on Amazon Elastic Compute Cloud (EC2) to Amazon DocumentDB for free (for six months per instance) with virtually no downtime using the AWS Database Migration Service (DMS). With DMS, customers can migrate from a MongoDB replica set or from a sharded cluster to Amazon DocumentDB. For more information about migrating both relational and non-relational databases to Amazon DocumentDB, see Migrating to Amazon DocumentDB.

Cost effective

- **Pay only for what you use**: There is no upfront commitment with Amazon DocumentDB. You pay an hourly charge for each instance that you launch, and when you're finished with an Amazon DocumentDB instance, you can delete or pause it. You do not need to overprovision storage as a safety margin, and you only pay for the storage you actually consume. To see more details, visit the Amazon DocumentDB pricing page.
- **Price predictability at any scale**: Amazon DocumentDB offers I/O-Optimized storage configuration for those seeking price predictability. Amazon DocumentDB I/O-Optimized offers up to 40% cost savings for I/O-intensive applications where I/O charges exceed 25% of the total Amazon DocumentDB I/O-Optimized database spend. With Amazon DocumentDB I/O-Optimized, you can effectively eliminate the uncertainty of variable I/O charges from your billing structure. Instead, you are billed only for compute, storage, and backup charges, ensuring price predictability and transparency.
- **Price-performance**: Amazon DocumentDB enables you to choose between Standard and I/O-Optimized storage configurations for your database cluster. The flexibility enables you to maximize price-performance by choosing the appropriate configuration based on your needs. If your application requires low to moderate I/O consumption, you should choose Standard configuration. If your use case demands I/O intensive workloads then you can opt for I/O-Optimized storage configuration.

Fully Managed

- **Automatic provisioning and setup**: 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 selected. Customers can launch a cluster and connect the application within minutes without additional configuration.
- **Monitoring and metrics**: Amazon DocumentDB provides Amazon CloudWatch metrics for the cloud database instances. Customers can use the AWS Management Console to view over 40 key operational metrics for the cluster, including compute, memory, storage, query throughput, MongoDB opcounters, and active connections.
- **Automatic software patching**: Amazon DocumentDB will keep customers database up-to-date with the latest patches. Customers can control if and when the cluster is patched via Database Engine Version Management.

Highly Secure and Compliant

- **Network isolation**: Amazon DocumentDB runs in Amazon Virtual Private Cloud (VPC), which allows customers to isolate the cluster in the virtual network and connect to on-premises IT infrastructure using industry- standard encrypted IPsec virtual private networks (VPNs). In addition, using Amazon DocumentDB's VPC configuration, customers can configure firewall settings and control network access to the cluster.
- **Authorization**: Amazon DocumentDB supports role-based access control (RBAC) with built-in roles and defined roles. RBAC enables customers to enforce least privilege as a best practice by restricting the actions that users are authorized to perform. Amazon DocumentDB is integrated with AWS Identity and Access Management (IAM) and provides customers the ability to control the actions that AWS IAM users and groups can take on specific Amazon DocumentDB resources, including clusters, instances, snapshots, and parameter groups. In addition, tag the Amazon DocumentDB resources, and control the actions that the IAM users and groups can take on groups of resources that have the same tag (and tag value).
- **Encryption**: Amazon DocumentDB allows customers to encrypt databases using keys created and controlled 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.

- **Compliance certifications**: Amazon DocumentDB was designed to meet the highest security standards and to make it easy for customers to verify our security and meet regulatory and compliance obligations. Amazon DocumentDB has been assessed to comply with PCI DSS, ISO 9001, 27001, 27017, and 27018, SOC 1, 2 and 3, and Health Information Trust Alliance Common Security Framework certification (HITRUST CSF), in addition to being HIPAA eligible.

Highly Available

- **Global clusters**: 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.
- **Multi-AZ deployments with read replicas**: If there is instance failure, Amazon DocumentDB automates failover to one of up to 15 Amazon DocumentDB replicas customers have created in any of three Availability Zones. If no Amazon DocumentDB replicas have been provisioned, in the case of a failure, Amazon DocumentDB will attempt to create a new instance for customers automatically.
- **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.
- **Automatic, continuous, incremental backups, and point-in-time restore**: Amazon DocumentDB's simple database backup capability enables point-in-time recovery for clusters. Customers can restore their cluster to any second during the retention period, up until the last five minutes. The automatic backup retention period can be configured up to thirty-five days. Automated backups are stored in Amazon Simple Storage Service (Amazon S3), which is designed for 99.999999999% durability. Amazon DocumentDB backups are automatic, incremental, and continuous and have no impact on cluster performance.
- **Cluster snapshots**: Cluster snapshots are user-initiated backups of cluster stored in Amazon S3 that will be kept until explicitly deleted. They leverage the automated incremental snapshots to reduce the time and storage required. Customers can create a new cluster from a Cluster Snapshot whenever desired.

Generative AI and machine learning

Amazon DocumentDB offers capabilities to enable machine learning (ML) and generative artificial intelligence (AI) models to work with data stored in Amazon DocumentDB in real

time. Customers no longer have to spend time managing separate infrastructure, writing code to connect with another service, and duplicating data from their primary database.

- **Vector search**: With vector search for Amazon DocumentDB, you can store, index, and search millions of vectors with millisecond response times. A vector is a numerical representation that represents the semantic meaning of unstructured data such as text, images, and video. You can store vectors from Amazon Bedrock, Amazon SageMaker, and other third party or propriety models. Learn how to get started by visiting our vector search for Amazon DocumentDB documentation.
- **No-code machine learning with Amazon DocumentDB and Amazon SageMaker Canvas**: Amazon DocumentDB integrates with Amazon SageMaker Canvas, making it easy to build generative applications using data stored in Amazon DocumentDB. The in-console integration removes the undifferentiated heavy lifting to connect and access Amazon DocumentDB to accelerate your AI/ML development with a low code no code (LCNC) experience. You can build AI/ML models for classic use cases such as regression and forecasting, or create generative AI solutions such as content generation, text extraction, and text summarization within SageMaker Canvas.

23.1.2 Benefits

- **Build applications faster**: Develop and evolve applications faster by storing, querying, indexing, and aggregating data, including vectors, in the same flexible JSON format that is generated in your applications.
- **Scale your compute power in seconds**: Scale to millions of document read and write requests per second with petabytes of storage by scaling compute and storage independently.
- **Improve productivity and lower total cost of ownership**: With a fully managed database that removes manual database management tasks without licensing fees, you can enhance productivity and cost of ownership.
- **High availability:** Gain enterprise capabilities such as high availability, durability, low-latency global reads, and built-in security best practices while using existing MongoDB APIs, drivers, and tools. Amazon DocumentDB automatically divides your storage volume into 10 GB segments spread across many disks. Each 10 GB chunk of your storage volume is replicated six ways, across three Availability Zones (AZs). Amazon DocumentDB is designed to transparently handle the loss of up to two copies of data without affecting write availability and up to three copies without affecting read availability. Amazon DocumentDB's storage volume is also self-healing. Data blocks and disks are continuously scanned for errors and repaired automatically.

23.2 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.

23.3 Pricing Overview

Please see the AWS UK G Cloud 14 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/documentdb/
- **Service quotas:** https://docs.aws.amazon.com/documentdb/latest/developerguide/limits.html
- **Service FAQs:** https://aws.amazon.com/documentdb/faqs/

23.5 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.

24 Amazon DynamoDB

24.1 Service Overview

Amazon DynamoDB is a fast and flexible nonrelational database service for any scale. DynamoDB enables customers to offload the administrative burdens of operating and scaling distributed databases to AWS so that they don't have to worry about hardware provisioning, setup and configuration, throughput capacity planning, replication, software patching, or cluster scaling.

24.1.1 Features

- **Key-value and document data models**: Amazon DynamoDB supports both key-value and document data models. As a NoSQL database, Amazon DynamoDB has a flexible schema so each item can have many different attributes. A flexible schema allows you to easily adapt as your business requirements change, without the burden of having to redefine the table schema as you would in relational databases.
- **Serverless that scales to zero**: With Amazon DynamoDB, there are no servers to provision, patch, or manage, and no software to install, maintain, or operate. Amazon DynamoDB on-demand pricing provides pay-as-you-go pricing, scales to zero, and automatically scales tables to adjust for capacity and maintains performance with zero administration.
- **ACID transactions**: Amazon DynamoDB is built for mission-critical workloads, including support for atomicity, consistency, isolation, and durability (ACID) transactions for applications that require complex business logic. Amazon DynamoDB supports 100 actions per transaction, improving developer productivity.
- **Active-active replication with global tables**: Amazon DynamoDB global tables provides active-active replication of your data across your choice of AWS Regions with 99.999% availability. Global tables are multi-active meaning you can writes and read from any replica, and your globally distributed applications can access data locally in the selected Regions to get single-digit millisecond read and write performance. Also, global tables automatically scale capacity to accommodate your multi-Region workloads. Global tables improve your application's multi-Region

resiliency and should be considered as part of your organization's business continuity strategy.

- **Amazon DynamoDB Streams as part of an event-driven architecture**: Amazon DynamoDB Streams is a change data capture capability. Whenever an application creates, updates, or deletes items in a table, Amazon DynamoDB Streams record a time-ordered sequence of every item level changes in near-real time making it ideal for event-driven architecture applications to consume and action the changes. All changes are de-duplicated and stored for 24 hours.
- **Secondary indexes**: Similar to all other database systems, you start by creating a table which is a collection of items. With Amazon DynamoDB, each item in the table has its own primary key. Many applications can also benefit from having one or more secondary keys to more efficiently search data using other attributes. Amazon DynamoDB offers the option to create both global and local secondary indexes, which lets you query the data in the table using a secondary or alternate key.
- **Fine grained access control**: With Amazon DynamoDB, there are no username or passwords. Amazon DynamoDB uses AWS identity and access management (IAM) to authenticate, create, and access resources. You can specify IAM policies and conditions that allow fine-grained access, restricting read or write access down to specific items and attributes in a table, based on the identity of that user. This enables customers to enforce security policies at the code level.
- **Encryption at rest**: Amazon DynamoDB encrypts all customer data at rest by default. Encryption at rest enhances the security of your data by using encryption keys stored in AWS Key Management Service (AWS KMS). With addition of AWS Database Encryption SDK, you can perform attribute-level encryption to further enforce granular access control on data within your table. Amazon DynamoDB enables you to build security-sensitive applications that meet strict encryption compliance and regulatory requirements.
- **Point-in-time recovery**: Point-in-time recovery (PITR) helps protect your Amazon DynamoDB tables from accidental write or delete operations. PITR provides continuous backups of your Amazon DynamoDB table data, and you can restore that table to any point in time up to the second during the preceding 35 days.
- **On-demand backup and restore**: On-demand backup and restore allows you to create full backups of your Amazon DynamoDB tables' data for data archiving, which can help you meet your corporate and governmental regulatory requirements. You can back up tables from a few megabytes to hundreds of terabytes of data and not affect performance or availability to your production applications. With AWS Backup integration, you can also copy on-demand backups cross-account and cross-Region, cost allocation tagging for backups, and transitioning backups to cold storage.
- **Private network connectivity**: Amazon DynamoDB supports gateway Virtual Private Cloud (VPC) endpoints and interface VPC endpoints for connections within a VPC or from on-premises data centers. You can configure private network connectivity from your on-premises applications to DynamoDB through interface VPC endpoints enabled with AWS PrivateLink. This helps simplify your private connectivity to DynamoDB and maintains compliance.
- **Read/write capacity modes**: Amazon DynamoDB provides capacity modes for each table: on-demand and provisioned. For workloads that are less predictable for which you are unsure whether you'll have high utilization, on-demand capacity mode takes care of managing capacity for you, and you pay only 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, Amazon 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, Amazon 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 Amazon DynamoDB APIs without changing code.
- **Standard Infrequent Access (Standard-IA) table class**: For data that is infrequently accessed, you can use the Amazon DynamoDB Standard-IA table class which helps reduce your Amazon DynamoDB costs by up to 60%. Standard-IA tables lower storage cost is designed for long-term storage of data that is infrequently accessed, such as application logs, historical gaming data, old social media posts, and more. It has the same availability, durability, and performance as Amazon DynamoDB Standard tables, which is the default and most cost-effective option for most workloads.
- **Auto scaling for improved cost efficiencies**: For tables using provisioned capacity, Amazon DynamoDB provides auto scaling of throughput and storage based on your previously set capacity by monitoring the performance usage of your application.
- **Reserved capacity**: If your application has predictable reads and writes throughput, you may consider reserved capacity, which can save you up to 54% over standard provisioned capacity rates for one-year term or up to 77% for three-year term.
- **Free tier**: Amazon DynamoDB offers a generous free tier for anyone seeking a risk-free way to experience Amazon DynamoDB's capabilities.
- **Microsecond latency with Amazon DynamoDB Accelerator**: Amazon DynamoDB Accelerator (DAX) is an Amazon DynamoDB-compatible caching service 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 Amazon DynamoDB tables by up to 10 times—taking the time required for reads from milliseconds to microseconds, even at millions of requests per second. Learn more about DAX capabilities and DAX pricing.
- **Bulk import/export from Amazon Simple Storage Service (S3)**: Bulk import/export from S3 helps you get more value from your data by removing the need to write code to move, transform, and copy your Amazon DynamoDB tables from one application, account, or region to another. Bulk import/export does not use your table's read or write capacity so you don't need to plan or provision additional capacity. The bulk import/export process is fully managed by Amazon DynamoDB.
- **Advanced streaming applications with Kinesis Data Streams for Amazon DynamoDB**: Amazon Kinesis Data Streams for Amazon DynamoDB captures item-level changes in your Amazon DynamoDB tables to power live dashboards, generate metrics, and deliver data into data lakes. Kinesis Data Streams enables you to build advanced streaming applications such as real-time log aggregation, real-time business analytics, and IoT data capture.
- **Change tracking with triggers**: Amazon DynamoDB integrates with AWS Lambda to provide triggers. Using triggers, you can automatically run a custom function when item-level changes in an Amazon DynamoDB table are detected. With triggers, you can build event-driven applications that react to data modifications in Amazon DynamoDB tables. The Lambda function can perform any actions you specify, such as sending a notification or initiating a workflow.
- **Monitoring and diagnosing system performance with Amazon Cloudwatch**: To easily monitor your database performance, Amazon DynamoDB is integrated with Amazon Cloudwatch which collects and processes raw database performance data. You can use Amazon CloudWatch to create customized views and dashboards of

metrics and alarms for your Amazon DynamoDB databases. This monitoring capability is offered by default and is complimentary.

24.1.2 Benefits

- **Performance at scale**: Deliver applications with consistent, high performance and nearly unlimited throughput and storage.
- **Designed for 99.999% SLA**: Unlock fast local reads and writes performance with multi-Region global tables database.
- **Scale-to-zero serverless**: Optimize costs with a fully manage serverless database that automatically scales to fit your needs.
- **Easy AWS integration**: Do more with your data using built-in tools to perform analytics, extract insights, and more.

24.2 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.

24.3 Pricing Overview

Please see the AWS UK G Cloud 14 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/dynamodb/
- Service **quotas:** https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/ServiceQuotas.html
- **Service FAQs:** https://aws.amazon.com/dynamodb/faqs/

24.5 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.

25 Amazon ECS Anywhere

25.1 Service Overview

Amazon ECS Anywhere is a feature of Amazon ECS that enables you to run and manage container-based applications on-premises, including on your own virtual machines (VMs) and bare metal servers. With ECS Anywhere, you do not need to install or operate local container orchestration software, thus reducing operational overhead. ECS Anywhere offers a completely managed solution that enables you to standardize container management across all of your environments.

25.1.1 Features

- **Consistent experience across your container-based applications**: 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 Amazon Systems Manager integration automatically and securely establishes trust between your on-premises hardware and AWS control plane.
- **Supported platforms and operating systems**: You can use ECS Anywhere with any VM (e.g., running on VMware, Microsoft Hyper-V, or OpenStack) or bare metal server running a supported operating system (OS). The ECS agent, software that allows a host to connect with the ECS control plane, is supported and tested for the long-term support (LTS) releases of Amazon Linux 2, Ubuntu, RHEL, SUSE, Debian, CentOS, and Fedora.

25.1.2 Benefits

- **Fully managed**: As a fully managed and highly scalable container orchestration solution, Amazon ECS makes it easy for you to run container-based applications on different types of compute capacity such as AWS Fargate and AWS Graviton2-powered instances as well as servers in your own data centers.
- **Easy cluster management**: ECS Anywhere extends the reach of Amazon ECS to provide you with a single management interface for all of your container-based applications, irrespective of the environment they're running in. As a result, you have a simple, consistent experience when it comes to cluster management, workload scheduling, and monitoring for both the cloud and on-premises.
- **No need to maintain container orchestration software**: With ECS Anywhere, you do not need to install and maintain any container orchestration software, thus removing the need for your team to learn specialized knowledge domains and skillsets for disparate tooling. ECS Anywhere makes it easy for you to run your applications in on-premises environments as long as desired and then migrate to the cloud with a single click at any time.

25.2 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.

25.3 Pricing Overview

Please see the AWS UK G Cloud 14 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/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/

25.5 Technical Requirements

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

26 Amazon EKS Anywhere

26.1 Service Overview

Amazon EKS Anywhere is container management software built by AWS that makes it easier to run and manage Kubernetes clusters on-premises and at the edge. Amazon EKS Anywhere is built on Amazon EKS Distro, which is the same reliable and secure Kubernetes distribution used by Amazon Elastic Kubernetes Service (Amazon EKS) in AWS Cloud. Amazon EKS Anywhere simplifies on-premises Kubernetes cluster management through the automation of undifferentiated heavy lifting such as infrastructure setup and Kubernetes cluster lifecycle operations. Amazon EKS Anywhere supports VMWare vSphere, bare metal, AWS Snowball Edge, Apache CloudStack and Nutanix for cluster infrastructure capacity. Amazon EKS Anywhere can run on Docker for dev/test deployments only. Amazon EKS Anywhere is built on the Kubernetes sub-project called Cluster API (CAPI), which provides declarative APIs and tooling to simplify the provisioning, upgrading, and operating of multiple Kubernetes clusters across different types of infrastructure.

26.1.1 Features

- **Operate Kubernetes clusters on your own infrastructure**: Amazon EKS Anywhere lets you create and operate Kubernetes clusters on your own infrastructure. Amazon EKS Anywhere builds on the strengths of Amazon EKS Distro and provides open-source software that's up to date and patched so you can have an on-premises Kubernetes environment that's more reliable than a self-managed Kubernetes offering.
- **Use Console to view Kubernetes clusers**: To view all of your Kubernetes clusters (including Amazon EKS Anywhere clusters) running anywhere through the EKS Connector, you can use the Amazon EKS console. EKS Connector is a software agent that runs on a Kubernetes cluster and enables the cluster to register with the Amazon EKS console. You can connect any conformant Kubernetes cluster to the Amazon EKS console through the EKS Connector, including Amazon EKS Anywhere clusters running on-premises, self-managed clusters on EC2, and other Kubernetes clusters running outside AWS. Once connected, you can use the Amazon EKS console to view all connected clusters and the underlying Kubernetes resources, regardless of where the cluster is running.
- **Supported infrastructure types**: Amazon EKS Anywhere supports different types of infrastructure including VMware vSphere, bare metal, AWS Snowball Edge, Apache CloudStack, and Nutanix. Amazon EKS Anywhere provides Bottlerocket, a Linux-based open-source operating system built by AWS, as the default node operating system. You can alternatively use Ubuntu and Red Hat Enterprise Linux (RHEL) as the node operating system.

26.1.2 Benefits

- **Deploy across hybrid environments**: Manage your Kubernetes clusters and applications in hybrid environments and run Kubernetes in your data centers.
- **Maintain data sovereignty**: Keep large datasets on premises and maintain data location legal requirements concerning data location.

- **Model machine learning (ML) workflows**: Efficiently run distributed training jobs using the latest Amazon Elastic Compute Cloud (EC2) GPU-powered instances, including Inferentia, and deploy training and inferences using Kubeflow.

26.2 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 Amazon S3. Users control this via User intervention and schedules. Users schedule and recover backups through a web interface.

26.3 Pricing Overview

Please see the AWS UK G Cloud 14 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/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/

26.5 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.

27 Amazon Elastic Block Store (EBS)

27.1 Service Overview

Amazon Elastic Block Store (Amazon EBS) provides block level storage volumes for use with EC2 instances. EBS volumes behave like raw, unformatted block devices. You can mount these volumes as devices on your instances. EBS volumes that are attached to an instance are exposed as storage volumes that persist independently from the life of the instance. You can create a file system on top of these volumes, or use them in any way you would use a block device (such as a hard drive). You can dynamically change the configuration of a volume attached to an instance.

27.1.1 Features

- SSD-backed volumes (IOPS-intensive)

  - **Provisioned IOPS SSD (io2 Block Express and io1) volumes**: Provisioned IOPS SSD volumes are designed to deliver a maximum of 256,000 IOPS, 4,000 MB/s of throughput, and 64 TiB in size per volume. Io2 Block Express is the latest generation of the Provisioned IOPS SSD volumes that delivers 4x higher throughput, IOPS, and capacity than io1 volumes, along with sub-millisecond latency. Io2 Block Express provides the highest block storage performance for the largest, most I/O-intensive, mission-critical deployments of Oracle, Microsoft SQL Server, SAP HANA, and SAS Analytics. When attached to Nitro System-based EC2 instances, io2 Block Express provides sub-millisecond latency, while io1 is designed to achieve single-digit millisecond latencies.
  - **General purpose SSD (gp3 and gp2) volumes**: General-purpose volumes are backed by solid-state drives (SSDs) and are suitable for a broad range of transactional workloads, virtual desktops, medium sized single instance databases, latency sensitive interactive applications, dev/test environments, and boot volumes. Amazon gp3 volumes are the latest generation of general-purpose SSD-based EBS volumes that enable customers to provision performance independent of storage capacity, while providing up to 20% lower pricing per GB than existing gp2 volumes. The new gp3 volumes deliver a baseline performance of 3,000 IOPS and 125 MB/s at any volume size. Customers looking for higher performance can scale up to 16,000 IOPS and 1,000 MB/s for an additional fee. Both gp3 and gp2 volumes are designed to offer single-digit millisecond latency and deliver the provisioned performance 99% of the time. We advise using io2 Block Express volumes if you require more IOPS than gp3 can offer, have a workload where low latency is essential, or require improved performance consistency or durability.
- HDD-backed volumes (MB/s-intensive)
  - **Throughput optimized HDD (st1) volumes**: ST1 is backed by hard disk drives (HDDs) and is ideal for frequently accessed, throughput intensive workloads with large datasets and large I/O sizes, such as MapReduce, Kafka, log processing, data warehouse, and ETL workloads. These volumes deliver performance in terms of throughput, measured in MB/s, and include the ability to burst up to 250 MB/s per TB, with a baseline throughput of 40 MB/s per TB and a maximum throughput of 500 MB/s per volume. ST1 is designed to deliver the expected throughput performance 99% of the time and has enough I/O credits to support a full-volume scan at the burst rate. To maximize the performance of st1, we recommend using EBS-optimized EC2 instances.
  - **Cold HDD (sc1) volumes**: SC1 is backed by hard disk drives (HDDs) and provides the lowest cost per GB of all EBS volume types. It is ideal for less frequently accessed workloads with large, cold datasets. Similar to st1, sc1 provides a burst model: these volumes can burst up to 80 MB/s per TB, with a baseline throughput of 12 MB/s per TB and a maximum throughput of 250 MB/s per volume. For infrequently accessed data, sc1 provides extremely inexpensive storage. SC1 is designed to deliver the expected throughput performance 99% of the time and has enough I/O credits to support a full-volume scan at the burst rate.
- **Data Protection**: 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. Amazon 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. 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.

- **Key Capabilities**: Amazon EBS 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. You can enable Multi-Attach on an EBS Provisioned IOPS volume to allow a volume to be concurrently attached to up to sixteen Nitro-based EC2 instances within the same Availability Zone. Amazon EBS Torn Write Prevention (TWP) ensures that full 16KiB write operations are persisted to the block storage. With this feature, you can turn off the double write operation that MySQL and MariaDB databases perform for increased database write throughput. These customers can increase the number of transactions processed per second (TPS) by up to 30% without having to over-provision or scale up their clusters, thereby reducing storage cost.

27.1.2 Benefits

- **Scale fast**: Scale fast for your most demanding, high-performance workloads, including mission-critical applications such as SAP, Oracle, and Microsoft products.
- **High performance**: Protect against failures with high availability, including replication within Availability Zone (AZs), and 99.999% durability with io2 Block Express volumes.
- **Optimize storage and cost**: 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.
- **Secure**: Encrypt your block storage resources without needing to build, maintain, and secure your own key management infrastructure. Prevent unauthorized access to your data by restricting public access and configuring locks on data backups.
- **Simple data protection**: Protect block storage data in the cloud, as well as on-premises block data using Amazon EBS Snapshots – a point-in-time copy that can be used to enable disaster recovery, migrate data across regions and accounts, and improve backup compliance. AWS further simplifies the lifecycle management of your snapshots through integration with Amazon Data Lifecycle Manager, which allows you to create policies so that you can automate multiple tasks including creation, deletion, retention, and sharing of snapshots.
- **Availability and durability**: Amazon EBS offers a higher durability volume (io2 Block Express), 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.

27.2 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.

27.3 Pricing Overview

Please see the AWS UK G Cloud 14 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/ebs/

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

27.5 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.
- AWS CLI Reference: Documents the AWS CLI commands for Amazon EBS.

28 Amazon Elastic Compute Cloud (EC2)

28.1 Service Overview

Amazon Elastic Compute Cloud (Amazon EC2) offers compute platform with over 750 instances and choice of the latest processor, storage, networking, operating system, and purchase model to help you best match the needs of your workload. AWS is the first major cloud provider to support Intel, AMD, and Arm processors, the only cloud with on-demand EC2 Mac instances, and the only cloud with 400 Gbps Ethernet networking. AWS offers the best price performance for machine learning training, and the lowest cost per inference instances in the cloud. More SAP, high performance computing (HPC), ML, and Windows workloads run on AWS than any other cloud.

Amazon EC2 provides a wide selection of instance types optimized 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.

Amazon EC2 instances are grouped into the following categories:

- **General purpose:** Provide a balance of compute, memory, and networking resources. These instances are ideal for applications that use these resources in equal proportions, such as web servers and code repositories.
- **Compute optimized:** Designed for compute intensive applications that benefit from high performance processors. These instances are ideal for batch processing workloads, media transcoding, high performance web servers, high performance computing (HPC), scientific modeling, dedicated gaming servers, ad server engines, and machine learning inference.
- **Memory optimized:** Designed to deliver fast performance for workloads that process large data sets in memory.
- **Storage optimized:** Designed for workloads that require high, sequential read and write access to very large data sets on local storage. They are optimized to deliver tens of thousands of low-latency, random I/O operations per second (IOPS) to applications.
- **Accelerated computing:** Use hardware accelerators, or co-processors, to perform functions, such as floating point number calculations, graphics processing, or data pattern matching, more efficiently than is possible in software running on CPUs.
- **High-performance computing:** Purpose built to offer the best price performance for running HPC workloads at scale on AWS. These instances are ideal for applications

that benefit from high-performance processors, such as large, complex simulations and deep learning workloads.

- **Previous generation:** AWS offers previous generation instance types for users who have optimized their applications around them and have yet to upgrade. We encourage you to use current generation instance types to get the best performance, but we continue to support previous generation instance types.

For details, see the AWS website at https://aws.amazon.com/ec2/instance-types/.

28.1.1 Features

- **Multiple Locations**: Amazon EC2 provides the ability to place instances in multiple locations. Amazon EC2 locations are composed of Regions and Availability Zones. Regions consist of one or more Availability Zones and are geographically dispersed. 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.
- **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.
- **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.
- **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.
- **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. Amazon cloud storage is reliable, scalable, and secure storage for your data. Amazon Elastic Block Store (Amazon EBS) is an easy to use, high performance block storage. Amazon Simple Storage Service (Amazon S3) is object storage with industry-leading scalability, availability, and security.
- **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.
- **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 modeling, 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.
- **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.
- **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.
- **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 varied 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.

28.1.2 Benefits

- **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.
- **Maintenance**: AWS regularly performs routine hardware, software, power, and network maintenance with minimal disruption across all EC2 instance types. This is achieved by a combination of technologies and methods across the entire AWS Global infrastructure, such as live update and live migration as well as redundant and concurrently maintainable systems. Non-intrusive maintenance technologies such as live update and live migration do not require instances to be stopped or rebooted. Customers are not required to take any action prior to, during or after live migration or live update. These technologies help improve application uptime and reduce your operational effort.
- **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. Cluster Compute, Cluster GPU, and High Memory Cluster instances have been specifically engineered to provide high-performance network capability and can be programmatically launched into clusters – allowing applications to get the

low-latency network performance required for tightly coupled, node-to-node communication. Cluster instances also provide significantly increased throughput making them well suited for customer applications that need to perform network-intensive operations.

- **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.
- **Flexibility**: Optimize performance and cost with flexible options like AWS Graviton-based instances, Amazon EC2 Spot instances, and AWS Savings Plans.
- **Security:** Provide secure compute for your applications. Security is built into the foundation of Amazon EC2 with the AWS Nitro System.
- **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.

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 14 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/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/

28.5 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.
- Amazon EC2 Instance Types: Describes the specifications for each instance type available in Amazon EC2.
- EC2 section of the AWS CLI Reference: Documents the AWS CLI commands for Amazon EC2.
- 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.

29 Amazon Elastic Container Registry (ECR)

29.1 Service Overview

Amazon ECR is a fully managed container registry that makes it easy for developers to share and deploy container images and artifacts. Amazon ECR is integrated with Amazon Elastic Container Service (Amazon ECS), Amazon Elastic Kubernetes Service (Amazon EKS), and AWS Lambda, simplifying your development to production workflow. Amazon ECR eliminates the need to operate your own container repositories or worry about scaling the underlying infrastructure. Amazon ECR hosts your images in a highly available and scalable architecture, allowing you to deploy containers for your applications reliably. Integration with AWS Identity and Access Management (IAM) provides resource-level control of each repository that lets you share images across your organization or with anyone in the world.

29.1.1 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.

29.1.2 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.

29.2 Backup/Restore and Disaster Recovery

This service allows for backup and recovery. The service can back up container images, image metadata, image scanning reports, and lifecycle policies.

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

29.5 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.
- Amazon ECR section of the AWS CLI Reference: Documents the Amazon ECR commands available in the AWS CLI.

30 Amazon Elastic Container Service (ECS)

30.1 Service Overview

Amazon ECS is a fully managed container orchestration service. ECS allows you to build, deploy, and manage containerized applications at any scale on AWS, in traditional Amazon Elastic Cloud Compute (EC2) instances or on a serverless compute plane with AWS Fargate. With simple API calls, you can launch and stop container-enabled applications, query the complete state of your cluster, and access many familiar features like security groups, Elastic Load Balancing (ELB), Amazon Elastic Block Store (EBS) volumes, and Identity Access Management (IAM) roles. You can use Amazon ECS to schedule container placement across your cluster based on your resource needs and availability requirements.

30.1.1 Features

- **Integrated and serverless by design with AWS Fargate**: AWS Fargate is built into Amazon ECS. You need todefine your application's requirements, select AWS Fargate as your launch type in the console or Command Line Interface (CLI), and AWS Fargate takes care of all the scaling and infrastructure management required to run your containers across flexible compute options, with automatic integrations to other supporting AWS services that your application needs.
- **Hybrid deployments**: With Amazon 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.
- **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.
- **Fully-managed container orchestration service**: Amazon ECS natively integrates with both AWS and third-party tools to make it easier for teams to focus on building the applications, not the environment.
- **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.
- **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.
- **Scheduling and Task Placement**: Amazon ECS includes multiple scheduling strategies that place containers across your clusters based on your resource needs (for example, CPU or RAM) and availability requirements. Using the available scheduling strategies, you can schedule batch jobs, long-running applications and services, and daemon processes. Amazon ECS task scheduling allows you to run

processes that perform work and then stop, such as batch processing jobs. 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.

- **Container auto-recovery**: 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 Block Store (Amazon EBS) is an easy-to-use, scalable, high-performance block-storage service that allows you to provision cost-performant storage for your Amazon ECS and AWS Fargate deployments. You can use all EBS configurations and capabilities with your container workloads deployed on Amazon ECS.
- **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.
- **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.
- **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.

30.1.2 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 PCI DSS Level 1, ISO 9001, ISO 27001, ISO 27017, ISO 27018, SOC 1, SOC 2, SOC 3, and HIPAA eligibility.

30.2 Backup/Restore and Disaster Recovery

This service allows for backup and recovery. The service can back up application container images, application state and data volumes, logs and metrics, deployment artifacts, and instances. 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 14 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/ecs/
- **Service quotas:** https://docs.aws.amazon.com/AmazonECS/latest/developerguide/service-quotas.html
- **Service FAQs:** https://aws.amazon.com/ecs/faqs/

30.5 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.
- API Reference: Documents the Amazon ECS Query API.
- Amazon ECS section of AWS CLI Reference: Documents the Amazon ECS commands available in the AWS Command Line Interface (AWS CLI).
- 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.

31 Amazon Elastic File System (EFS)

31.1 Service Overview

Amazon Elastic File System (EFS) is designed to provide serverless, fully elastic file storage that lets you share file data without provisioning or managing storage capacity and performance. Using the AWS Management Console, you can create file systems that are accessible to Amazon Elastic Compute Cloud (EC2) instances, Amazon container services (Amazon Elastic Container Service [ECS], Amazon Elastic Kubernetes Service [EKS], and AWS Fargate), and AWS Lambda functions through a file system interface (using standard operating system file I/O APIs). Amazon EFS file systems can automatically scale from gigabytes to petabytes of data without needing to provision storage.

31.1.1 Features

- **Fully managed**: Amazon EFS is a fully managed service providing network file sharing (NFS) shared file system storage for Linux workloads. Amazon EFS makes it quick and efficient to create and configure file systems. You needn't worry about managing file servers or storage, updating hardware, configuring software, or performing backups. In seconds, you can create a fully managed file system using the AWS Management Console, the AWS Command Line Interface (CLI), or an AWS SDK.
- **Highly available and durable**: Amazon EFS offers two file system types that you can choose from based on your durability and availability needs. EFS Regional file

systems (recommended) offer the highest levels of durability and availability by storing data with and across multiple Availability Zones (AZs). EFS One Zone file systems store data redundantly within a single AZ, so data in these file systems will be unavailable and might be lost during a disaster or other fault within the AZ. Amazon EFS is designed for 99.999999999 percent (11 nines) of durability Amazon EFS storage classes.over a given year.

- **Fully elastic and highly scalable**: By default, Amazon EFS file system storage and throughput capacity is elastic, growing and shrinking automatically to suit your workloads needs. With elastic capacity, provisioning is unnecessary, and you're billed only for what you use. Amazon EFS file systems can deliver the throughput, IOPS, and low latency necessary for a broad range of workloads, independent of the file system size. Amazon EFS is designed to be highly scalable, growing to petabytes of storage, hundreds of thousands of I/O operations per second (IOPS), and tens of gigabytes per second of throughout, with massively parallel access from various AWS compute instances.
- **Throughput modes**: With the default EFS Elastic Throughput mode, your file system throughput automatically scales with your workload activity and you pay only for what you use. Use the Elastic Throughput mode if you're unsure of your application's peak throughput needs or if your application's throughput usage is very spiky. If you know your workload's peak throughput requirements and you expect your workload to consume a higher share of your application's peak throughput capacity, use the Provisioned Throughput mode.
- **Storage classes**: Amazon EFS offers three storage classes: EFS Standard, EFS Infrequent Access, and EFS Archive. Data that is frequently accessed tends to have higher performance needs, so EFS provides an SSD-powered EFS Standard class designed to deliver sub-millisecond latencies. For data that's infrequently accessed, you can use EFS's two cost-optimized storage classes that provide low double-digit millisecond latencies: EFS Infrequent Access (IA), designed for data accessed only a few times a quarter, and EFS Archive, designed for data accessed less than few times a year. EFS IA offers up to 95% lower cost than EFS Standard for infrequently-accessed data. EFS Archive is a further cost-optimized experience for colder data, offering up to 50% lower cost than EFS Infrequent Access.
- **Lifecycle management and Intelligent-Tiering**: By enabling EFS Lifecycle Management, you can automatically tier files between storage classes based on your access patterns. You can create a custom lifecycle management policy to transition files between storage classes, or use the default, recommended policy which will tier files from EFS Standard to EFS IA after 30 consecutive days without access and to EFS Archive after 90 consecutive days without access. You can also enable EFS Intelligent-Tiering to transition files from EFS IA and EFS Archive back to EFS Standard for subsequent faster, sub-millisecond access.
- **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 (Amazon EC2) instances, as well as AWS container and serverless compute services. Amazon EFS also simultaneously supports on-premises servers 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 via AWS Direct Connect or AWS VPN services.
- **Containers and serverless file storage**: In addition to traditional EC2 instances, Amazon EFS runs on AWS containers and serverless compute services that require shared storage for latency-sensitive and IOPS-heavy workloads at any scale. In a single step, Amazon EFS provides applications running on Amazon Elastic Container

Service (Amazon ECS), Amazon Elastic Kubernetes Service (EKS), AWS Fargate, and AWS Lambda with access to shared file systems for stateful workloads.

- **Amazon EFS Replication**: Amazon EFS Replication copies your file system data into another file system in the AWS Region of your choice. Organizations in regulated industries are often mandated by compliance requirements to store a secondary copy of their data, hundreds of miles away from the primary, to plan for Disaster Recovery (DR) events. Using EFS Replication, you can perform DR workflows such as failover and failback without requiring additional infrastructure or custom processes. EFS Replication transfers only incremental data to synchronize your file systems and is designed to provide a recovery point objective (RPO) and a recovery time objective (RTO) of minutes. You can use the Amazon EFS console, AWS Command Line Interface (CLI), AWS CloudFormation, and APIs to activate replication on an existing file system.
- **AWS Backup**: Amazon EFS Backups are powered by AWS Backup, which is a fully managed backup service that centrally manages and automates backups of your Amazon EFS file systems, removing the need for costly custom solutions and manual processes. AWS Backup goes beyond backing up EFS and centralizes the backup of data across other AWS services in the cloud and on premises. As applications move to the cloud, their data can become distributed across multiple services, making it difficult to manage and consolidate backup activity without creating custom scripts and manual processes. Using AWS Backup, you can centrally configure and audit AWS resources, automate backup scheduling, set retention policies, and monitor backup activity.
- **Security and compliance**: Control network access to your file systems using Amazon Virtual Private Cloud (VPC) security group rules, and application access to your file systems using AWS Identity and Access Management (IAM) policies and Amazon EFS Access Points. AWS is certified under numerous certification programs to help you achieve your compliance goals. See this list of compliance programs in scope for Amazon EFS for more information.
- **Encryption**: Amazon EFS provides a comprehensive encryption solution to help you secure both your stored data and data in flight. Data at rest is transparently encrypted using encryption keys managed by the AWS Key Management Service (KMS), eliminating the need to build and maintain a key management infrastructure. Encryption of data in transit uses industry-standard Transport Layer Security (TLS) to help you secure network traffic, without the need to modify your applications. Refer to the user documentation on encryption for more information about encrypting file system data.
- **AWS DataSync**: AWS DataSync is a managed data transfer service that makes it faster and simpler to move data between on-premises storage and Amazon EFS. Use DataSync to transfer active datasets over the internet or AWS Direct Connect up to 10 times faster than open-source tools, without the need to modify your applications. Use the service for one-time data migrations, ongoing workflows with periodic synchronization, or replication for data protection and recovery. DataSync automatically manages many tasks known to slow down migrations or burden IT operations, including infrastructure management, encryption, data validation, and data transfer orchestration.
- **AWS Transfer Family**: AWS Transfer Family provides fully managed support for file transfers directly into and out of Amazon EFS. With support for Secure File Transfer Protocol (SFTP), File Transfer Protocol over SSL (FTPS), and File Transfer Protocol (FTP), Transfer Family helps you seamlessly migrate your file transfer workflows to AWS.

31.1.2 Benefits

- **Quick and easy to create**: Create and configure shared file systems simply and quickly for AWS compute services—no provisioning, deploying, patching, or maintenance required.
- **Elastic and scalable**: Scale workloads on-demand to petabytes of storage and gigabytes per second of throughput out of the box.
- **Pay as you go**: Reduce TCO with automatic lifecycle management to cost-optimized Infrequent Access and Archive storage classes designed to lower costs up to 97%.
- **Secure and reliable**: Securely and reliably access your files with a fully managed file system designed for 99.999999999 percent (11 9s) durability and up to 99.99 percent (4 9s) of availability.

31.2 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.

31.3 Pricing Overview

Please see the AWS UK G Cloud 14 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/efs/
- **Service quotas:** https://docs.aws.amazon.com/efs/latest/ug/limits.html
- **Service FAQs:** https://aws.amazon.com/efs/faq/

31.5 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.

32 Amazon Elastic Kubernetes Service (EKS)

32.1 Service Overview

Amazon Elastic Kubernetes Service (Amazon EKS) is a managed Kubernetes service to run Kubernetes in the AWS cloud and on-premises data centers. In the cloud, 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. 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. On-premises, EKS provides a consistent, fully-supported Kubernetes solution with integrated tooling and simple deployment to AWS Outposts, virtual machines, or bare metal servers.

32.1.1 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.
- **Hybrid Deployments**: You can use EKS on AWS Outposts to run containerized applications requiring particularly low latencies to on-premises systems. AWS Outposts is a fully managed service that extends AWS infrastructure, AWS services, APIs, and tools to virtually any connected site. With EKS on Outposts, you can manage containers on-premises with the same ease as you manage your containers in the cloud. You can attach nodes running in AWS Local Zones or AWS Wavelength to EKS, giving you more choices for AWS-managed infrastructure at the edge. Amazon EKS Distro packages up the same open-source Kubernetes software distribution used in Amazon EKS on AWS for use on your own on-premises infrastructure. Manage EKS Distro clusters with your own tooling or with Amazon EKS Anywhere. With Amazon EKS Anywhere, you can easily create and operate Kubernetes clusters (building with the software in Amazon EKS Distro) on-premises, including on your own virtual machines (VMs) and bare metal servers. EKS Anywhere saves you the complexity of building and supporting your own tooling to manage Kubernetes clusters.
- **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.

- **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. This globally routable IPv6 address can be used to directly communicate with any IPv6 endpoint in your Amazon VPC, on-premises network, or the public internet. Further, EKS configures networking so that pods can still communicate with IPv4 based endpoints outside the cluster, enabling you to adopt the benefits of IPv6 using Kubernetes without requiring that all dependent services deployed across your organization are migrated to IPv6.
- **EKS Pod Identity**: EKS Pod Identity simplifies the work customers need to do to setup applications on EKS clusters to access AWS services. EKS cluster administrators get a simplified workflow for obtaining IAM credentials required for authenticating Kubernetes applications to access AWS resources such as S3 buckets, DynamoDB tables, and more. EKS Pod Identity makes it easy to use IAM roles across multiple clusters, and simplifies IAM policy management by supporting reuse of policies across IAM roles.
- **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. Cloud Map works with external-dns, an open-source Kubernetes connector that automatically propagates internal service locations to the Cloud Map service registry as Kubernetes services launch, and removes them upon termination. Kubernetes-based services are discoverable via Cloud Map, which provides a unified service registry for all container workloads.
- **Service Mesh**: Service mesh standardizes how every microservice within your application communicates, making it easy to build and run complex microservices applications.
- **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. EKS uses the Amazon VPC container network interface (CNI), allowing Kubernetes pods to receive IP addresses from the VPC. Amazon EKS works with the Project Calico network policy engine to provide fine-grained networking policies for your Kubernetes workloads. Use the Kubernetes network policy API to control access on a per-service basis..
- **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.
- **Compliance**: Amazon EKS is certified by multiple compliance programs for regulated and sensitive applications. Amazon EKS is compliant with SOC, PCI, ISO, FedRAMP-Moderate, IRAP, C5, K-ISMS, ENS High, OSPAR, HITRUST CSF, and is a HIPAA eligible service.

- **Container Image Signature Verification**: Amazon EKS is compatible with container image signature verification to enable deploying container workloads with approved images and artifacts. You can verify images (or any other OCI artifact like Software Bill of Materials) signed by AWS Signer, a fully managed signing solution, before deploying images in your Amazon EKS clusters. AWS supports open-source based image signing and verification solutions so you can easily sign artifacts stored in your registry, and verify them using open source policy-as-code or admission controllers.
- **Load balancing**: Amazon EKS supports using Elastic Load Balancing including Application Load Balancer (ALB), Network Load Balancer (NLB), and Classic Load Balancer. You can run standard Kubernetes cluster load balancing or any Kubernetes-supported ingress controller with your Amazon EKS cluster.
- **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.
- **Cost monitoring**: Amazon EKS simplifies the process of understanding the costs associated with your Kubernetes usage, both at the cluster level, and the individual application level.
- **Cost allocation tagging**: Amazon EKS automatically adds an AWS cost allocation tag to every EC2 instance that joins a cluster. This frees you from having to enforce a custom tagging policy across your organization to gain insights into cluster level costs. After you activate the EKS cluster name cost allocation tag in the AWS Billing Console, you can use AWS Cost and Usage reports track your EC2 costs associated with EKS clusters.
- **Kubecost**: Amazon EKS supports Kubecost which enables you to monitor costs broken down by Kubernetes resources including pods, nodes, namespaces, and labels. Kubernetes platform administrators and finance leaders can use Kubecost to visualize a breakdown of their Amazon EKS associated charges, allocate costs, and charge back to organizational units such as application teams. You can provide your internal teams and business units with transparent and accurate cost data based on their actual AWS bill and get customized recommendations for cost optimization based on their infrastructure environment and usage patterns within their clusters.
- **Logging**: Amazon EKS is integrated with AWS CloudTrail to provide visibility into EKS management operations, including audit history. You can use CloudTrail to view API calls to the Amazon EKS API. Amazon EKS also delivers Kubernetes control plane logs to Amazon CloudWatch for analysis, debugging, and auditing.
- **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 centers 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.

32.1.2 Benefits

- **AWS Service Integration**: Leverage built-in integrations with AWS services such as EC2, VPC, IAM, EBS and more.
- **Kubernetes cost savings**: Reduce costs with efficient compute resource provisioning and automatic Kubernetes application scaling.
- **Automated Kubernetes control plane security**: Ensure a more secure Kubernetes environment with security patches automatically applied to your cluster's control plane.

32.2 Backup/Restore and Disaster Recovery

This service allows for backup and recovery. The service can back up machine images and EBS snapshots in Amazon S3 and also use the open-source tool, Velero to create application backups. 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 14 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/eks/index.html
- **Service quotas:** https://docs.aws.amazon.com/general/latest/gr/eks.html
- **Service FAQs:** https://aws.amazon.com/eks/faqs/

32.5 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 how to use Amazon EKS features with other AWS services.
- Amazon EKS best practices guides: Describes best practices for using Amazon EKS.

33 Amazon Elastic MapReduce (EMR)

33.1 Service Overview

Amazon EMR is the industry-leading cloud big data solution for petabyte-scale data processing, interactive analytics, and machine learning using open-source frameworks such as Apache Spark, Apache Hive, and Presto.

33.1.1 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.
- **Provision clusters in minutes**: You can launch an EMR cluster in minutes. You don't need to worry about infrastructure provisioning, cluster setup, configuration, or tuning. EMR takes care of these tasks allowing you to focus your teams on developing differentiated big data applications.
- **Easily scale resources to meet business needs**: You can easily set scale out and scale in using EMR Managed Scaling policies and let your EMR cluster automatically manage the compute resources to meet your usage and performance needs. This improves cluster utilization and saves on costs. EMR Studio is an integrated development environment (IDE) that makes it easy for data scientists and data engineers to develop, visualize, and debug data engineering and data science applications written in R, Python, Scala, and PySpark. EMR Studio provides fully managed Jupyter Notebooks, and tools like Spark UI and YARN Timeline Service to simplify debugging.
- **Single-click high availability**: You can easily configure high availability for multi-master applications such as YARN, HDFS, Apache Spark, Apache HBase, and Apache Hive with a single click . When you enable multi-master support in EMR, EMR will configure these applications for High Availability, and in the event of failures, will automatically fail-over to a standby master so that your cluster is not disrupted, and place your master nodes in distinct racks to reduce risk of simultaneous failure. Hosts are monitored to detect failures, and when issues are detected, new hosts are provisioned and added to the cluster automatically.
- **EMR Managed Scaling**: Automatically resizes your cluster for best performance at the lowest possible cost. With EMR Managed Scaling you specify the minimum and maximum compute limits for your clusters and Amazon EMR automatically resizes them for best performance and resource utilization. EMR Managed Scaling continuously samples key metrics associated with the workloads running on clusters.
- **Easily reconfigure running clusters**: You can now modify the configuration of applications running on EMR clusters including Apache Hadoop, Apache Spark, Apache Hive, and Hue without re-starting the cluster. EMR Application Reconfiguration allows you to modify applications on the fly without needing to shut down or re-create the cluster. Amazon EMR will apply your new configurations and

gracefully restart the reconfigured application. Configurations can be applied through the Console, SDK, or CLI.

- **Serverless option**: Amazon EMR Serverless is a new option in Amazon EMR that makes it easy and cost-effective for data engineers and analysts to run applications built using open source big data frameworks such as Apache Spark, Hive or Presto, without having to tune, operate, optimize, secure or manage clusters.
- **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. Amazon EMR pricing is per-second with a one-minute minimum, and starts at $.015 per instance hour for a small instance ($131.40 per year). Amazon EMR makes it easy to use Spot instances so you can save both time and money. The EMR File System (EMRFS) allows EMR clusters to efficiently and securely use Amazon S3 as an object store for Hadoop. You can store your data in Amazon S3 and use multiple Amazon EMR clusters to process the same data set. Each cluster can be optimized for a particular workload, which can be more efficient than a single cluster serving multiple workloads with different requirements. You can use the AWS Glue Data Catalog as a managed metadata repository to store external table metadata for Apache Spark and Apache Hive.
- **Flexible data stores**: With Amazon EMR, you can leverage multiple data stores, including Amazon S3, the Hadoop Distributed File System (HDFS), and Amazon DynamoDB.
- **Support for open source applications**: Amazon EMR versioned releases allow you to 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 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. Integration with AWS Lake Formation allows you to define and manage fine-grained authorization policies in AWS Lake Formation to access databases, tables, and columns in AWS Glue Data Catalog. You can enforce the authorization policies on jobs submitted through Amazon EMR Notebooks and Apache Zeppelin for interactive EMR Spark workloads, and send auditing events to AWS CloudTrail. By enabling this integration, you also enable federated Single Sign-On to EMR Notebooks or Apache Zeppelin from enterprise identity systems compatible with Security Assertion Markup Language (SAML) 2.0.
- **Consistent Hybrid Experience**: AWS Outposts is a fully managed service that extends AWS infrastructure, AWS services, APIs, and tools to virtually any data center, co-location space, or on-premises facility for a truly consistent hybrid experience. Amazon EMR on AWS Outposts allows you to, allowing you to deploy and manage EMR clusters in your data center using the same AWS Management Console, Software Development Kit (SDK), and Command Line Interface (CLI) used for EMR.

33.1.2 Benefits

- **Big data analytics with scalable data pipelines**: Run big data applications and petabyte-scale data analytics faster, and at less than half the cost of on-premises solutions.
- **Open-source frameworks 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. Get up to 2X faster time-to-insights with performance-optimized and open-source API-compatible versions of Spark, Hive, and Presto. Easily develop, visualize, and debug your applications using EMR Notebooks and familiar open-source tools in EMR Studio.
- **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.

33.2 Backup/Restore and Disaster Recovery

This service allows for backup and recovery. The service can back up block volumes, take automated daily snapshots and store them using Amazon S3. 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 14 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/emr/
- **Service quotas:** https://docs.aws.amazon.com/general/latest/gr/emr.html
- **Service FAQs:** https://aws.amazon.com/emr/faqs/

33.5 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.
- API Reference: Describes the Amazon EMR API operations, including sample requests, responses, and errors for the supported web services protocols.

34 Amazon ElastiCache

34.1 Service Overview

Amazon ElastiCache is a web service that streamlines deployment and running of Memcached or Redis protocol-compliant caches in the cloud. ElastiCache improves application performance by allowing you to retrieve information from a fast, managed, in-memory system instead of relying on slower disk-based systems. The service simplifies and offloads the management, monitoring, and operation of in-memory environments, enabling

your engineering resources to focus on developing applications. Using ElastiCache, you can not only improve load and response times to user actions and queries but also reduce the cost associated with scaling web applications.

ElastiCache automates common administrative tasks required to operate a distributed in-memory key-value environment. Using ElastiCache, you can add a caching or in-memory layer to your application architecture in a matter of minutes with a few steps in the AWS Management Console. ElastiCache is designed to automatically maintain high availability, and provides a 99.99% availability Service Level Agreement (SLA). ElastiCache is protocol compliant with Memcached and Redis, so code, applications, and popular tools that you use today with your existing Memcached or Redis environments will seamlessly work with the service. There are no up-front investments required, and you pay only for the resources you use.

34.1.1 Features

- **Get started in under a minute**: With Amazon ElastiCache Serverless, you can create a highly available cache in under a minute without infrastructure provisioning or configuration. You can create an ElastiCache Serverless cache in a few steps by specifying a cache name in the AWS Management Console, AWS Software Development Kit (SDK), or AWS Command Line Interface (CLI).
- **No capacity planning**: ElastiCache Serverless removes the complex, time-consuming process of capacity planning by continuously monitoring a cache's compute, memory, and network use and instantly scaling vertically and horizontally to meet demand.
- **Pay-for-use billing model**: With our pay-for-use billing model, you do not have to worry about how quickly ElastiCache Serverless scales back capacity after you scale down a workload. You only pay for the data you store and the compute your application uses.
- **Quickly get started**: With ElastiCache Serverless, you can simply create a new serverless cache in under a minute using the console, AWS CLI, or AWS SDKs, without needing to manage infrastructure or capacity. If you are designing your own cluster, resources are preconfigured with the appropriate parameters and settings, and cache parameter groups enable granular control for fine-tuning of your Redis or Memcached environment.
- **Fully managed Redis and Memcached**: ElastiCache is a fully managed service. We automate time-consuming management tasks—such as capacity planning, software patch management, failure detection, and recovery—allowing you to pursue higher value application development. You get built-in access to the underlying in-memory database environment, making it straightforward to use ElastiCache with your existing Redis and Memcached tools and applications. With ElastiCache Serverless, all minor version updates, performance enhancements, and security patches are automatically applied with no configuration required and without application disruption.
- **Add a cache to your relational database**: You can use the console for Amazon Relational Database Service (Amazon RDS) and Amazon Aurora to create an ElastiCache cluster and attach it to your relational database., By doing so, you can accelerate application performance with faster reads and reduce costs.
- **Monitoring**: Amazon CloudWatch metrics provide insights to your ElastiCache resources at no additional charge. You can use the console to view over 40 key operational metrics for your instances, including compute, utilized memory, cache hit ratio, active connections, replication, and commands.

- **ElastiCache publishes messages about notable events**. ElastiCache Serverless events including new cache creation, deletions, and cache configuration updates are sent to Amazon EventBridge. When working with self-designed cache clusters, ElastiCache sends events to Amazon Simple Notification Service (Amazon SNS).
- **Tagging**: Benefit from the ability to tag your ElastiCache resources and Redis or Memcached snapshots for tracking and billing purposes. You can use AWS Cost Explorer to attribute costs to resources and Resource Groups to create and maintain collections of resources that share a common set of tags. To learn more about tagging your ElastiCache resources, refer to the documentation on ElastiCache for Redis tagging and ElastiCache for Memcached tagging.
- **JSON support**: ElastiCache for Redis provides built-in support for JavaScript Object Notation (JSON) documents in addition to data structures included in open source Redis. You can simplify application development by using the built-in commands designed and optimized for JSON documents. ElastiCache for Redis supports partial JSON document updates, as well as powerful searching and filtering using the JSONPath query language. JSON support is available when using ElastiCache for Redis version 6.2 and above.
- **Microsecond response times**: ElastiCache helps improve application performance and increase throughput for read-heavy workloads by removing the need to access disk-based databases for frequently accessed data. ElastiCache can scale to millions of operations per second with microsecond response times.
- **High throughput and latency**: ElastiCache for Redis version 7.1 delivers up to 100% more throughput and 50% lower P99 latency compared to ElastiCache for Redis version 7.0. You can achieve over 1 million requests per second per node, or 500 million requests per second per cluster, on r7g.4xlarge nodes or larger.
- **Scale clusters to match demand**: ElastiCache Serverless automatically and elastically scales to meet application performance demands. ElastiCache Serverless continuously monitors the memory, compute, and network bandwidth used on the cache by your application. It enables the cache to scale up in place, while also scaling out in parallel, to ensure the cache can support the traffic needs of your application. Learn more about scaling ElastiCache for Redis clusters and scaling ElastiCache for Memcached clusters.
- **Application auto scaling**: When designing your own cache, ElastiCache for Redis auto scaling gives you the ability to automatically increase or decrease the desired shards or replicas in your ElastiCache for Redis service to maintain steady, predictable performance at the lowest possible cost. ElastiCache for Redis uses CloudWatch metrics to determine when it is time to scale up or down.
- **High availability and multi-availability zones**: ElastiCache offers a 99.99% Service Level Agreement (SLA) when using a multi-availability zone (multi-AZ) or serverless configuration. ElastiCache Serverless automatically stores data redundantly across multiple AZs, with no user configuration required. When designing your own cache cluster, you can take advantage of multiple AWS AZs by creating replicas in multiple AZs to achieve high availability and scale read traffic. In the case of primary node loss, AWS automatically detects the failure and failover to a read replica to provide higher availability without the need for manual intervention.
- **Cross-Region disaster recovery with Global Datastore**: Global Datastore in ElastiCache for Redis provides fully managed, fast, reliable, and secure replication across AWS Regions. With Global Datastore, you can write to your ElastiCache for Redis cluster in one Region and have the data available to be read from two other cross-Region replica clusters to enable low-latency reads and disaster recovery across AWS Regions. In the unlikely event of Regional degradation, one of the

healthy cross-Region replica clusters can be promoted to become the primary cluster with full read and write capabilities.

- **Instance monitoring and repair**: ElastiCache continuously monitors the health of your instances. In the case a node experiences failure or a prolonged degradation in performance, ElastiCache will automatically restart or replace the node and associated processes.
- **Backup, restore, and export**: ElastiCache for Redis helps protect your data by creating snapshots of your clusters. You can set up automatic snapshots or initiate manual backups in a few steps in the console or through simple API calls. Using these snapshots, or any Redis RDB–compatible snapshot stored on Amazon Simple Storage Service (Amazon S3), you can then seed new ElastiCache for Redis clusters. You can also export your snapshots to an Amazon S3 bucket of your choice for disaster recovery, analysis, or cross-Region backup and restore.
- **Network isolation**: ElastiCache allows you to run your resources in Amazon Virtual Private Cloud (Amazon VPC). Amazon VPC allows you to isolate your ElastiCache resources by specifying the IP ranges you wish to use for your nodes, and to connect to other applications inside the same Amazon VPC.
- **Encryption in transit and at rest**: ElastiCache supports encryption in transit, which allows you to encrypt all communications between clients and your ElastiCache server, as well as within the ElastiCache service boundary. ElastiCache also supports encryption at rest, which allows you to encrypt your disk usage and backups in Amazon S3. Learn more about encryption and ElastiCache for Redis data security and Elasticache for Memcached data security.
- **Redis authentication and access control**: ElastiCache for Redis supports authentication with AWS Identity and Access Management (IAM) authentication using IAM identities, Redis AUTH, and role-based access control (RBAC). With IAM Authentication, you can authenticate a connection to ElastiCache for Redis using AWS IAM identities to strengthen your security model and simplify many administrative security tasks.
- **Compliance**: ElastiCache supports compliance with programs such as SOC 1, SOC 2, SOC 3, ISO, MTCS, C5, PCI, HIPAA, and FedRAMP.
- AWS PrivateLink: You can use AWS PrivateLink to privately access ElastiCache from your Amazon VPC. PrivateLink allows you to privately access ElastiCache API operations without an internet gateway, NAT device, VPN connection, or AWS Direct Connect connection.
- **Pay only for what you use**: With ElastiCache, you only pay for the resources you consume with no upfront costs or long-term commitments. You are charged for data stored and compute consumed with ElastiCache Serverless and hourly based on the number of nodes, node type, and pricing model selected when designing your own cluster.
- **Cost optimize your relational workloads**: You can optimize your relational database costs with in-memory caching using ElastiCache. You can save up to 55% in cost and gain up to 80x faster read performance using ElastiCache with Amazon RDS for MySQL (compared to Amazon RDS for MySQL alone).
- **Data tiering**: You can use data tiering for ElastiCache for Redis as a lower-cost way to scale your clusters up to hundreds of terabytes of capacity. Data tiering provides a price-performance option for Redis workloads by using lower-cost solid state drives (SSDs) in each cluster node in addition to storing data in memory.
- **Reserved nodes**: ElastiCache reserved nodes provide you with a significant discount over on-demand usage when you commit to a one-year or three-year term. With reserved nodes, you can make a no upfront, partial upfront, or all upfront

payment to create a reservation to run your node in a specific Region. These reservations are available in one-year or three-year increments and offer a significant discount off of the ongoing hourly usage charge.

34.1.2 Benefits

- **Fast response times**: Realize microsecond response times across hundreds of millions of operations per second.
- **Cost-optimized performance**: Add a cache for frequently read data to maximize resources and lower total cost of ownership.
- **No capacity management**: Create a highly available, serverless cache in under a minute and instantly scale to meet demand.
- **Business continuity**: Achieve a 99.99% SLA with multi–Availability Zone (AZ) deployments, and bolster disaster recovery with cross-Region replication.
- **Redis- & Memcached-compatible**: Build applications quickly using popular open-source Redis & Memcached APIs.

34.2 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.

34.3 Pricing Overview

Please see the AWS UK G Cloud 14 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/elasticache/
- **Service quotas:** https://docs.aws.amazon.com/general/latest/gr/elasticache-service.html

34.5 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.
- API Reference: Describes all the API operations for Amazon ElastiCache in detail. Also provides sample requests, responses, and errors for the supported web services protocols.
- ElastiCache section of the AWS CLI Reference: Describes the command line interface for Amazon ElastiCache in detail. Provides basic syntax, options, and usage examples for each command.

35 Amazon EventBridge

35.1 Service Overview

Amazon EventBridge is a serverless service that uses events to connect application components together, making it easier for you to build scalable event-driven applications. Event-driven architecture is a style of building loosely-coupled software systems that work together by emitting and responding to events. Event-driven architecture can help you boost agility and build reliable, scalable applications.

Use EventBridge to route events from sources such as home-grown applications, AWS services, and third-party software to consumer applications across your organization. EventBridge provides simple and consistent ways to ingest, filter, transform, and deliver events so you can build applications quickly..

35.1.1 Features

- **Amazon EventBridge Pipes**: EventBridge Pipes is a serverless integration resource that helps you build point-to-point integrations by providing a simpler and consistent way to integrate event producers with event consumers without writing additional code. You can integrate with over 14 AWS services.
- **EventBridge Scheduler**: Amazon EventBridge Scheduler is a serverless scheduler that makes it easy to create, execute, and manage scheduled tasks at scale. You can reliably schedule millions of events and tasks that can invoke over 200 AWS services as a target. EventBridge Scheduler provides a central location to manage all of your scheduled jobs across the cloud.
- **Improve availability of event-driven applications with Global Endpoints**: Global endpoints is an easier and reliable way that you can improve the availability of your event-driven applications on AWS. Global endpoint makes it easier for you to build robust and reliable applications by automatically failing over their event ingestion to a secondary Region during service disruptions, without the need for manual intervention.
- **Archive, replay, and reprocess past events with Event Replay**: With Event Replay, you can reprocess past events back to an event bus or a specific EventBridge rule. This feature helps you debug your applications quickly, extend them by hydrating targets with historic events, and recover from errors.
- **Easily keep track of events with Schema Registry**: The EventBridge Schema Registry stores event schemas in a registry that other developers can more easily search and access in your organization. Thus, you don’t have to find events and their structure manually. The registry also helps you generate code bindings for programming languages such as Java, Python, or TypeScript directly in your integrated development environment (IDE) so that the event can be used as an object in your code.
- **Fully managed and scalable event bus**: EventBridge is a serverless, fully managed, and scalable event bus. It helps you build reliable and scalable applications at a faster rate. Amazon EventBridge Event Bus enables you to easily communicate between applications and services via events. It reduces the operational overhead without having to provision, patch, and manage servers.
- **Effortlessly integrate with third-party SaaS applications**: Your AWS applications can take action based on events that SaaS applications generate. EventBridge is natively integrated with SaaS applications from many providers including Shopify, BuildKite, Datadog, OneLogin, PagerDuty, Saviynt, Segment.io, 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; they will appear on your event bus.

- **Over 200 built-in event sources and targets**: EventBridge is directly integrated with over 200 event sources and over 20 targets, including AWS Lambda, Amazon Simple Queue Service (SQS), Amazon Simple Notification Service (SNS), AWS Step Functions, Amazon Kinesis Data Streams, and Amazon Kinesis Data Firehose, with additional sources and targets planned. All mutating API calls (all calls except Describe*, List*, and Get*) across AWS services generate events through AWS CloudTrail.
- **Increase developer agility with decoupled event publishers and subscribers**: EventBridge makes it easier for you to build event-driven application architectures. Applications or microservices to: Publish events to the event bus without awareness of subscribers; Subscribe to events without awareness of the publisher; Send events from your own applications to an event bus via the service's PutEvents API; Receive events through any of the many supported AWS target services; Decoupling allows teams to work independently, leading to faster development and improved agility.
- **Only process relevant events with Event Filtering**: Event Filtering reduces your costs by only processing events that your application needs. 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. Rules allow you to transform the structure of your event before passing it on to a target. 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.
- **Automatic response to operational changes on AWS services**: 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 helps you respond quickly to operational changes and take corrective action. You need to 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 a Lambda function to remediate an issue or notify an Amazon Simple Service Notification Service (SNS) topic to alert an operator.
- **Event Bus Monitoring, Auditing, and Reports via Amazon CloudWatch**: 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 analyze events that are initiated in your environment.
- **Build with security and compliance in mind to protect your data**: 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. EventBridge is GDPR, SOC, ISO, DoD CC SRG, and FedRamp compliant and is also HIPAA eligible.
- **Flexible pricing with pay per use**: Amazon EventBridge offers flexible pricing, you only pay for events published to your event bus, events ingested for your Schema Discovery, Event Replay, and API Destinations making it cost-effective solution for all organizations. There are no additional charges for rules or event delivery. There are no minimum fees or upfront commitments. All state change events published by AWS services are free.

- **Reliable event delivery with automatic retries**: 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. EventBridge also provides a 99.99% availability service level agreement (SLA), verifying your applications can access the service reliably.

35.1.2 Benefits

- **Increase developer agility**: Remove the need to coordinate across service teams with decoupled microservices using AWS, SaaS apps, or your own custom apps.
- **Monitor and audit applications**: Monitor and audit your AWS environments, and respond to operational changes in your applications in real time to prevent infrastructure vulnerabilities.
- **Extend functionality with SaaS integrations**: Connect your applications to other software-as-a-service (SaaS) SaaS applications by sending a custom event to EventBridge, and then send it through API Destinations to Zendesk CRM.
- **Scheduling in your Applications**: Use EventBridge Scheduler in your applications and platforms to provide scheduling services to your customers with reminders, delayed actions, or prompts to continue where they left off.

35.2 Backup/Restore and Disaster Recovery

This service allows for backup and recovery. The service can retrieve past events from archives using filters for replay as needed. Users can recover event data stored in CloudWatch Logs. Users control this via a manual or scheduled API call. Users schedule and recover backups through a web interface.

35.3 Pricing Overview

Please see the AWS UK G Cloud 14 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/eventbridge/
- **Service quotas:** https://docs.aws.amazon.com/eventbridge/latest/userguide/eb-quota.html
- **Service FAQs:** https://aws.amazon.com/eventbridge/faqs/

35.5 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.
- EventBridge section of AWS CLI Reference: Documents the Amazon EventBridge commands available in the AWS Command Line Interface (AWS CLI).

36 Amazon File Cache

36.1 Service Overview

Amazon File Cache provides a high-speed cache on AWS that makes it easier to process file data, regardless of where it's stored. Amazon File Cache serves as temporary, high-performance storage for data on premises or on AWS. The service allows you to make dispersed datasets available to file-based applications on AWS with a unified view and high speeds.

36.1.1 Features

- **Low-latency and high throughput data access for cached datasets**: Access data in Amazon File Cache from Amazon Elastic Compute Cloud (EC2) instances and Amazon Elastic Container Service (ECS) and Amazon Elastic Kubernetes Service (EKS) container instances with sub-millisecond average first-byte latencies, up to hundreds of GB/s of throughput, and millions of IOPS.
- **Cache on-premises datasets**: You can create a high-speed cache that's linked to one or more on-premises NFS file systems with just a few clicks in the AWS Management Console.
- **Cache AWS datasets**: You can create a high-speed cache that's linked to one or more AWS file systems or Amazon S3 buckets with just a few steps in the AWS Management Console. File Cache transparently presents data from your AWS file systems or Amazon S3 buckets as a unified set of files and directories and allows you to write results back to your datasets.
- **Reduce administrative burden**: You can create a high-performance cache with just a few steps in the AWS Management Console, command line interface (CLI), or API. With Amazon File Cache, you don't have to worry about managing the cache hardware and software or tuning performance—these time-consuming administrative tasks are automated as a benefit of this fully managed service.
- **Automate capacity management**: Amazon File Cache monitors free space on the cache and automatically releases less frequently used data when your cache begins to fill up so you can run your workload without worrying about running out of storage.
- **Network isolation**: Access your cache from endpoints in your Amazon Virtual Private Cloud (VPC) that isolates your cache in your own virtual network. You can configure security group rules and control network access to your Amazon File Cache resources.
- **Resource-level permissions**: Amazon File Cache is integrated with AWS Identity and Access Management (IAM). This integration allows you to control the actions that your AWS IAM users and groups can take to manage your caches, such as creating and deleting caches. You can also tag your Amazon File Cache resources and control the actions that IAM users and groups can take based on those tags.

36.1.2 Benefits

- **Accelerate cloud bursting**: Accelerate cloud bursting workloads using cache storage designed to deliver sub-millisecond latencies, up to hundreds of GB/s of throughput, and millions of IOPS.
- **For on-premises and in-cloud data sources**: Simplify access to multiple on-premises and in-cloud data sources with a unified view in a single namespace.
- **Amazon S3 integration**: Unlock insights from data in Amazon S3 object storage using file-based applications without changing your workflow.

36.2 Backup/Restore and Disaster Recovery

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

36.3 Pricing Overview

Please see the AWS UK G Cloud 14 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/fsx/
- **Service quotas:** https://docs.aws.amazon.com/fsx/latest/FileCacheGuide/limits.html
- **Service FAQs:** https://aws.amazon.com/filecache/faqs/

36.5 Technical Requirements

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

- Amazon File Cache User Guide: Describes key concepts of Amazon File Cache and provides instructions for launching an Amazon File Cache resource and using the features of Amazon File Cache.

37 Amazon FinSpace

37.1 Service Overview

Amazon FinSpace is a fully managed data management and analytics service that makes it easy to store, catalog, and prepare financial industry data at scale, which reduces the time it takes for financial services industry (FSI) customers to find and access all types of financial data for analysis from months to minutes.

Financial services organizations analyze 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. When performing analysis, it can take months to find the right data, get permissions to access the data in a compliant way, and prepare it for analysis.

37.1.1 Features

- **Secure and scalable data management for kdb**: Easily store kdb database files, encrypted with a KMS key you provide. Use of cloud object store technology seamlessly scales the kdb database as data volumes grow. Integrated, and configurable, high performance caching accelerates important queries, while managing the cost of what data is cached. Built-in database versioning allows you to travel back in time to query your data as it existed at a previous point in time.
- **Launch managed kdb clusters integrated with AWS services**: Using simple APIs, customers can launch and configure kdb clusters. Managed kdb Insights eliminates much of the operational overhead and provides pay as you go compute for kdb Insights applications. Managed kdb Insights is integrated with Terraform, to enable integration with CI/CD workflows, and AWS services, such as CloudWatch, CloudTrail, and IAM Permissions.

- **Autoscaling and High Availability to keep up with business demand**: Managed kdb Insights includes configurable autoscaling to ensure that kdb applications can keep up with the most volatile market conditions. You can also configure Multi-AZ to ensure kdb applications are highly available during the more important business hours.
- **Supports existing kdb configurations to simplify and accelerate migration**: Configure kdb ticker plants, real time clusters, historical clusters, and gateways to match existing environments to easily migrate existing kdb workloads and Q code. In addition to configuration as code through programmatic AWS APIs, customers can set up, access and manage the kdb environment (database and clusters) via the familiar interface of the FinSpace console.

37.1.2 Benefits

- **Simplifies the data handling process**: FinSpace simplifies the process of discovering data, gaining permissions to access data, and transforming data to be ready for analysis—thus saving months of prep work done by financial services customers today.
- **Migrate kdb applications to AWS**: Quickly gain benefits of cloud scale and availability by lifting and shifting your kdb applications to fully managed service on AWS.
- **Rapidly deploy new kdb infrastructure**: Reduce time to deploy new kdb applications, with autoscaling and high availability, from weeks to hours.
- **Test trading algorithms on petabytes of historical data**: Use on-demand compute to validate trading strategies against petabytes of historical market data (i.e. backtesting).

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

37.5 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.

- Data API Reference: Describes all the API operations for data resources in Amazon FinSpace.
- Management API Reference: Describes all the API operations for managing an Amazon FinSpace environment.

38 Amazon Fraud Detector

38.1 Service Overview

Amazon Fraud Detector is a fully managed service enabling customers to identify potentially fraudulent activities and catch more online fraud faster.

38.1.1 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.

38.1.2 Benefits

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

38.2 Backup/Restore and Disaster Recovery

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

38.3 Pricing Overview

Please see the AWS UK G Cloud 14 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/frauddetector/
- **Service quotas:** https://docs.aws.amazon.com/frauddetector/latest/ug/limits.html
- **Service FAQs:** https://aws.amazon.com/fraud-detector/faqs/

38.5 Technical Requirements

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

- Overview: Understand what Amazon Fraud Detector is, how it works, and how to use it for your business use case.
- Set up for Amazon Fraud Detector: Before you get started with Amazon Fraud Detector, set up the interfaces you plan to use.
- Start using Amazon Fraud Detector: Use the tutorial to get started with Amazon Fraud Detector using the AWS Management Console or the AWS SDK for Python.

39 Amazon FSx for Lustre

39.1 Service Overview

Amazon FSx for Lustre makes it easy and cost effective to launch, run, and scale the world's most popular high-performance file system.

The open source Lustre file system is designed for applications that require fast storage – where you want your storage to keep up with your compute. Lustre 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.

As a fully managed service, Amazon FSx brings Lustre to the masses, allowing you to use it for any workload where storage speed matters. Amazon FSx eliminates the traditional complexity of setting up and managing high-performance Lustre file systems, allowing you in minutes to spin up, run, and scale a battle-tested high-performance file system. It also provides multiple deployment options so you can optimize cost for your needs.

39.1.1 Features

- **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.
- **Eliminate administrative burden and scale capacity and performance 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.
- **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 can 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.
- **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.
- **Release inactive data from your file system**: To optimize your available storage capacity, you can release inactive data from your file system after the files are exported to Amazon S3. When a file is released, the file data is removed from the file system (and retained on S3) and the metadata remains on your file system. If a user or application accesses a released file, the data is automatically and transparently loaded back onto your file system from your S3 bucket.

39.1.2 Benefits

- **Accelerate compute workloads**: Accelerate compute workloads with shared storage that provides sub-millisecond latencies, up to hundreds of GBs/s of throughput, and millions of IOPS.

- **High-performance file system**: Access and process Amazon S3 data from a high-performance file system by linking your file systems to S3 buckets.
- **Range of deployment options**: Optimize price and performance with a range of deployment options, including storage type, performance tier, and replication level.

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. The service supports point-in-time (PITR) recovery of data. Users control this via a manual or scheduled API call. Users schedule and recover backups through a web interface.

39.3 Pricing Overview

Please see the AWS UK G Cloud 14 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 NetApp ONTAP

40.1 Service Overview

Amazon FSx for NetApp ONTAP provides fully managed shared storage in the AWS Cloud with the popular data access and management capabilities of ONTAP.

40.1.1 Features

- Support for petabyte-scale datasets in a single namespace
- Up to tens of gigabytes per second (GBps) of throughput per file system
- Multi-protocol access to data using the Network File System (NFS), Server Message Block (SMB), and Internet Small Computer Systems Interface (iSCSI) protocols
- Highly available and durable Multi-AZ and Single-AZ deployment options
- Automatic data-tiering that reduces storage costs by automatically transitioning infrequently accessed data to a lower-cost storage tier based on your access patterns

- Data compression, deduplication, and compaction to reduce your storage consumption
- Support for NetApp's SnapMirror replication feature
- Support for NetApp's on-premises caching solutions: NetApp Global File Cache and FlexCache
- Support for access and management using native AWS or NetApp tools and API operations
  - AWS Management Console, AWS Command Line Interface (AWS CLI), and SDKs
  - NetApp ONTAP CLI, REST API, and BlueXP
- Support for the following data protection and security features:
  - Encryption of file system data and backups at rest using AWS KMS keys
  - Encryption of data in transit using SMB Kerberos session keys
  - On-demand antivirus scanning
  - Authentication and authorization using Microsoft Active Directory
  - File access auditing
  - NetApp's SnapLock feature with support for both Compliance and Enterprise volumes

40.1.2 Benefits

- **Migrate apps and data with ease**: Migrate and build applications quicker with ONTAP's data access and data management capabilities on AWS.
- **Enable multiprotocol access**: Make your data available to a broad set of workloads and users via industry-standard NFS, SMB, and iSCSI protocols.
- **Automate storage scaling and tiering**: Support your growing data sets with storage capacity that grows and shrinks automatically.
- **Increase performance**: Take advantage of SSD performance at a fraction of the cost with built-in storage efficiency and tiering technologies.

40.2 Backup/Restore and Disaster Recovery

FSx for ONTAP provides highly available and durable storage with fully managed backups and support for cross-Region disaster recovery. For customers who use NetApp ONTAP on-premises, FSx for ONTAP is an ideal solution to migrate, back up, or burst your file-based applications from on-premises to AWS without the need to change your application code or how you manage your data.

40.3 Pricing Overview

Please see the AWS UK G Cloud 14 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/latest/ONTAPGuide/what-is-fsx-ontap.html
- **Service quotas:** https://docs.aws.amazon.com/fsx/latest/ONTAPGuide/limits.html
- **Service FAQs:** https://aws.amazon.com/fsx/netapp-ontap/faqs/

40.5 Technical Requirements

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

41 Amazon FSx for OpenZFS

41.1 Service Overview

Amazon FSx for OpenZFS is a fully managed file storage service that lets you launch, run, and scale fully managed file systems built on the open-source OpenZFS file system. FSx for OpenZFS makes it easy to migrate your on-premises file servers without changing your applications or how you manage data, and to build new high-performance, data-intensive applications on the cloud.

41.1.1 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 as low as 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 delivers up to 21 gigabytes per second (GB/s) of throughput and over 1 million IOPs for frequently-accessed cached data. For data accessed from persistent disk storage, FSx for OpenZFS file systems delivers up to 10 GB/s and up to 350,000 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 customize your file system for your specific workload needs. You can scale your storage capacity provisioned throughput and SSD IOPs 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 Zstandard and LZ4 compression technologies, which can help you reduce your storage costs for high-performance FSx for OpenZFS file systems and backup storage.
- **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 from AWS compute instances and containers**: Amazon FSx for OpenZFS file systems supports thousands of simultaneous clients running in Amazon Elastic Compute Cloud (Amazon EC2), Amazon Elastic Container Service (Amazon ECS), Amazon Elastic Kubernetes Service (Amazon EKS), VMware Cloud on AWS, Amazon WorkSpaces, and Amazon AppStream 2.0 instances.
- **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, in-place data cloning, and on-demand data replication across file systems, all available natively via the FSx console or API.
- **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.
- **Integrated with AWS services**: With Amazon FSx for OpenZFS, you can easily control who can administer your file systems, volumes, and backups—such as create, update, and delete—using AWS Identity and Access Management (IAM). You can monitor storage and performance metrics using Amazon CloudWatch, monitor and secure API calls using AWS CloudTrail and IAM, and manage file systems using AWS CloudFormation.
- **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.
- **Multi-AZ deployments**: Amazon FSx offers a multi-Availability Zone (Multi-AZ) deployment option, designed to provide continuous availability to data in the event of a file server failure, connectivity loss, or even an AZ-wide disruption.
- **Hot standby file systems with on-demand replication**: Amazon FSx for OpenZFS supports on-demand data replication across file systems. By scheduling regular, incremental replication to a second file system, you can implement an active–passive disaster recovery solution with an RTO as low as 5 minutes. Because on-demand replication only transfers incremental changes, you can achieve an RPO as low as 15 minutes for workloads driving up to 2 GB/s of data churn.
- **Security**: 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. 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.

41.1.2 Benefits

- **Easily migrate Linux file servers**: Easily migrate Linux file servers to AWS with file storage that matches or exceeds capabilities you rely on today and is accessible via the NFS protocol.
- **Simplify building and testing applications**: Simplify building and testing applications using OpenZFS instant data snapshots, data cloning, and on-demand data replication across file systems.
- **Accelerate workloads**: Accelerate workloads with high-performance storage delivering over 1 million IOPs and latencies of a few hundred microseconds.
- **Optimize costs**: Optimize costs with a few clicks by scaling throughput levels and enabling storage efficiency capabilities such as data compression.

41.2 Backup/Restore and Disaster Recovery

This service allows for backup and recovery. The service can take automatic daily backups of the file system, near-instant snapshots using the OpenZFS snapshot functionality, and cloning of volumes from snapshots. The service allows point-in-time recovery of data stored in Amazon S3. 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.

41.3 Pricing Overview

Please see the AWS UK G Cloud 14 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/index.html
- **Service quotas:** https://docs.aws.amazon.com/fsx/latest/OpenZFSGuide/limits.html
- **Service FAQs:** https://aws.amazon.com/fsx/openzfs/faqs/

41.5 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 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.

42 Amazon FSx for Windows File Server

42.1 Service Overview

Amazon FSx for Windows File Server provides fully managed, highly reliable, and scalable file storage that is accessible over the industry-standard Service Message Block (SMB) protocol. It is built on Windows Server, delivering a wide range of administrative features such as user quotas, end-user file restore, and Microsoft Active Directory (AD) integration.

To support a wide spectrum of workloads, Amazon FSx provides high levels of throughput and IOPS, and consistent sub-millisecond latencies.

Amazon FSx is accessible from Windows, Linux, and MacOS compute instances and devices. Thousands of compute instances and devices can access a file system concurrently. Amazon FSx File Gateway provides low-latency, on-premises access to fully managed file shares in the cloud.

42.1.1 Features

- **High-speed, consistent file storage in the cloud**: Amazon FSx for Windows File Server is designed to deliver fast, predictable, and consistent performance by providing multiple GB/s of throughput and hundreds of thousands of IOPS per file system. It is built on the latest AWS compute, disk, and networking technologies, providing fast, consistent sub-millisecond latencies for file operations with even lower latencies for frequently accessed data.
- **Throughput and IOPS performance**: Amazon FSx for Windows File Server delivers up to 21 gigabytes per second (GB/s) of throughput and millions of IOPs for frequently-accessed cached data. For data accessed from persistent disk storage, FSx for Windows file systems delivers up to 12 GB/s and up to 350,000 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**: Amazon FSx for Windows File Server 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. And FSx for Windows supports SMB Multichannel, enabling you to get up to the full throughput and IOPS levels for your file system even for just a single client.
- **Simple, flexible administration**: With Amazon FSx for Windows File Server, 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. You also have access to a rich set of Windows Server administrative features including end-user file restore, user quotas, and Access Control Lists (ACLs).
- **Deeply integrated with AWS services**: Amazon FSx for Windows File Server integrates with your on-premises Microsoft Active Directory or AWS Microsoft Managed AD to let you control user access; with AWS CloudTrail to let you monitor and log administration actions; with Amazon CloudWatch to let you monitor file system storage and performance; with AWS CloudFormation to let you model, provision, and manage file systems efficiently; with AWS Backup to let you create policy-driven backup plans; with Amazon Kendra to let you index and search documents stored on your file systems; and with Amazon ECS to enable persistent, shared storage for containerized applications.
- **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.
- **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.
- **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.
- **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 Continously 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.
- **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.
- **Cross-region and cross-account backup compliance**: To provide additional layers of data protection and meet business continuity, disaster recovery, and compliance requirements, you can copy your Amazon FSx file system backups across AWS Regions, AWS accounts, or both.
- **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.

- **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), your level of SSD IOPS above the default 3 IOPS that are included for every GB of SSD storage (per IOPS-mo), and your throughput capacity (priced per MBps-month). You are billed hourly for your backup storage (priced per GB-month).
- **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.
- **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.
- **Flexible throughput and IOPS selection**: Amazon FSx for Windows File Server offers multiple throughput capacity levels that you can choose from, allowing you to cost-optimize for the performance your workloads require. You can also optionally provision higher levels of IOPS as needed, independently from the storage and throughput capacity of your file system, allowing you to pay only for the IOPS you need.

42.1.2 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.

42.2 Backup/Restore and Disaster Recovery

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. Users control this via a manual or scheduled API call. Users schedule and recover backups through a web interface.

42.3 Pricing Overview

Please see the AWS UK G Cloud 14 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/fsx/
- **Service quotas:** https://docs.aws.amazon.com/fsx/latest/WindowsGuide/limits.html
- **Service FAQs:** https://aws.amazon.com/fsx/windows/faqs/

42.5 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.

43 Amazon GuardDuty

43.1 Service Overview

Amazon GuardDuty is a security monitoring service that analyzes and processes Foundational data sources, such as AWS CloudTrail management events, AWS CloudTrail event logs, VPC flow logs (from Amazon EC2 instances), and DNS logs. It also processes Features such as Kubernetes audit logs, RDS login activity, S3 logs, EBS volumes, Runtime monitoring, and Lambda network activity logs. It uses threat intelligence feeds, such as lists of malicious IP addresses and domains, and machine learning to identify unexpected, potentially unauthorized, and malicious activity within your AWS environment. This can include issues like escalation of privileges, use of exposed credentials, or communication with malicious IP addresses, domains, presence of malware on your Amazon EC2 instances and container workloads, or discovery of unusual patterns of login events on your database.

43.1.1 Features

- **Accurate, account-level threat detection**: 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 geolocation at an atypical time of day. For programmatic AWS accounts, GuardDuty checks for unusual API calls, such as attempts to obscure account activity by disabling CloudTrail logging or taking snapshots of a database from a malicious IP address.

- **Continuous monitoring across AWS accounts without added cost and complexity**: GuardDuty continuously monitors and analyzes your AWS account and workload event data found in CloudTrail, VPC Flow Logs, and DNS logs. There is no additional security software or infrastructure to deploy and maintain for the foundational protections in GuardDuty. By associating your AWS accounts together, you can aggregate threat detection instead of working on an account-by-account basis. In addition, you do not have to collect, analyze, and correlate large volumes of AWS data from multiple accounts. Focus on responding quickly, keeping your organization secure, and continuing to scale and innovate on AWS.
- **Threat detections developed and optimized for the cloud**: GuardDuty helps you access built-in detection techniques developed and optimized for the cloud. AWS Security continuously maintains and improves these detection algorithms. The primary detection categories include:
  - **Reconnaissance:** This activity suggests reconnaissance by an attacker, such as unusual API activity, suspicious database login attempts, intra-VPC port scanning, unusual failed login request patterns, or unblocked port probing from a known bad IP.
  - **Instance compromise:** This activity indicates an instance compromise, such as cryptocurrency mining, backdoor command and control (C&C) activity, runtime activity for Amazon EC2 (Preview), malware using domain generation algorithms (DGA), outbound denial of service activity, unusually high network traffic volume, unusual network protocols, outbound instance communication with a known malicious IP, temporary Amazon EC2 credentials used by an external IP address, and data exfiltration using DNS.
  - **Account compromise:** Common patterns indicative of account compromise include API calls from an unusual geolocation or anonymizing proxy, attempts to disable AWS CloudTrail logging, changes that weaken the account password policy, unusual instance or infrastructure launches, infrastructure deployments in an unusual region, credential theft, suspicious database login activity, and API calls from known malicious IP addresses.
  - **Bucket compromise:** This activity indicates a bucket compromise, such as suspicious data access patterns indicating credential misuse, unusual Amazon S3 API activity from a remote host, unauthorized S3 access from known malicious IP addresses, and API calls to retrieve data in S3 buckets from a user with no prior history of accessing the bucket or invoked from an unusual location. GuardDuty continuously monitors and analyzes CloudTrail S3 data events (for example, GetObject, ListObjects, DeleteObject) to detect suspicious activity across all your S3 buckets.
  - **Malware detection:** GuardDuty begins a malware detection scan when it identifies suspicious behavior indicative of malicious software in EC2 instance or container workloads. GuardDuty generates temporary replicas of Amazon EBS volumes attached to such EC2 instance or container workloads and scans the volume replicas for trojans, worms, crypto miners, rootkits, bots, and more that might be used to compromise the workloads, repurpose resources for malicious use, and gain unauthorized access to data. GuardDuty Malware Protection generates contextualized findings that can validate the source of the suspicious behavior. These findings can be routed to the proper administrators and initiate automated remediation.

  - **Container compromise:** Activity identifying possible malicious or suspicious behavior in container workloads is detected by continuously monitoring and profiling EKS clusters by analyzing its EKS audit logs and container runtime activity in EKS or ECS.

- **Threat severity levels for efficient prioritization**: 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. An example would be a large amount of traffic returned to a remote host hiding behind the Tor network or activity that deviates from normally observed behavior. A High severity level indicates that the resource in question (for example, an EC2 instance or a set of IAM user credentials) is compromised and is actively being used for unauthorized purposes.
- **Threat response and remediation automation**: GuardDuty offers HTTPS APIs and command line interface (CLI) tools, as well as integration with Amazon EventBridge to support automated security responses to security findings. For example, you can automate the response workflow by using EventBridge as an event source to invoke a Lambda function.
- **Scalable threat detection**: GuardDuty is designed to automatically manage resource utilization based on the overall activity levels within your AWS accounts, workloads, and data. GuardDuty adds detection capacity only when necessary and reduces utilization when capacity is no longer needed. You now have a cost-effective architecture that maintains the security processing power that you need while minimizing expenses. You only pay for the detection capacity that you use, when you use it. GuardDuty gives you security at scale, no matter your size.
- **One-step deployment with no additional software or infrastructure to deploy and manage**: With one action in the AWS Management Console or a single API call, you can activate GuardDuty on a single account. With a few more steps in the console, you can activate GuardDuty across multiple accounts. GuardDuty supports multiple accounts through AWS Organizations integration as well as natively within GuardDuty. Once turned on, GuardDuty immediately starts analyzing continuous streams of account and network activity in near real time and at scale. There are no additional security software, sensors, or network appliances to deploy or manage. Threat intelligence is pre-integrated into the service and is continuously updated and maintained.
- **Broad, container-aware protection**: GuardDuty provides comprehensive protection for container workloads across your AWS compute estate that would otherwise be difficult and complex to achieve. Whether you're running workloads with server-level control on EC2 or serverless modern application workloads on ECS with Fargate, GuardDuty detects potentially malicious and suspicious activity, gives you container-level context with runtime monitoring, and helps you identify security coverage gaps in your container workloads across your AWS environment.

43.1.2 Benefits

- **Continuous monitoring**: Keep your accounts, workloads, and data secure by continuously monitoring for potential threats across your AWS environment.
- **ML-powered threat detection**: Expose threats quickly using anomaly detection, machine learning (ML), behavioral modeling, and threat intelligence feeds from AWS and leading third parties.
- **Respond faster to threats**: Mitigate threats sooner by initiating automated responses.

- **Scalable threat detection**: Quickly and easily scale threat detection across your AWS environment.
- **Protect a broad range of AWS compute workloads**: Safeguard your accounts, data, and resources across various AWS compute types, spanning Amazon Elastic Compute Cloud (Amazon EC2), serverless workloads, and container workloads—including those on AWS Fargate.

43.2 Backup/Restore and Disaster Recovery

This service allows for backup and recovery. Users can configure backup policies and audit controls. Users control this via a manual or scheduled API call. Users schedule and recover backups through a web interface.

43.3 Pricing Overview

Please see the AWS UK G Cloud 14 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/guardduty/
- **Service quotas:** https://docs.aws.amazon.com/guardduty/latest/ug/guardduty_limits.html
- **Service FAQs:** https://aws.amazon.com/guardduty/faqs/

43.5 Technical Requirements

Please refer to 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.
- Amazon GuardDuty section of the AWS CLI Reference: Provides syntax and examples for the AWS CLI commands for use with Amazon GuardDuty.

44 Amazon Inspector

44.1 Service Overview

Amazon Inspector is an automated vulnerability management service that continually scans Amazon Elastic Compute Cloud (EC2), AWS Lambda functions, and container images in Amazon ECR and within continuous integration and continuous delivery (CI/CD) tools, in near-real time for software vulnerabilities and unintended network exposure.

44.1.1 Features

- **Automated vulnerability management for compute workloads**: Amazon Inspector is a comprehensive vulnerability management service that spans various resources, such as Amazon EC2, Lambda functions, 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 facilitate data exfiltration.

- **Simplified one-click onboarding and integration with AWS Organizations**: Start Amazon Inspector across multiple accounts with one step 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 seamlessly 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 Amazon EC2 instances, Lambda functions, and container images in Amazon ECR. It promptly initiates scans for software vulnerabilities and unintended network exposure. All workloads are continually rescanned when a new CVE is published or when there are changes in the workloads, including installation of new software in an EC2 instance.
- **Integration with AWS Systems Manager Agent**: Amazon Inspector uses the widely deployed AWS Systems Manager Agent (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.
- **Agentless vulnerability assessments for Amazon EC2 (in preview)**: Amazon Inspector offers continuous monitoring of your Amazon EC2 instances for software vulnerabilities without installing an agent or additional software. Amazon Inspector takes a snapshot of the EBS volume to extract data about the system and configuration of the instances to perform vulnerability assessments. With this capability, you can expand your vulnerability assessment coverage across your EC2 infrastructure with Amazon Inspector agentless scanning for EC2 instances (preview) that do not have SSM Agents installed or configured.
- **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.
- **Amazon Inspector risk score for findings**: Amazon Inspector generates a highly contextualized Amazon 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.
- **Automatic closure of remediated findings**: Amazon Inspector automatically detects if a vulnerability has been patched or remediated. Once detected, it automatically changes the state of the finding to “Closed” without manual intervention.
- **Detailed coverage monitoring**: Amazon Inspector provides a comprehensive, near real-time overview of organization-wide environment coverage, so you can avoid gaps in coverage. It provides metrics and detailed information on accounts, as well as Amazon EC2 instances, Amazon ECR repositories, and container images that are actively being scanned by Amazon Inspector. Additionally, it highlights the resources not being actively monitored and provides guidance on how to include them.
- **Integration with AWS Security Hub and Amazon 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.
- **Integrating vulnerability mapping and generative AI powered remediation to layers in Lambda functions**: Amazon Inspector scans the custom proprietary application code within a Lambda function for code security vulnerabilities such as injection flaws, data leaks, weak cryptography, or missing encryption based on AWS security best practices. Upon detecting code vulnerabilities within the Lambda function or layer, Amazon Inspector generates actionable security findings that provide several details, such as security detector name, impacted code snippets, and

remediation suggestions to address vulnerabilities. Using generative AI and automated reasoning, Amazon Inspector provides in-context code patches for multiples classes of vulnerabilities, reducing the effort required to fix code vulnerabilities. By addressing vulnerabilities at the foundational layers, you can help improve security of all downstream Lambda functions.

- **Manage software bill of materials (SBOM) exports**: Amazon Inspector offers automated and centralized management of software bill of materials (SBOM) exports. It enables the easy export of a consolidated SBOM for all monitored resources to a preconfigured Amazon S3 bucket, supporting industry-standard formats. You can download the SBOM artifact, perform Amazon Athena queries, or create Amazon QuickSight dashboards to gain valuable insights and visualize trends.
- **Integration with developer tools**: Amazon Inspector integrates with developer tools like Jenkins and TeamCity for container image assessments. It allows developers to assess their container images within these CI/CD tools, pushing security earlier in the software development lifecycle. The findings are available in the CI/CD tool's dashboard, allowing you to take immediate automated actions in response to critical security issues, such as blocking builds or image pushes to container registries. Your CI/CD tools can be hosted anywhere, in AWS, on-premises, or hybrid clouds, providing consistency for developers to use a single solution across all your development pipelines.
- **Support for CIS Benchmark assessments**: Amazon Inspector supports the Center for Internet Security's CIS Benchmarks. You can run Amazon Inspector to perform on-demand and targeted assessments against OS-level CIS configuration benchmarks for Amazon EC2 instances across your AWS Organization. Amazon Inspector CIS assessments support both level 1 and 2 configuration benchmark checks across operating systems, including Amazon Linux 2, Windows 2019, and Windows 2022.

44.1.2 Benefits

- **Detect software vulnerabilities**: Detect software vulnerabilities and unintended network exposure in AWS workloads such as Amazon EC2, AWS Lambda functions, and container images in Amazon ECR and within continuous integration and continuous delivery (CI/CD) tools, in near-real time.
- **Manage SBOM exports centrally**: Incorporate security earlier in the development cycles and centrally manage software bill of materials (SBOM) exports for all monitored resources.
- **Prioritize remediation**: Use the Amazon Inspector risk score to prioritize remediation reducing mean time to remediate (MTTR).
- **Maximize vulnerability assessment coverage**: Seamlessly scan EC2 instances switching between agent-based and agentless scanning

44.2 Backup/Restore and Disaster Recovery

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

44.3 Pricing Overview

Please see the AWS UK G Cloud 14 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/inspector/
- **Service quotas:** https://docs.aws.amazon.com/general/latest/gr/inspector.html
- **Service FAQs:** https://aws.amazon.com/inspector/faqs/

44.5 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.

45 Amazon IVS

45.1 Service Overview

Amazon Interactive Video Service (Amazon IVS) is a managed live streaming solution, built on the same technology that powers Twitch, that makes low-latency or real-time live video available to viewers around the world. Simply send your live streams to Amazon IVS, and the service does everything you need to make them available to your viewers, letting you focus on building engaging and interactive experiences alongside live video. You can customize and enhance the audience experience through the player and broadcast SDKs, the Chat APIs, and timed metadata APIs, giving you the tools to build a more valuable relationship with your viewers on your own websites and applications.

45.1.1 Features

45.1.1.1 Real-Time Streaming

- **Set Up New Live Video Streams in Minutes**: Amazon IVS provides all the components needed for a real-time live video streaming solution. Hosts publish and viewers subscribe to live video stream using the broadcast SDK on application you can build for iOS Android, and web.
- **Optimized for Live Video Streaming**: Build applications for latency-sensitive use cases, helping your viewers stay connected and engaged with Real-time Streaming. Deliver live streams with a latency that can be under 300 milliseconds from host to viewer.
- **Build engaging audience experiences**: Create community relationships with your audience using scalable stream chat APIs to build chat rooms to accompany your live streaming video.
- **Multi-Platform Broadcast SDK, Send live video from iOS and Android devices**: Use iOS and Android applications to publish and subscribe to live video with the mobile broadcast SDK and from a browser using the web broadcast SDK. Websites and applications using the SDK can access the device camera, microphone, screen

sources, or use custom audio and video sources. Streams published from devices will use optimized encoder settings and network congestion handling.

- **Combine video and audio from multiple hosts into one low-latency stream**: Use a virtual resource called a stage to enable hosts to exchange audio and video in real time, then broadcast that stage to an Amazon IVS low-latency channel to reach a larger audience.
- **Pre-configured layered encoding**: Layered encoding, sometimes referred to as "Simulcast" is a standardized technique where multiple renditions, or layers, of the same video source are published and the rendition that best fits viewers needs based on available bandwidth is sent to their application.
- **High concurrency**: Deliver large-scale interactive events with support for up to 10,000 viewers and up to 12 hosts.
- **Cost-Effective**: Amazon IVS Real-Time Steaming provides simple, pay-as-you-go pricing based on participant hours, which is the duration of time each host or viewer is connected to a stage resource.

45.1.1.2 Low-latency Streaming

- **Set Up New Live Video Streams in Minutes**: Amazon IVS 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.
- **Ingest, transcode, and distribute over managed infrastructure**: All components from stream ingest and transcode to delivery and playback with the Amazon IVS player SDK are optimized to reduce latency. Actual latency for viewers is usually under five seconds and can be less than three seconds.
- **Build engaging audience experiences**: Create community relationships with your audience using scalable stream chat APIs to build chat rooms to accompany your live streaming 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.
- **Send live video from web, iOS, and Android devices**: Use iOS and Android applications to send live video to Amazon IVS with the mobile broadcast SDK and send direct from a browser using the web broadcast SDK. Websites and applications using the SDK can access 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.
- **Playback on web, iOS, and Android platforms**: 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.
- **Combine real-time video and audio from multiple hosts into one live stream**: Use a virtual resource called a stage to enable hosts to exchange audio and video in real time, then broadcast that stage to an Amazon IVS channel to reach a larger audience. Learn more.
- **Save live video outputs to Amazon S3 for archiving and on-demand playback**: 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.

- **Pre-configured adaptive bitrate transcoding**: With a standard or advance channels, 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.
- **Automatic scale for size of audience and number of channels**: 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.
- **Pay-as-you-go pricing**: 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).

45.1.2 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.

45.2 Backup/Restore and Disaster Recovery

This service allows for backup and recovery. Data streams are stored durably in Amazon S3 for later retrieval. Streaming can automatically failover to another instance without any data loss. Developers can take data stream snapshots for recovery.

45.3 Pricing Overview

Please see the AWS UK G Cloud 14 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/ivs/index.html
- **Service quotas:** https://docs.aws.amazon.com/ivs/latest/userguide/service-quotas.html

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

45.5 Technical Requirements

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

- Low-Latency Streaming User Guide: Introduces you to and helps you get started with Amazon IVS Low-Latency Streaming. Provides instructions on using various features with the console, API, and command-line interface. Includes information about using the Amazon IVS broadcast SDK on various platforms.
- Low-Latency Streaming API Reference: Describes in detail all Amazon IVS Low-Latency Streaming 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.
- Real-Time Streaming User Guide: Introduces you to and helps you get started with Amazon IVS Real-Time Streaming. Provides instructions on using various features with the console, API, and command-line interface. Includes information about using the Amazon IVS broadcast SDK on various platforms.
- Real-Time Streaming API Reference: Describes in detail all Amazon IVS Real-Time Streaming 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..

46 Amazon Kendra

46.1 Service Overview

Amazon Kendra is a highly accurate and easy-to-use enterprise search service that's powered by machine learning (ML). It allows developers to add search capabilities to their applications so their end users can discover information stored within the vast amount of content spread across their company. This includes data from manuals, research reports, FAQs, human resources (HR) documentation, and customer service guides, which may be found across various systems such as Amazon Simple Storage Service (S3), Microsoft SharePoint, Salesforce, ServiceNow, RDS databases, or Microsoft OneDrive. When you type a question, the service uses ML algorithms to understand the context and return the most relevant results, whether that means a precise answer or an entire document. For example, you can ask a question such as "How much is the cash reward on the corporate credit card?" and Amazon Kendra will map to the relevant documents and return a specific answer (such as "2%"). Kendra provides sample code so you can get started quickly and easily integrate highly accurate search into your new or existing applications.

46.1.1 Features

- **Generative AI**: Create secure, generative AI-powered conversational experiences for your users on top of your enterprise content. Amazon Kendra provides an optimized Kendra Retriever API that allows you to use Amazon Kendra's high-accuracy semantic ranker as an enterprise retriever for your Retrieval Augmented Generation (RAG) workflow. The Kendra Retriever API finds and retrieves passages from your enterprise content that are most semantically relevant to the user's question and have optimized granularity to maximize the quality of your RAG payload - without the need for you to have expertise in accurate semantic retrieval. These optimized

passages can then be sent, along with the user's question, to an LLM to get a generative response.

- **Intelligent search**: Amazon Kendra uses ML to deliver more-relevant answers from unstructured data. Search for general keywords (such as "health benefits") or ask natural language questions ("How long is maternity leave?") and Amazon Kendra will use reading comprehension to give specific answers ("14 weeks").
- **Incremental learning**: Amazon Kendra uses ML to continuously optimize search results based on end-user search patterns and feedback. For example, when users search for "How do I change my health benefits?" multiple human resources (HR) benefit documents will compete for a top spot. 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. It 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.
- **Connectors**: Using connectors with Amazon Kendra is quicker and easier—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. Amazon Kendra offers native connectors for popular data sources such as Amazon Simple Storage Service (S3), Microsoft SharePoint, Salesforce, ServiceNow, Google Drive, Confluence, and many more. If a native connector is not available, Amazon Kendra offers a custom data source connector as well as a host of partner-supported connectors.
- **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.
- **Experience Builder**: You can now deploy a fully functional and customizable search experience with Amazon Kendra in a few steps, without any coding or ML experience. Experience Builder delivers an intuitive visual workflow to quickly build, customize, and launch your search application securely on the cloud.
- **Search Analytics Dashboard**: Amazon Kendra Search Analytics Dashboard helps you to better understand quality and usability metrics across your search applications. The dashboard helps administrators and content creators understand how easily the end users are finding relevant search results, the quality of the search results, and gaps in the content.
- **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.
- **Query autocompletion**: Amazon Kendra includes the functionality to autocomplete an end user's search query. Query autocompletion not only helps you reduce typing by about 25%, but it also helps guide you toward more precise and commonly asked questions. These questions typically result in answers that are more relevant and useful.

46.1.2 Benefits

- **Unified search experience**: Quickly implement a unified search experience across multiple structured and unstructured content repositories.
- **Highly accurate answers**: Use natural language processing (NLP) to get highly accurate answers without the need for machine learning (ML) expertise.

- **Customized search results**: Fine-tune your search results based on content attributes, freshness, user behavior, and more.
- **Fast**: Deliver ML-powered instant answers, FAQs, and document ranking as a fully managed service.

46.2 Backup/Restore and Disaster Recovery

This service allows for backup and recovery. The service encrypts your data both in transit and at rest. You can create snapshots of your Amazon Kendra index. Point-in-time recovery (PITR) restores index to state within the last 5 minutes. The service configures to multiple availability zones by default for high availability. Users can control backups via scheduled API call. Users schedule and recover backups through a web interface.

46.3 Pricing Overview

Please see the AWS UK G Cloud 14 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/kendra/index.html
- **Service quotas:** https://docs.aws.amazon.com/kendra/latest/dg/quotas.html
- **Service FAQs:** https://aws.amazon.com/kendra/faqs/

46.5 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.

47 Amazon Keyspaces (for Apache Cassandra)

47.1 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.

47.1.1 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. Learn more about supported CQL APIs.

- **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.
- **Multi-Region Replication**: With Multi-Region Replication, you can replicate your data with automated, fully-managed, active-active replication across the AWS Regions of your choice. Multi-Region Replication uses storage-based asynchronous replication and data is typically propagated across the Regions in your keyspace within a second. With Multi-Region Replication you can improve both availability and resiliency while also benefiting from low latency local reads and writes for global applications.
- **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.
- **Encryption at rest and in transit**: Amazon Keyspaces encrypts all customer data at rest by default. Encryption at rest enhances the security of your data by using encryption keys stored in AWS Key Management Service. You can choose to encrypt your data at rest with an AWS owned customer master key (default) or with a

customer managed customer master key, giving you full control over how your data is encrypted.

- **Access management**: Amazon Keyspaces is integrated with AWS Identity and Access Management to help you manage access to your tables and data.
- **Secure network connectivity**: Amazon Keyspaces supports secure networking by using TLS. You can use AWS PrivateLink to provide secure connectivity between your resources in Amazon Keyspaces and Amazon Virtual Private Cloud (Amazon VPC).
- **Availability SLA**: Amazon Keyspaces offers an availability SLA of 99.99% within an AWS Region, providing you with a strong availability commitment and no scheduled downtime.

47.1.2 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.

47.2 Backup/Restore and Disaster 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 automatically with per-second granularity, and you can restore that table to any second in the preceding 35 days. You can enable PITR or initiate backup-and-restore operations using the AWS Management Console or a single API call.

47.3 Pricing Overview

Please see the AWS UK G Cloud 14 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/keyspaces/index.html

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

47.5 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.
- API Reference: Describes supported API operations in the AWS SDK for Amazon Keyspaces in detail. Also provides sample requests, responses, and errors for the supported web services protocols.
- Amazon Keyspaces Code Examples Repository: Contains code examples and other types of examples to help accelerate your development of Amazon Keyspaces applications. The examples demonstrate common service implementations and infrastructure patterns using various programming languages and supported technologies.

48 Amazon Kinesis Data Streams

48.1 Service Overview

With Kinesis Data Streams, you can build custom applications that process or analyze streaming data for specialized needs. You can add various types of data such as clickstreams, application logs, and social media to a Kinesis data stream from hundreds of thousands of sources. Within seconds, the data will be available for your applications to read and process from the stream.

48.1.1 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 Managed Service for Apache Flink, 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 Managed Service for Apache Flink, and Amazon Kinesis Data Firehose to build complete applications quickly.

48.1.2 Benefits

- **Easy to provision**: Kinesis Data Streams manages the infrastructure, storage, networking, and configuration needed to stream your data at the level of your data throughput. You don't have to worry about provisioning, deployment, or ongoing maintenance of hardware, software, or other services for your data streams. By default, Kinesis Data Streams scales capacity automatically, freeing you from provisioning and managing capacity. You can choose provisioned mode if you want to provision and manage throughput on your own.
- **High availability**: Kinesis Data Streams synchronously replicates data across three Availability Zones, providing high availability and data durability.
- **Accelerated log and data feed intake**: Instead of waiting to batch the data, you can have your data producers push data to a Kinesis data stream as soon as the data is produced, preventing data loss in case of producer failure. For example, system and application logs can be continuously added to a data stream and be available for processing within seconds.
- **Real-time metrics and reporting**: You can extract metrics and generate reports from Kinesis data stream data in real time. For example, your Amazon Kinesis application can work on metrics and reporting for system and application logs as the data is streaming in, rather than waiting to receive data batches.
- **Real-time data analytics**: With Kinesis Data Streams, you can run real-time streaming data analytics. For example, you can add clickstreams to your Kinesis data stream and have your Kinesis application run analytics in real time, allowing you to gain insights from your data in minutes instead of hours or days.
- **Log and event data collection**: Collect log and event data from sources such as servers, desktops, and mobile devices. You can then build applications using Amazon Lambda or Amazon Managed Service for Apache Flink to continuously process the data, generate metrics, power live dashboards, and emit aggregated data into stores such as Amazon Simple Storage Service (Amazon S3).
- **Power event-driven applications**: Quickly pair with AWS Lambda to respond or adjust to immediate occurrences within the event-driven applications in your environment, at any scale.

48.2 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/.

48.3 Pricing Overview

Please see the AWS UK G Cloud 14 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/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/

48.5 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.
- Amazon Kinesis Data Streams section of the AWS CLI Reference: Describes the command line interface for Kinesis Data Streams. Provides basic syntax, options, and usage examples for each command.

49 Amazon Lex

49.1 Service Overview

Amazon Lex is a fully managed artificial intelligence (AI) service with advanced natural language models to design, build, test, and deploy voice and text conversational interfaces in applications. Lex integrates with AWS Lambda, used 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 contact centers, chat and text platforms, and IoT devices. Lex provides rich insights and pre-built dashboards to track metrics for your

49.1.1 Features

- **Conversational FAQ:** Amazon Lex is committed to infusing Generative AI into all parts of the builder and end-user experiences to help increase containment while resolving increasingly complex use cases with confidence. Amazon Lex uses Generative AI and Large Langue Models (LLMs) to enhance the builder and customer experience. Builders can easily enable conversational answers to commonly asked customer questions by leveraging a new intent type that queries an authorized knowledge source and utilizes foundation models from Amazon Bedrock to provide an accurate response. This retrieval augmented generation (RAG) based solution provides a customized, conversational request and response framework that allows customers to get automated answers quickly and reliably, and improve self-service further.
- **Assisted Slot Resolution**: This feature leverages large language models (LLM) to resolve slot values when the native NLU is unable to. Customers can provide free form responses to basic questions and the LLM's enhanced reasoning capabilities will resolve the response to a specified format understood by the slot definitions.
- **Descriptive Bot Builder**: This Gen AI based feature allows bot developers to create a full bot based on a user's prompt. Simply provide a description in natural language to generate a baseline bot which can be further refined.
- **Sample Utterance Generation**: Builders can now automatically generate variations of sample utterances to improve intent classification accuracy with minimal effort. Using LLMs, we will provide sample utterances for training an intent based on the name, description, and existing utterances present in the intent.

- **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 is able to learn the multiple ways users can express their intent based on a few sample utterances provided by the bot builder. 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 & Multi-turn dialog**: 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.
- **Visual Conversation Builder**: The Visual Conversation Builder in the Amazon Lex console is a drag-and-drop conversation builder that accelerates bot building. Simply connect conversation nodes and easily iterate and test conversation designs in a no-code environment. It empowers any user to quickly build sophisticated and natural automated interactions, view conversation intent at a glance, and get visual feedback as changes are made.
- **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.
- **Automated Chatbot Designer**: Amazon Lex V2 offers an Automated Chatbot Designer that simplifies bot design by utilizing existing conversation transcripts. The designer analyzes the transcripts and proposes an initial bot design with intents and slot types. You can then customize the design by adding prompts, testing the bot, and deploying it.
- **Test Workbench**: Test workbench enables you to author and execute test sets to measure bot performance as you add new use cases and updates. After changes, you can ensure your bot meets the performance criteria by having Lex automatically generate audio and text based test sets from previous user interactions. Lex will then provide aggregated results and present detailed insights into speech transcription, intent matching, and slot resolution.
- **Analytics**: Lex Analytics gives you access to prebuilt dashboards detailing key metrics such as the number of total conversations and intent recognition rates. Analytics will help you better understand where in the conversation people are failing and how users navigate across intents.
- **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.

- **Integration with Amazon Bedrock**: Amazon Bedrock is a service that makes foundation models from Amazon and leading AI startups available through an API, so you can choose from various models to power generative-AI capabilities. Amazon Lex leverages Bedrock to call upon these foundation models to improve the builder experience and the end user experience.
- **Integration with Amazon Polly**: 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. You can use Polly to respond to your users in speech interactions. 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.
- **Integration with AWS Lambda**: Amazon Lex natively supports integration with AWS Lambda for data retrieval, updates, and business logic execution. The serverless compute capacity allows effortless execution of business logic at scale while you focus on developing bots. From Lambda, you can use AWS Lambda to easily integrate with your existing enterprise applications and databases. You just write your integration code and AWS Lambda automatically runs your code when needed to send or retrieve data from any external system. You can also access various AWS services, such as Amazon DynamoDB for persisting conversation state and Amazon SNS for notifying end users.
- **Integration with Amazon Kendra**: Customer service conversations often involve finding specific information to answer certain questions. Amazon Kendra provides you with a highly accurate and easy-to-use intelligent search service powered by machine learning. You can add a Kendra search intent to find the most accurate answers from unstructured documents and FAQs. You simply define the search index parameters in the intent as part of the bot definition to expand its informational capabilities.
- **Integration with the Contact Center applications**: This includes: Amazon Connect—Amazon Lex is natively integrated with Amazon Connect, AWS' omnichannel cloud contact center enabling developers to build conversational bots that can handle customer queries over chat or phone; Genesys Cloud CX—Genesys Cloud CX is a cloud contact center solution that unifies customers and agent experiences across multiple channels such as phone, text and chat. You can deploy your voice and text bots on the Genesys Cloud platform to enable self-service experiences and improve customer engagement. Please refer to Genesys Cloud integration for more information; Amazon Chime SDK—The Amazon Chime SDK is a set of real-time communications components that developers can use to quickly add audio calling, video calling, and screen sharing capabilities to their own web, mobile or telephony applications. Amazon Chime SDK integrates with Amazon Lex so you can easily enable conversational experiences powered by Amazon Lex in contact centers that use Session Initiation Protocol (SIP) for voice communication; AWS Contact Center Intelligence (CCI)–Amazon Lex is used by several AWS CCI partners, so you can seamlessly create self-service customer service virtual agents, informational bots or application bots. Amazon Lex partners include Infosys, Quantiphi, and Xapp.ai.

49.1.2 Benefits

- **Create Seamless Voice and Text Interfaces**: Design and build sophisticated "voice and text" conversational interfaces of your choice.

- **Effortless Omni-Channel Experiences**: Deploy omni-channel experiences with pre-built integrations to contact center solutions.
- **Rapid Market Entry**: Accelerate time to market with automated chat bot designer, test workbench and more.
- **Budget-Friendly Conversations**: Pay only for speech and text requests with no upfront costs or minimum fees.

49.2 Backup/Restore and Disaster Recovery

The service supports backup of interaction schemas and language models.

49.3 Pricing Overview

Please see the AWS UK G Cloud 14 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/lex/index.html
- **Service quotas:** https://docs.aws.amazon.com/lex/latest/dg/gl-limits.html
- **Service FAQs:** https://aws.amazon.com/lex/faqs/

49.5 Technical Requirements

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

- Lex 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.
- Lex 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.
- Lex V2 API Reference: Describes the Amazon Lex V2 API operations and parameters in detail.

50 Amazon Lightsail

50.1 Service Overview

Amazon Lightsail is a virtual private server (VPS) provider and is the easiest way to get started with AWS for developers, small businesses, students, and other users who need a solution to build and host their applications on cloud. Lightsail provides developers compute, storage, and networking capacity and capabilities to deploy and manage websites and web applications in the cloud. Lightsail includes everything you need to launch your project quickly–virtual machines, containers, databases, CDN, load balancers, DNS management for a low, predictable monthly price.

50.1.1 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. As you're creating your instance, Lightsail lets you click-to-launch a simple operating system (OS), a pre-configured application, or development stack—such as WordPress, Windows, Plesk, LAMP, Nginx, and more. Every Lightsail instance comes with a built-in firewall allowing you to allow or restrict traffic to your instances based on source IP, port, and protocol.

- **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. Lightsail Containers enables customers to run Docker containers on the cloud right from their developer workflows. Lightsail creates containers from the Docker images pushed by the developers, while we take care of the infrastructure management complexities.
- **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. Additionally, Lightsail load balancers include integrated certificate management, providing free SSL/TLS certificates you can provision and add to a load balancer in just a few clicks. You can request and manage certificates directly from the Lightsail console – and we manage renewals on your behalf.
- **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. As you're creating your instance, Lightsail lets you click-to-launch a simple operating system (OS), a pre-configured application, or development stack - such as WordPress, Windows, Plesk, LAMP, Nginx, and more. Every Lightsail instance comes with a built-in firewall allowing you to allow or restrict traffic to your instances based on source IP, port and protocol. Once your instance is running you can create custom alarms and be notified by email or text message when something is out of the ordinary.
- **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. Upgrading is simple, you'll take a snapshot of your instance and follow the step-by-step process in the Lightsail console to export your snapshot to EC2. You can then use the Upgrade to EC2 wizard to get your new EC2 instance up and running.

- **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.

50.1.2 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.

50.2 Backup/Restore and Disaster Recovery

The services supports backup of Amazon Lightsail instances, disks, and databases. Users can manually take snapshots of their Lightsail instances, disks, and databases, or use automatic snapshots. Users schedule backups through a web interface. Users can recover backups themselves via scheduled API call.

50.3 Pricing Overview

Please see the AWS UK G Cloud 14 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/lightsail/
- **Service quotas:** https://docs.aws.amazon.com/general/latest/gr/lightsail.html
- **Service FAQs:** https://aws.amazon.com/lightsail/faq/

50.5 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.

51 Amazon Location Service

51.1 Service Overview

Amazon Location Service is a fully managed service that makes it easy for developers to add location functionality, such as maps, points of interest, geocoding, routing, tracking, and geofencing to their applications, without sacrificing data security, user privacy, data quality, or cost.

51.1.1 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, HERE, and GrabMaps. The map tiles from Esri, HERE, and GrabMaps are trusted by many customers worldwide and have been continually 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, HERE, and GrabMaps 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, HERE, and GrabMaps. This enables your application to obtain accurate estimates of travel time based on up-to-date roadway information and live traffic information. Amazon Location Service Routes can help you achieve business goals such as faster delivery and reduced fuel consumption.
- **Trackers**: 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. You can also link Trackers with Amazon Location Service Geofences to automatically evaluate location updates from your devices against your geofences.
- **Geofences**: 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. This event-driven setup makes it easy for your application to scale in size and sophistication as needed.
- **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. This helps you shield sensitive information from third parties, protect your user privacy, and reduce organizational security risks.
- **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**: API calls made from any of the Amazon Location Service SDKs are secured through the AWS Signature Version 4 signing process for authentication. You can use one of the Amazon Location Service SDKs with your AWS IAM or Amazon Cognito authentication. These SDKs will sign the API calls and use these established technologies to help keep your data secure.

- **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. With AWS CloudFormation, create templates with Amazon Location Service to consistently and quickly provision resources. Using Amazon CloudWatch, you can monitor your metrics on service usage and health, including requests, latency, and faults, so you don't have to build your own health monitoring tools. Amazon Location Service is integrated with AWS CloudTrail for logging and continuously monitoring your account activity. Integration with Amazon EventBridge enables an event-driven application architecture so you can use AWS Lambda functions to activate parts of your application and workflows.
- **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. In the Amazon Location Service Console, developers can learn and get started quickly with a visual and interactive learning tool that enables them to explore the service's different geolocation capabilities.

51.1.2 Benefits

- **Integrate geospatial data**: Easily integrate geospatial data into your applications for different use cases or geographies with data from multiple providers.
- **Accelerate application development**: Quickly deploy and accelerate application development with seamless integrations with AWS services.
- **Enhance security posture**: Protect user privacy, shield your sensitive information, and reduce your application's security risks. 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.
- **Cost-effective, high-quality geospatial data from trusted data providers**: Amazon Location Service requires no up-front commitment, and no minimum fee. You are only charged for what you use. Amazon Location Service offers a free tier for the first three months of usage. Location data is billed based on each request your application makes to the service. Beyond the Amazon Location Service free tier, you pay for the requests your application makes to the service.

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

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

51.5 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.
- API reference: Provides detailed descriptions of the Amazon Location Service API operations and data types for maps, places, geofences, trackers, and routes in detail.
- CLI reference: Describes Amazon Location Service commands for maps, places, geofences, trackers, and routes in the AWS CLI. Provides syntax, options, and usage examples for each command.
- Sample apps for tutorials: The GitHub repository for Amazon Location Service sample apps described in the developer guide tutorials.

52 Amazon Lookout for Equipment

52.1 Service Overview

Amazon Lookout for Equipment uses the data from your sensors to detect abnormal equipment behavior, so you can take action before machine failures occur and avoid unplanned downtime. Customers that want to build ML models to monitor the health or efficiency of their equipment can directly upload their historical sensor data to Amazon Lookout for Equipment and automatically build a ML model that learns the normal behavior patterns and alerts to abnormal behavior. Customers can set up Amazon Lookout for Equipment to read real-time data from their equipment and detect the current behavior of the asset.

52.1.1 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.

52.1.2 Benefits

- **Ease of use**: Lookout for Equipment removes the complexities of data science so you can focus on your data and use case without needing deep technical experience in machine learning.
- **Fast proof of concepts**: Lookout for Equipment automates away many of steps in building models on industrial data so you can focus getting the right output and fast.
- **Cost Effectiveness**: Lookout for Equipment is priced based on usage, so training and inferencing can be set up and automated for cost effective model deployment,
- **Accuracy**: While not all equipment issues can be seen in the data, Lookout for Equipment can find key abnormal equipment behavior with best in class detection algorithms
- **Scale**: Lookout for Equipment can be used across thousands of pieces of equipment located across the world.

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 14 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/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/

52.5 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.

53 Amazon Lookout for Metrics

53.1 Service Overview

Amazon Lookout for Metrics uses machine learning (ML) to make it easier for customers to accurately detect anomalies in their metrics. After a customer uploads their data, Lookout for Metrics will automatically inspect it and uses ML to create accurate anomaly detection models. When anomalies are detected, Lookout for Metrics will group related anomalies together and provide a severity score so customers can diagnose issues quickly or maximize opportunities.

53.1.1 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.

53.1.2 Benefits

- **Automates anomaly detection on data using ML**: Lookout for Metrics supports a collection of machine learning algorithms. No one algorithm works for all kinds of data, so Lookout for Metrics inspects the data and applies the right ML algorithm to the right data to detect anomalies accurately. Because Lookout for Metrics takes care of the heavy lifting associated with selecting, training, building, and deploying the models, customers do not need any machine learning experience to get started quickly.
- **Helps customers get started with little to no data**: Customers can provide historical data for training the models if they choose. In the event customers choose not to or don't have historical data, Lookout for Metrics can learn as it goes. The amount of time for Lookout for Metrics to learn and detect anomalies will vary based on the customer's data.
- **Provides actionable results**: Lookout for Metrics will group concurrent anomalies into logical groups and send a single alert for the anomaly group rather than separate alerts, enabling customers to see the full-picture. The anomalies are also ranked in order of severity so that customers know which ones to focus on immediately.
- **Continuously improves via human-in-the-loop feedback**: Lookout for Metrics allows customers to provide feedback on the detected anomalies, which is used to improve the accuracy and performance of the models continuously.

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 14 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/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/

53.5 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.

54 Amazon Lookout for Vision

54.1 Service Overview

Amazon Lookout for Vision is a machine learning (ML) service that helps increase industrial production quality and reduce operational costs by identifying visual defects in objects. With Amazon Lookout for Vision you can identify visual defects at scale and decrease dependency on manual inspection.

54.1.1 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 labeling**: 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. You can also select multiple images and apply the label to the entire selection with one click. If you organize your images in Amazon Simple Storage Service (Amazon S3) into Normal and Anomaly folders, you can use the Auto Labeler feature to automatically label each image when creating a dataset. If you have a large data set, you can also use Amazon SageMaker Ground Truth to efficiently label your images at scale.
- **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. For every image in the test dataset, you can see a side-by-side comparison of the model's prediction vs. the label assigned. You can also review detailed performance metrics such as precision/recall metrics, F1 score, and confidence scores.

- **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. Your trained models can be deployed on any NVIDIA Jetson edge appliance or x86 compute platform running Linux with an NVIDIA GPU accelerator. You can use AWS IoT Greengrass V2 to deploy, and manage your edge compatible customized models on your fleet of devices. AWS IoT Greengrass is an open source Internet of Things (IoT) edge runtime and cloud service that helps you build, deploy, and manage IoT applications on your devices.
- **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.

54.1.2 Benefits

- **Detect damages to parts**: You can detect defects related to welds, dents, cracks, bubbles, surface color, and shape of objects.
- **Identify missing components**: You can identify missing components or parts related to absence, presence, placement of objects within a certain tolerance.
- **Uncover process issues**: You can detect a defect that has a repeating pattern, which can indicate a process issue.

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 14 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/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/

54.5 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.

55 Amazon Macie

55.1 Service Overview

Amazon Macie is a data security service that discovers sensitive data using machine learning and pattern matching, provides visibility into data security risks, and enables automated protection against those risks.

55.1.1 Features

- **Ongoing evaluation of your Amazon S3 security posture**: Amazon Macie continually evaluates your Amazon S3 environment and provides a summary of your data security posture across all of your accounts. You can search, filter, and sort S3 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. Macie then automatically samples and analyzes objects in your S3 buckets, inspecting them for sensitive data such as personally identifiable information (PII), builds an interactive data map of where your sensitive data in S3 resides across accounts, and provides sensitivity score for each bucket. The interactive data map can guide your decisions to perform deeper investigations of specific S3 buckets by running targeted, sensitive data discovery jobs with Macie.
- **Targeted sensitive data discovery**: 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 targeted 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.
- **Discover proprietary or unique 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.
- **Securely review and validate sensitive data found in an Amazon S3 object**: Macie allows for one-selection, temporary retrieval of up to 10 examples of sensitive data found in S3. This capability helps you more easily view and understand which contents of an S3 object were identified to be sensitive, so you can review, validate, and quickly take action as needed. All sensitive data examples captured are encrypted using customer-managed AWS Key Management Service (KMS) keys and are temporarily viewable within the Macie console after being retrieved.
- **Create and manage allow lists to specify text or text patterns**: Macie's allow list feature can help you reduce alert volume due to data text or formats in your

environment that do not require action. An allow list defines specific text or a text pattern that you want Macie to ignore when it inspects S3 objects for sensitive data. If text matches an entry or pattern in an allow list, Macie doesn't report the text in sensitive data findings or sensitive data discovery results, even if the text matches the criteria of a managed data identifier or a custom data identifier.

- **One-selection deployment with no upfront data source integration**: With one selection in the AWS Management Console or a single API call, you can enable Amazon Macie in a single account. With a few more selections in the console, you can enable Macie across multiple accounts. Once enabled, Macie generates an ongoing 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. Macie supports multiple accounts through AWS Organizations integration. Security and sensitive data discovery findings are aggregated in the Macie administrator account and sent to Amazon EventBridge. 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.

55.1.2 Benefits

- **Automate sensitive data discovery at scale**: Macie uses machine learning and pattern matching to discover sensitive data at scale in a cost-efficient way. Macie automatically detects a large and growing list of sensitive data types, including personally identifiable information (PII) such as names, addresses, and credit card numbers.
- **Gain cost-efficient visibility into sensitive data stored in Amazon S3**: Macie gives you constant visibility of your data stored in Amazon Simple Storage Service (Amazon S3).
- **Simplified setup**: Macie's setup is simplified with one selection 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 selections.
- **Strengthen your data security posture**: Discover sensitive data across your S3 environment to increase visibility and automated remediation of data security risks.

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

55.5 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.

56 Amazon Managed Blockchain

56.1 Service Overview

Amazon Managed Blockchain (AMB) is a fully managed service designed to help you build resilient Web3 applications on both public and private blockchains. Use AMB Access for instant and serverless access to multiple blockchains. Build your Web3-ready applications without having to worry about deploying specialized blockchain infrastructure and keeping them connected to the blockchain network. With AMB Query, you can use developer-friendly APIs to access real-time and historical data from multiple blockchains. The standardized blockchain data can be integrated with AWS services, with no specialized blockchain infrastructure or ETL (extract, transform, and load) needed. All AMB features scale securely for institutional grade and mainstream consumer application builds.

56.1.1 Features

- **Bitcoin (BTC)**: Bitcoin is a decentralized blockchain network that enables secure peer-to-peer transactions of value denominated in the network's native cryptocurrency, Bitcoin. The Bitcoin network is used by individuals, financial institutions, fintech companies, governments, and more as a medium of exchange, a commodity for investment, or a publicly verifiable and immutable ledger for inscribed data. With AMB Access, you can access a pool of Bitcoin mainnet and testnet nodes through regional endpoints, allowing you to write transactions, read data from the ledger, and invoke JSON-RPC methods available on the Bitcoin Core node client. With serverless Bitcoin endpoints, you can focus on building your applications instead of investing in undifferentiated work such as provisioning, maintaining, and load-balancing Bitcoin nodes.
- **Ethereum (ETH)**: Ethereum is a decentralized and programmable blockchain network on which users around the world can transact, collaborate, and build applications. The Ethereum virtual machine (EVM) helps developers create powerful and composable decentralized applications (dApps) in the form of smart contracts. Build Ethereum dApps on mainnet and select testnets with Ethereum full nodes using the go-ethereum (Geth) execution client and the Lighthouse consensus client. You can use your dedicated (single-tenant) Ethereum node(s) to invoke JSON-RPC APIs for both the Execution and Consensus layers to build and test smart contracts, perform fungible or non-fungible token transactions, or query data from the Ethereum blockchain.
- **Polygon (POL)**: Polygon is an Ethereum Virtual Machine (EVM) compatible, Layer-2 scaling solution. Known for its high throughput and low fees, Polygon is a preferred blockchain for several Web2 and Web3 builders building decentralized applications (DApps). Through AMB Access Polygon (Preview), you can get serverless access to Polygon PoS mainnet and Mumbai testnet, invoke JSON-RPC APIs to build and test smart contracts, perform fungible or non-fungible token transactions, or retrieve data from the Polygon blockchain.

- **Hyperledger Fabric**: You can deploy a private, permissioned blockchain with the Hyperledger Fabric framework in minutes without complex configuration tasks. Once your network is created, you can launch one or more blockchain peer nodes for high availability using the AWS Management Console. Whether your use case demands a large consortium of members to join the network or you simply wish to create a development network within your organization, AMB simplifies the process. You can invite other organizations to join your blockchain network through their AWS accounts, or you can create additional members (organizations) in your AWS account to build a simulated network for testing.
- **Token APIs**: AMB Query provides useful APIs that simplify building blockchain applications, including queriesrelated to token balances, transaction details, and on-chain events.
- **Transaction APIs**: Customers can use the AMB Query Transaction APIs to query information related to transaction details and on-chain events.
- **Contract APIs**: Customers can use the AMB Query Contract APIs to query information related to smart contracts. For example, customers can use the AMB Query ListAssetContracts API to list ERC-20, ERC-721, and ERC-1155 smart contracts and the GetAssetContract API to get the properties of a smart contract, such as the type and deployer address.

56.1.2 Benefits

- **Reliable blockchain infrastructure for public and private blockchains**: Amazon Managed Blockchain (AMB) Access provides public blockchains nodes for Ethereum, Polygon (Preview) and Bitcoin, as well as the ability to create private blockchain networks with the Hyperledger Fabric framework. Choose from various methods to engage with public blockchains, including serverless multi-tenant APIs to fully managed (dedicated) public blockchain nodes.
- **Quickly ship applications**: Amazon Managed Blockchain (AMB) Query provides serverless access to standardized, multi-blockchain datasets with developer-friendly APIs. You can use AMB Query to more quickly ship applications that require data from one or more public blockchains, without the overhead of parsing blockchain data, tracing smart contracts, and maintaining specialized indexing infrastructure.

56.2 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. Users schedule backups through a web interface or recover backups using manual or scheduled APIs.

56.3 Pricing Overview

Please see the AWS UK G Cloud 14 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/managed-blockchain/
- **Service quotas:** https://docs.aws.amazon.com/general/latest/gr/managedblockchain.html
- **Service FAQs:** https://aws.amazon.com/managed-blockchain/faqs/

56.5 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.

57 Amazon Managed Service for Apache Flink

57.1 Service Overview

With Amazon Managed Service for Apache Flink—a framework and engine for building streaming applications., you can transform and analyze streaming data in real time with Apache Flink. Apache Flink is an open source framework and engine for processing data streams. Amazon Managed Service for Apache Flink reduces the complexity of building, managing, and integrating Apache Flink applications with other AWS services.

Amazon Managed Service for Apache Flink takes care of everything required to continuously run streaming applications and scales automatically to match the volume and throughput of your incoming data. With Amazon Managed Service for Apache Flink, there are no servers to manage, there is no minimum fee or setup cost, and you only pay for the resources your streaming applications consume.

57.1.1 Features

- **Open source**: Amazon Managed Service for Apache Flink includes open source libraries such as Apache Flink, Apache Beam, Apache Zeppelin, AWS SDK, and AWS service integrations. Apache Flink is a framework and engine for building highly available and accurate streaming applications. Apache Beam is a unified model for defining streaming and batch data processing applications that are run across multiple runtime engines. AWS SDKs help take the complexity out of coding for many AWS services by providing APIs in your preferred language, and they include AWS libraries, code samples, and documentation.
- **Flexible APIs**: Amazon Managed Service for Apache Flink supports Flink's flexible APIs in Java, Scala, Python, and SQL that are specialized for different use cases including stateful event processing, streaming ETL (extract, transform, and load), and real-time analytics. With prebuilt operators and analytics capabilities, you can build an Apache Flink streaming application in hours instead of months, and the libraries are extensible, so you can perform real-time processing for various use cases.
- **AWS service integrations**: You can set up and integrate a data source or destination with minimal code. Use the Amazon Managed Service for Apache Flink libraries to integrate with the following AWS services: Amazon S3, Amazon MSK, Amazon OpenSearch Service, Amazon DynamoDB, Amazon Kinesis Data Streams, Amazon Kinesis Data Firehose, Amazon CloudWatch, AWS Glue Schema Registry, Advanced integration capabilities.
- **Exactly-once processing**: Using Amazon Managed Service for Apache Flink, you can build applications where processed records affect the results exactly once, referred to as exactly-once processing. Even in the case of an application disruption, such as internal service maintenance or user-initiated application update, the service ensures all data is processed and there is no duplicate data.

- **Stateful processing**: The service stores previous and in-progress computations, or state, in running application storage. Compare real-time and past results over any time period and achieve fast recovery during application disruptions. State is always encrypted and incrementally saved in running application storage.
- **Durable application backups**: Create and delete durable application backups through a simple API call. Immediately restore your applications from the latest backup after a disruption, or restore your application to an earlier version.
- **ML integration**: Amazon Managed Service for Apache Flink supports machine learning (ML) algorithms. You can create real-time applications for classification, clustering, evaluation, feature engineering recommendations, regressions, and statistics.
- **AWS Glue Schema Registry compatibility**: Amazon Managed Service for Apache Flink is compatible with the AWS Glue Schema Registry. The Schema Registry helps you improve data quality and safeguard against unexpected changes using compatibility checks that govern schema evolution for your schemas on Amazon Managed Service for Apache Flink workloads connected to Apache Kafka, Amazon MSK, or Amazon Kinesis Data Streams, as either a source or sink connector.

57.1.2 Benefits

- **Easy to build applications**: Build and run Apache Flink applications, without setting up infrastructure and managing resources and clusters.
- **High speed processing**: Process gigabytes of data per second with sub-second latencies and respond to events in real time.
- **Highly available**: Deploy highly available and durable applications with Multi-AZ deployments and APIs for application lifecycle management.

57.2 Backup/Restore and Disaster Recovery

Amazon Managed Service for Apache Flink automatically backs up your running application's state using checkpoints and snapshots. Checkpoints save the current application state and enable Amazon Managed Service for Apache Flink applications to recover the application position to provide the same semantics as a failure-free execution. Checkpoints use running application storage. Checkpoints for Apache Flink applications are provided through Apache Flink's checkpointing functionality. Snapshots save a point-in-time recovery point for applications and use durable application backups. Users schedule backups through a web interface or recover backups using manual or scheduled APIs.

57.3 Pricing Overview

Please see the AWS UK G Cloud 14 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/managed-flink/
- **Service quotas:** https://docs.aws.amazon.com/managed-flink/latest/java/limits.html
- **Service FAQs:** https://aws.amazon.com/managed-service-apache-flink/faqs/

57.5 Technical Requirements

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

- Developer Guide for Apache Flink Applications: Provides a conceptual overview of Managed Service for Apache Flink applications, including detailed instructions for using the various features.
- API Reference: Describes the API operations for Managed Service for Apache Flink (formerly known as Kinesis Data Analytics, and provides sample requests, responses, and errors for the supported web services protocols.
- Managed Service for Apache Flink section of the AWS CLI Reference: Describes the command line interface for Managed Service for Apache Flink (formerly known as Kinesis Data Analytics) in detail. Provides basic syntax, options, and usage examples for each command.

58 Amazon Managed Service for Grafana

58.1 Service Overview

Amazon Managed Grafana is a fully managed multicloud, cross-project service with rich, interactive data visualizations to help customers analyze, monitor, and alarm on metrics, logs, and traces across multiple data sources. You can create interactive dashboards and share them with anyone in your organization with an automatically scaled, highly available, and enterprise-secure service.

58.1.1 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.
- **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.
- **Extend your experience with Grafana community plugins**: With Amazon Managed Grafana, you can install Grafana community plugins to connect to additional third-party data sources such as Apache Cassandra, Sentry or to leverage new visualization panels such as Flowcharting, Polystat, etc. You can discover, install and manage version updates for your plugins, directly from the Plugin Catalog in your Amazon Managed Grafana workspace.
- **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 IAM Identity Center (successor to 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.

- **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.
- **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 IAM Identity Center (successor to 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.
- **No servers to manage**: 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.
- **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.
- **Upgrade to Grafana Enterprise directly from the AWS Console**: You can optionally upgrade to Grafana Enterprise via an AWS Marketplace purchase from the Amazon Managed Grafana console. This gives you access to additional enterprise plugins for a wide variety of third-party ISVs, including AppDynamics, Atlassian Jira, Datadog, Dynatrace, Gitlab, Honeycomb, MongoDB, New Relic, Oracle Database, Salesforce, SAP HANA, ServiceNow, VMware Tanzu Observability by Wavefront, and Snowflake.

58.1.2 Benefits

- **Build workspaces**: Build, package, and deploy workspaces that are provisioned, set up, scaled, and maintained for you.
- **Operational data**: Visualize, analyze, and correlate operational data across multiple sources, and query across multiple AWS accounts and Regions.
- **Integrate with AWS security**: Integrate with AWS security services to meet your corporate security and compliance requirements
- **Migrate Grafana**: Migrate from your self-managed Grafana environment, so there’s no need to start from scratch.

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 14 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/grafana/index.html
- **Service quotas:** https://docs.aws.amazon.com/grafana/latest/userguide/AMG_quotas.html
- **Service FAQs:** https://aws.amazon.com/grafana/faqs/

58.5 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.

59 Amazon Managed Service for Prometheus

59.1 Service Overview

Amazon Managed Service for Prometheus is a serverless monitoring service for metrics compatible with open source Prometheus, making it easier for you to securely monitor and alert on container environments. As a fully managed service, Amazon Managed Service for Prometheus helps with monitoring containers on AWS, on premises, and in hybrid and multicloud environments; it is a cross-project solution based on the popular open-source Cloud Native Computing Foundation (CNCF) Prometheus project. Amazon Managed Service for Prometheus is powered by Cortex, an open-source CNCF project that adds horizontal scalability to ingest, store, query, and alert on Prometheus metrics.

Amazon Managed Service for Prometheus makes it easy to get started with monitoring applications across Amazon Elastic Kubernetes Service, Amazon Elastic Container Service, and AWS Fargate, and self-managed Kubernetes clusters. Amazon Managed Service for Prometheus automatically scales as your monitoring needs grow. It offers highly available, multi-Availability Zone deployments, and integrates AWS security and compliance capabilities. Amazon Managed Service for Prometheus offers native support for the Prometheus Query Language (PromQL) as well as the ability to ingest Prometheus metrics from over 150+ Prometheus exporters maintained by the open source community.

59.1.1 Features

- **Setup and configuration**: Getting started with Amazon Managed Service for Prometheus is easy. With just a few clicks in the AWS console, you can create an Amazon Managed Service for Prometheus workspace, which is a logical space dedicated to the storage and querying of Prometheus metrics. 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. Amazon Managed Service for Prometheus collector is a serverless feature that you can use to securely discover and collect Prometheus metrics from your Amazon EKS applications and infrastructure using a Prometheus-compatible configuration.

- **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.
- **No collection agents required**: For Amazon EKS workloads, you can configure Amazon Managed Service for Prometheus collector to collect Prometheus metrics from Amazon EKS applications and infrastructure without the need to build, package, or deploy any agents in-cluster. With Amazon Managed Service for Prometheus collector, you don't have to worry about scaling, right-sizing, patching, or operating any agents to collect Prometheus metrics. Amazon Managed Service for Prometheus collector is pre-configured to provide high availability, multi-AZ collection for Prometheus metrics.
- **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**: 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.
- **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.
- **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. Additionally, Amazon Managed Service for Prometheus collector, an agentless scraper, can be utilized to collect Prometheus metrics from Amazon EKS. Metrics can be ingested from any clusters running on AWS and hybrid environments, with on-demand scaling to meet your growing needs.
- **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.
- **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.

59.1.2 Benefits

- **Monitor applications across self-managed Kubernetes**: Use Prometheus query language (PromQL) to filter, aggregate, ingest, and query millions of unique time series metrics from your self-managed Kubernetes clusters.
- **Secure**: Integrate with AWS security services to meet your corporate security and compliance requirements.
- **Simple to manage**: Automatically scale as your ingestion and query needs grow, and maintain consistent response times for large container deployments. Reduce the operational costs of configuring, upgrading, and scaling standalone Prometheus servers.

59.2 Backup/Restore and Disaster Recovery

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.

59.3 Pricing Overview

Please see the AWS UK G Cloud 14 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/prometheus/
- **Service quotas:** https://docs.aws.amazon.com/prometheus/latest/userguide/AMP_quotas.html
- **Service FAQs:** https://aws.amazon.com/prometheus/faqs/

59.5 Technical Requirements

Please refer to https://docs.aws.amazon.com/prometheus/ 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 that are available as part of the AWS SDKs. Use these APIs to manage workspaces, rule groups, and scrapers.
- Prometheus-compatible API Reference: Describes the Prometheus-compatible APIs provided by Amazon Managed Service for Prometheus. Use the Prometheus-compatible APIs to manage alerts, query metrics, and write metrics to your workspace.

60 Amazon Managed Streaming for Apache Kafka

60.1 Service Overview

Amazon Managed Streaming for Apache Kafka (Amazon MSK) is an AWS streaming data service that manages Apache Kafka infrastructure and operations, making it easy for developers and DevOps managers to run Apache Kafka applications and Kafka Connect connectors on AWS, without the need to become experts in operating Apache Kafka. Amazon MSK operates, maintains, and scales Apache Kafka clusters, provides enterprise-grade security features out of the box, and has built-in AWS integrations that accelerate development of streaming data applications.

60.1.1 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.
- **Amazon MSK serverless**: MSK serverless is a cluster type for Amazon MSK that makes it easy for you to run Apache Kafka clusters without having to manage compute and storage capacity. MSK Serverless automatically provisions and scales resources while also managing Apache Kafka partitions, so you can stream data without having to worry about right-sizing or scaling clusters.
- **High availability is default**: All clusters are distributed across multiple AZs (three is the default), are supported by Amazon MSK's service-level agreement, 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.
- **Reliable and automated data replication across MSK clusters**: MSK Replicator is a feature of Amazon MSK that enables you to reliably replicate data across Amazon MSK clusters in just a few clicks without requiring expertise to setup open-source tools, writing code, or managing infrastructure. MSK Replicator automatically provisions and scales underlying resources, so that you can replicate data on-demand and pay only for what you use. With MSK Replicator, you can build highly available and fault-tolerant multi-region applications for increased resiliency. You can

also use MSK Replicator to provide lower latency data access in different geographic regions or to distribute data to your partners.

- **Highly secure offering private connectivity**: Your Apache Kafka cluster runs in an Amazon Virtual Private Cloud (VPC) managed by Amazon MSK. Kafka clients in your own Amazon VPC can access the cluster privately through a cross-account elastic network interface that Amazon MSK deploys in your VPC. If your Kafka clients are spread across one or more VPCs or AWS accounts, you can still connect privately to your cluster by using the multi-VPC private connectivity feature. This feature eliminates the operational overhead of self-managing a PrivateLink solution and scales seamlessly as the Amazon MSK cluster scales, enabling you to maintain private connectivity to the cluster without making additional configuration changes. Multi-VPC private connectivity also eliminates the challenges with managing non-overlapping IPs, complex peering and routing tables associated with other VPC connectivity solutions as it allows for overlapping IPs across connecting VPCs.
- **Cross-account access control**: Use a cluster policy for your Amazon MSK cluster to define which IAM principals have cross-account permissions to set up private connectivity to your Amazon MSK cluster. When used with IAM client authentication, you can also use the cluster policy to granularly define Kafka data plane permissions for the connecting clients.
- **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.
- **Connectivity over the internet**: Amazon MSK offers an option to securely connect to the brokers of Amazon MSK clusters running Apache Kafka 2.6.0 or later versions over the internet. By enabling Public Access, authorized clients external to a private Amazon Virtual Private Cloud (VPC) can stream encrypted data in and out of specific Amazon MSK clusters.
- **Cross-account access control**: Use a cluster policy for your Amazon MSK cluster to define which cross-account IAM principals have permissions to set up cross-account private connectivity to your Amazon MSK cluster. When used with IAM client authentication, you can also use the cluster policy to granularly define Kafka data plane permissions for the connecting clients.
- **Deeply integrated**: Amazon MSK integrations include: AWS IAM for Apache Kafka and service-level API access control, Amazon Managed Service for Apache Flink for running fully managed Apache Flink applications to process streaming data within Apache Kafka, Amazon Managed Service for Apache Flink 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.

- **Graviton**: AWS Graviton3 processors are the latest generation of custom-designed AWS Graviton processors built on the AWS Nitro System. The Graviton3 processors based M7g instances deliver higher storage throughput and increased network throughput compared to similar sized M5 instances at a lower cost.
- **Open source**: 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 $2.00 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.
- **Tiered storage**: With tiered storage, you can store virtually unlimited data in MSK without the need to provision and manage storage capacity with tiered storage. You can enable tiered storage with a few clicks for new or existing clusters and pay for what you use. You can first store data in a performance optimized primary storage tier and let MSK automatically tier data into the new low-cost tier for longer retention. The feature is supported in all AWS regions where MSK is present. To learn how to get started tiered storage, visit our Amazon MSK Developer Guide.
- **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.
- **Visible**: 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.

60.1.2 Benefits

- **Ease of operations**: Eliminate operational overhead, including the provisioning, configuration, and maintenance of highly available Apache Kafka and Kafka Connect clusters. Use applications and tools built for Apache Kafka out of the box (no code changes required), and scale cluster capacity automatically.
- **Secure**: Easily deploy secure, compliant, and production-ready applications using native AWS integrations.
- **Cost effective**: 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 using features such as Tiered Storage and Graviton instances.

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

60.5 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.
- Amazon MSK API Reference 1.0: Describes the API resources and operations for Amazon MSK. Also provides sample requests, responses, and errors for the supported web services protocols.

- Amazon MSK Connect API Reference: Describes the API resources and operations for Amazon MSK Connect. Also provides sample requests, responses, and errors for the supported web services protocols.

61 Amazon Managed Workflows for Apache Airflow

61.1 Service Overview

Amazon Managed Workflows for Apache Airflows (MWAA), is a managed Apache Airflow service used to extract business insights across an organization by combining, enriching, and transforming data through a series of tasks called a workflow. Managed Workflows free you from managing, configuring, and scaling the Airflow environment while you orchestrate data processing workflows and manage their execution through AWS-backed logging and monitoring capabilities. You can run your existing Airflow workflows on Amazon MWAA and interact with their environment programmatically using the AWS console, API, and command line interface (CLI).

61.1.1 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.

61.1.2 Benefits

- **Scalability**: Deploy Apache Airflow at scale without the operational burden of managing underlying infrastructure.
- **Security and monitoring**: Run Apache Airflow workloads in your own isolated and secure cloud environment. Monitor environments through Amazon CloudWatch integration to reduce operating costs and engineering overhead.
- **Connectivity**: Connect to AWS, cloud, or on-premises resources through Apache Airflow providers or custom plugins.

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 14 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/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/

61.5 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.
- API Reference: Describes the API operations for Amazon MWAA.
- Migration Guide: Describes migrating to a new Amazon MWAA environment from an existing Amazon MWAA environment, or from a self-managed Apache Airflow deployment.

62 Amazon MemoryDB for Redis

62.1 Service Overview

Amazon MemoryDB for Redis is a Redis-compatible, durable, in-memory database service that delivers ultra-fast performance. MemoryDB enables you to achieve microsecond read latency, single-digit millisecond write latency, high throughput, and Multi-AZ durability for modern applications, like those built with microservices architectures. These applications require low latency, high scalability, and use Redis' flexible data structures and APIs to make development agile and easy. MemoryDB stores your entire dataset in memory and leverages a distributed transactional log to provide both in-memory speed and data durability, consistency, and recoverability. You can use MemoryDB as a fully managed, primary database, enabling you to build high-performance applications without having to separately manage a cache, durable database, or the required underlying infrastructure.

62.1.1 Features

- **Redis compatibility**: Redis is a fast, open source, in-memory, key-value data store. Redis offers flexible APIs, commands, and data structures like streams, sets, and lists, to build agile and versatile applications. MemoryDB maintains compatibility with open source Redis and supports the same set of Redis data types, parameters, and commands that you are familiar with. This means that the code, applications, drivers, and tools you already use today with Redis can be used with MemoryDB so you can quickly build applications.

- **Ultra-fast performance**: MemoryDB stores your entire dataset in memory to deliver microsecond read latency, single-digit millisecond write latency, and high throughput. It can handle more than 13 trillion requests per day and support peaks of 160 million requests per second. Developers building with microservices architectures require ultra-high performance as these applications can involve interactions with many service components per user interaction or API call. With MemoryDB, you enable extreme low latency to deliver real-time performance for end users.
- **Multi-AZ durability**: In addition to storing your entire data set in memory, MemoryDB uses a distributed transactional log to provide data durability, consistency, and recoverability. MemoryDB stores data across multiple AZs so you can achieve fast database recovery and restart. You can use MemoryDB as a single, primary database service for your workloads requiring low-latency and high throughput instead of separately managing a cache for speed and an additional relational or nonrelational database for reliability.
- **Scalability**: You can scale your MemoryDB cluster to meet fluctuating application demands: horizontally by adding or removing nodes, or vertically by moving to larger or smaller node types. MemoryDB supports write scaling with sharding and read scaling by adding replicas. Your cluster continues to stay online and support read and write operations during resizing operations.
- **Easy to use**: Getting started with MemoryDB is easy. Just launch a new MemoryDB cluster using the AWS Management Console, or you can use the AWS CLI or SDK. MemoryDB database instances are pre-configured with parameters and settings appropriate for the node type select. You can launch a cluster and connect your application within minutes without additional configuration.
- **Monitoring and metrics**: MemoryDB provides Amazon CloudWatch metrics for your database instances. You can use the AWS Management Console to view over 35 key operational metrics for your cluster including compute, memory, storage, throughput, active connections, and more.
- **Automatic software patching**: MemoryDB automatically keeps your clusters up-to-date with new updates, and you can easily upgrade your clusters to the latest versions of Redis.
- **Networking**: MemoryDB runs in Amazon VPC, which allows you to isolate your database in your own virtual network and connect to your on-premises IT infrastructure using industry-standard, encrypted IPsec VPNs. In addition, using MemoryDB's VPC configuration, you can configure firewall settings and control network access to your database instances.
- **Encryption**: With MemoryDB, data at-rest is encrypted using keys you create and control through AWS Key Management Service (KMS). And, clusters created with AWS Graviton2 node types include always-on 256-bit DRAM encryption. MemoryDB supports encryption in-flight using Transport Layer Security (TLS).
- **API-level Permissions**: Using the AWS Identity and Access Management (IAM) features integrated with Amazon MemoryDB, you can control the actions that your AWS IAM users and groups can take on Amazon MemoryDB resources. For example, you can configure your IAM rules to help ensure that certain users only have read-only access, while an Administrator can create, modify, and delete resources.
- **Authentication and authorization**: MemoryDB uses Redis Access Control Lists (ACLs) to control both authentication and authorization for your cluster. ACLs enable you to define different permissions for different users in the same cluster.
- **Integration with Kubernetes**: AWS Controllers for Kubernetes (ACK) for Amazon MemoryDB enables you to define and use MemoryDB resources directly from your

Kubernetes cluster. This lets you take advantage of MemoryDB to support your Kubernetes applications without needing to define MemoryDB resources outside of the cluster or run and manage in-memory database capabilities within the cluster.

- **JSON support**: Amazon MemoryDB provides native support for JavaScript Object Notation (JSON) documents in addition to the data structures included in open source Redis, at no additional cost. You can simplify developing applications by using the built-in commands designed and optimized for JSON documents. MemoryDB supports partial JSON document updates, as well as powerful searching and filtering using the JSONPath query language. JSON support is available when using Redis 6.2 and above.
- **Vector search**: Amazon MemoryDB for Redis enables machine learning (ML) and generative artificial intelligence (AI) models to work with data stored in Amazon MemoryDB in real-time and without moving your data. With Amazon MemoryDB, you can store, search, index, and query vector embeddings within Redis data structures.
- Vectors are numerical representations of unstructured data, such as text, images, and video, created from ML models that help capture the semantic meaning of the underlying data. You can store vector embeddings from ML and AI models, such as those from Amazon Bedrock and Amazon SageMaker in your Amazon MemoryDB database.
- **Cost optimization**: MemoryDB offers data tiering as a lower cost way to scale your clusters up to hundreds of terabytes of capacity. Data tiering provides a price-performance option for MemoryDB by utilizing lower-cost solid state drives (SSDs) in each cluster node in addition to storing data in memory. It is ideal for workloads that access up to 20% of their overall dataset regularly, and for applications that can tolerate additional latency when accessing data on SSDs.

62.1.2 Benefits

- **Scalability**: Scale to hundreds of millions of requests per second and over one hundred terabytes of storage per cluster.
- **Availability**: Stores data durably using a Multi-AZ transaction log for 99.99% availability and near instantaneous recovery without data loss.
- **Security**: Secure your data with encryption at rest and in-transit, private VPC endpoints, and multiple authentication mechanisms, including IAM authentication.
- **Integration with AWS services**: Build applications quickly with Redis data structures, a rich open source API, and easily integrate with other AWS services.

62.2 Backup/Restore and Disaster Recovery

Users can create snapshots to back up the data and metadata of your MemoryDB cluster. You can manually create a snapshot, or you can use MemoryDB’s automated snapshot scheduler to take a new snapshot each day at a time you specify. You can choose to retain your snapshot for up to 35 days after it is created, and MemoryDB. Snapshots are stored in Amazon S3 which is designed for 99.999999999% (11 9's) durability. Also, you can choose to take a final snapshot of your cluster when you are deleting the cluster. Furthermore, you can export MemoryDB snapshots from the service to your Amazon S3 bucket.

Users schedule backups through a web interface or recover backups using manual or scheduled APIs.

62.3 Pricing Overview

Please see the AWS UK G Cloud 14 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/memorydb/
- **Service quotas:** https://docs.aws.amazon.com/memorydb/latest/devguide/quota-limits.html
- **Service FAQs:** https://aws.amazon.com/memorydb/faqs/

62.5 Technical Requirements

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

- Developer Guide: Helps you understand the components and features provided and how to use them. Learn how to access MemoryDB for Redis through a web-based GUI, command line tools, and the MemoryDB for Redis API.
- API Reference: Describes all the API operations for MemoryDB for Redis in detail. Also provides sample requests, responses, and errors for the supported web services protocols.
- MemoryDB for Redis section of the AWS CLI Reference: Describes the command line interface for Amazon MemoryDB for Redis in detail. Provides basic syntax, options, and usage examples for each command.

63 Amazon Monitron

63.1 Service Overview

Amazon Monitron is an end-to-end system that detects abnormal behavior in industrial machinery, so you can implement predictive maintenance and reduce unplanned downtime. Amazon Monitron captures vibration and temperature data from your equipment via wireless sensors, which can be set up in minutes using the Amazon Monitron mobile app with no ML experience or development work required. The 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.

63.1.1 Features

- **End-to-end system of wireless sensors and gateways**: Amazon Monitron's wireless sensors can easily attach to your equipment with adhesive, without the need for expensive and cumbersome cabling. These low-cost sensors (including Ex-rated sensors) are purpose-built to capture vibration and temperature data in order to monitor the condition of your rotating equipment. Amazon Monitron also includes Wi-Fi and ethernet gateways to transfer the sensor data to AWS. Amazon Monitron Sensors and Gateways are pre-configured to work with the Amazon Monitron service.
- **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.

63.1.2 Benefits

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

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 14 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/Monitron/index.html
- **Service quotas:** https://docs.aws.amazon.com/Monitron/latest/admin-guide/quotas.html
- **Service FAQs:** https://aws.amazon.com/monitron/faqs/

63.5 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.
- User Guide: Describes how admins and technicians use the Amazon Monitron mobile and web apps to monitor equipment for potential failures.

64 Amazon MQ

64.1 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 in the cloud. You get direct access to the ActiveMQ and RabbitMQ consoles and industry standard APIs and protocols for messaging, including JMS, NMS, AMQP 1.0 and 0.9.1, STOMP, MQTT, and WebSocket. You can easily move from any message broker that uses these standards to Amazon MQ because you don't have to rewrite any messaging code in your applications.

64.1.1 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. 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.
- **Monitoring**: Amazon MQ is integrated with Amazon CloudWatch and AWS CloudTrail. With CloudWatch you can monitor metrics on your brokers, queues, and topics.
- **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.
- **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.
- **RabbitMQ Features**: High Availability, Throughput, and Message Durability: Amazon MQ runs on the same highly reliable infrastructure used by other Amazon Web Services. Amazon MQ for RabbitMQ clusters, backed by Amazon EBS, use multi-AZ replication for high-availability and message durability.
- **Message Routing**: Messages in RabbitMQ brokers are routed through exchanges before arriving at queues. RabbitMQ features several built-in exchange types for typical routing logic.
- **Broad Client Language Support**: Develop using your favorite programming languages, including: Python, .NET, PHP, Python, JavaScript, Ruby, Java, and Go
- **Other RabbitMQ Features**: For a full description of RabbitMQ's features, please see the RabbitMQ website.
- **ActiveMQ Features**: High Availability, Throughput, and Message Durability: Amazon MQ runs on the same highly reliable infrastructure used by other Amazon Web Services. Amazon MQ for ActiveMQ provides durability-optimized brokers backed by Amazon Elastic File System (Amazon EFS) to support high availability and message durability.

- **Industry-standard APIs and Protocols**: ActiveMQ supports a wide range of clients including Java Message Service (JMS) 1.1, .NET Message Service (NMS), and a range of other languages including Node.js, Go, Python, Ruby, and C++. ActiveMQ also supports wire-level protocols including AMQP, STOMP, OpenWire, WebSocket, and MQTT.
- **JMS Messaging Features**: ActiveMQ provides all the standard JMS features including point-to-point (message queues), publish-subscribe (topics), request/reply, persistent and non-persistent modes, JMS transactions, and distributed (XA) transactions.

64.1.2 Benefits

- **Migrate with flexible configurations**: As a fully managed service for open-source message brokers, Amazon MQ helps to migrate your applications more easily with updated endpoints and industry-standard APIs and protocols.
- **Low maintenance**: Offload time-consuming tasks such as administration, maintenance operations, and security management for message brokers.
- **Availability**: Ensure high availability of your applications and durability of messages across AWS Availability Zones.

64.2 Backup/Restore and Disaster Recovery

The services provides backup and recovery. Amazon MQ for ActiveMQ offers a Cross-Region data replication (CRDR) feature for recovery that allows for asynchronous message replication from the primary broker in a primary AWS Region to the replica broker in a replica Region. Users can create primary and replica brokers for asynchronous data replication from the primary. Primary and replica brokers offer cross-Region data recovery.

64.3 Pricing Overview

Please see the AWS UK G Cloud 14 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/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/

64.5 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.
- REST API Reference: Describes all REST APIs for Amazon MQ and provides HTTP methods, example request and response schemas, and a full list of parameters.

- Amazon MQ in the AWS CLI Reference: Describes the AWS CLI commands that you can use to work with message brokers.

65 Amazon Neptune

65.1 Service Overview

Amazon Neptune is a service that includes a graph database engine, graph analytics database engine, graph machine learning (ML), and visualization tools, which can be used individually or together. The Neptune service make it easy to work with graph data on AWS. With Amazon Neptune Database, you can scale your graphs with more than 100,000 queries per second for the most demanding applications using a serverless graph database designed for superior scalability and availability. With Amazon Neptune Analytics, you can get insights and find trends by quickly processing large amounts of graph data. You can get results in seconds by invoking popular graph analytic algorithms.

65.1.1 Features

- **Serverless option**: Amazon Neptune Database Serverless is an on-demand deployment option that automatically adjusts database capacity based on an application's needs. Neptune Database Serverless can scale graph database workloads instantly to hundreds of thousands of queries. Neptune Serverless adjusts capacity to provide just the right amount of database resources that the application needs, and you pay only for the consumed capacity, saving up to 90% in database costs compared to peak capacity.
- **High throughput and low latency for graph queries**: Neptune is a purpose-built, high-performance graph database. 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 Database, you can use either Gremlin, openCypher, or SPARQL to run powerful queries that are easy to write and perform well. With Neptune Analytics, you can use openCypher.
- **Easy scaling of database compute resources**: With a few steps in the AWS Management Console, you can scale the compute and memory resources powering your production cluster up or down. With Neptune Database, you can scale by creating new replica instances of the desired size or by removing instances. Compute scaling operations typically complete in a few minutes.
- **Storage that automatically scales**: Neptune Database uses a distributed and shared storage architecture that will automatically grow as your database storage needs grow. Neptune data is stored in a cluster volume that has multi-AZ high availability. When a Neptune DB cluster is created, it is allocated a single segment of 10 GiB. As the volume of data increases and exceeds the currently allocated storage, Neptune automatically expands the cluster volume by adding new segments.
- **Low-latency read replicas**: With Neptune Database, you can increase read throughput to support high-volume application requests by creating up to 15 database read replicas.
- **Instance monitoring and repair**: The health of your Neptune database and its underlying EC2 instance is continually 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.
- **Multi-AZ deployments with read replicas**: On instance failure, Neptune automates failover to 1 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**: For Neptune Database, each 10 GiB chunk of your database volume is made durable across three Availability Zones. Neptune Database 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 Database storage is also self-healing—data blocks and disks are continually scanned for errors and replaced automatically.
- **Automatic, continuous, incremental backups and point-in-time restore**: Backup capability in Neptune Database 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 5 minutes. Your automatic backup retention period can be configured up to 35 days. Automated backups are stored in Amazon S3, which is designed for 99.999999999% durability. Neptune backups are automatic, incremental, and continual 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 use the automated incremental snapshots to reduce the time and storage required. You can create a new instance from a database snapshot whenever you desire.
- **Global database**: Amazon Neptune Global Database is designed for globally distributed applications, allowing a single Neptune database to span multiple Regions. It replicates the graph data with little impact to database performance, enables fast local reads with low latency in each Region, and provides disaster recovery in case of region-wide outages.
- **Network isolation**: Neptune Database runs in Amazon Virtual Private Cloud (Amazon VPC), which allows you to isolate your database in your own virtual network and connect to your on-premises IT infrastructure using industry standard, encrypted IPsec VPNs. In addition, by using the Neptune VPC configuration, you can configure firewall settings and control network access to your database instances.
- **Resource-level permissions**: Amazon Neptune is integrated with AWS Identity and Access Management (IAM) and provides you with the ability to control the actions that your IAM users and groups can take on specific Neptune resources including Database Instances, Database Snapshots, Database Parameter Groups, Database Event Subscriptions, and Database Options Groups.
- **Fine-grained access control**: Neptune provides fine-grained access to users retrieving Neptune data plane APIs with IAM for performing graph-data actions, such as reading, writing, and deleting data from the graph, and non-graph-data actions, such as starting and monitoring Neptune ML activities and checking the status of ongoing data plane activities.
- **Encryption**: Neptune supports encryption in transit with TLS version 1.2. Neptune allows you to encrypt your databases using keys you create and control through AWS Key Management Service (AWS 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.
- **Advanced auditing**: Amazon Neptune allows you to log database events with minimal impact on database performance. Logs can later be analyzed for database management, security, governance, regulatory compliance and other purposes. You can also monitor activity by sending audit logs to Amazon CloudWatch.
- **ML**: Amazon Neptune ML is powered by SageMaker, which uses GNNs, an ML technique purpose-built for graphs, to make fast and more accurate predictions using

graph data. With Neptune ML, you can improve the accuracy of most predictions for graphs by over 50% when compared to making predictions using non-graph methods.

- **Generative AI**: Vector search makes it easy for you to build ML augmented search experiences and generative artificial intelligence (generative AI) applications. You should use vector search if you want to build generative AI applications that combine data in an application domain and similarity search on vector embeddings. Vector search over graph data gives you an overall lower total cost of ownership and simpler management overhead because you do not need to manage separate data stores, build pipelines, or worry about keeping the data stores in sync. Customers building generative AI applications can use vector search to augment their large language models (LLMs) by integrating graph queries for domain-specific context with the results from low-latency, nearest neighbor similarity search on embeddings imported from LLMs hosted in Amazon Bedrock, GNNs in GraphStorm, or other sources. Neptune is integrated with LangChain, an open-source Python framework that makes it easier to develop generative AI applications using LLMs.

65.1.2 Benefits

- **Easier to use**: You can get started with Neptune by launching a new Neptune database instance or Neptune Analytics graph using the AWS Management Console. Neptune database instances are preconfigured with parameters and settings appropriate for the database instance class you have selected. You can launch a database instance and connect your application within minutes without additional configuration.
- **Easier to operate**: Neptune makes it easier to operate a high-performance graph database. With Neptune, you do not need to create custom indexes over your graph data. Neptune provides timeout and memory usage limitations to reduce the impact of queries that consume too many resources.
- **Monitoring and metrics**: Neptune provides Amazon CloudWatch metrics for your database instances. You can use the AWS Management Console to view over 20 key operational metrics for your database instances, including compute, memory, storage, query throughput, and active connections.
- **Automatic software patching**: Neptune will keep your database up to date with the latest patches. You can control if and when your instance is patched through Database Engine Version Management.
- **Database event notifications**: Neptune can notify you by email or SMS of important database events such as automated failover. You can use the AWS Management Console to subscribe to different database events associated with your Neptune databases.
- **Fast database cloning**: Neptune supports quick, efficient cloning operations, where entire multi-terabyte database clusters can be cloned in minutes. You can clone a Neptune database with just a few steps in the AWS Management Console, without impacting the production environment. The clone is distributed and replicated across three Availability Zones.
- **Pay only for what you use**: There is no up-front commitment with Neptune; you pay an hourly charge for each instance that you launch or the database resources you consume for serverless. When you’re finished with a Neptune database instance, you can delete it. You do not need to overprovision storage as a safety margin, and you only pay for the storage you actually consume.
- **Developer productivity**: Amazon Neptune supports the Property Graph model using the open source Apache TinkerPop Gremlin traversal language and provides a

Gremlin Websockets server that supports TinkerPop version 3.3. With Neptune, you can quickly build fast Gremlin traversals over property graphs. Amazon Neptune supports the W3C's Semantic Web standards of RDF 1.1 and SPARQL 1.1 (Query and Update), and provides an HTTP REST endpoint that implements the SPARQL Protocol 1.1. Neptune supports building graph applications using openCypher, currently one of the most popular query languages for developers working with graph databases. Developers, business analysts, and data scientists like openCypher's SQL-inspired syntax because it provides a familiar structure to compose queries for graph applications. Neptune Analytics supports algorithms for path finding, detecting communities (clustering), identifying important data (centrality), and quantifying similarity.

65.2 Backup/Restore and Disaster Recovery

The service provides data backup and 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.

65.3 Pricing Overview

Please see the AWS UK G Cloud 14 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/neptune/index.html
- **Service quotas:** https://docs.aws.amazon.com/general/latest/gr/neptune.html
- **Service FAQs:** https://aws.amazon.com/neptune/faqs/

65.5 Technical Requirements

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

- Neptune Database User Guide: Provides an overview of the Amazon Neptune graph database, including instructions for using its features with the Gremlin, openCypher, and SPARQL query languages. It also documents Neptune administrative and data-access API operations.
- Neptune Analytics User Guide: Provides an overview of the Neptune Analytics graph database, including instructions for using its features with the openCypher query language. It also documents Neptune Analytics administrative and data-access API operations.

66 Amazon OpenSearch Service

66.1 Service Overview

Amazon OpenSearch Service is a managed service that makes it easy for you to perform interactive log analytics, real-time application monitoring, website search, and more. OpenSearch is an open source, distributed search and analytics suite derived from Elasticsearch. Amazon OpenSearch Service offers the latest versions of OpenSearch,

support for 19 versions of Elasticsearch (1.5 to 7.10 versions), as well as visualization capabilities powered by OpenSearch Dashboards and Kibana (1.5 to 7.10 versions). Amazon OpenSearch Service currently has tens of thousands of active customers with hundreds of thousands of clusters under management processing trillions of requests per month. See the Amazon OpenSearch Service FAQ for more information.

66.1.1 Features

- **Managed OpenSearch**: Focus on analysis instead of spending time managing your deployment, and adjusting deployment configurations as requirements change—while using the power of open source search. 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.
- **Secure**: Meet and maintain high security for authentication, authorization, encryption, audit, and regulatory compliance. You can securely connect your applications to your managed Elasticsearch (version 7.10 and previous) or OpenSearch environment from your Amazon Virtual Private Cloud (Amazon VPC) or through the public internet, configuring network access using VPC security groups or IP-based access policies. You can also securely authenticate users and control access using Amazon Cognito, AWS Identity and Access Management (IAM), or basic authentication with a username and password. 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 preconfigured thresholds.
- **Observability**: Deliver log and trace analytics solutions on data at scale while developing interactive queries and visualizing results with high adaptability and speed. OpenSearch Service provides out-of-the-box support for over 2,200 open-source Sigma security rules to detect potential security threats by filtering through the security findings. You can even customize or use default Sigma rules to rapidly detect potential security threats and send alerts to a preselected destination. Use out-of-the-box support for multiple log sources. including Windows, NetFlow, AWS CloudTrail, DNS, and more.
- **Cost-conscious**: Integrate fast, scalable full-text search capabilities. Manage growing analytics costs for hot, UltraWarm, and cold tiers. All features are included without upsell. UltraWarm helps you cost-effectively expand the data that you want to analyze on OpenSearch Service. You can gain valuable insights on data that previously may have been deleted or archived. With UltraWarm, you can now economically retain more of your data to interactively analyze it whenever you want.

66.1.2 Benefits

- **Ease of operations**: Increase operational excellence by using a popular open source solution, managed by AWS. Optimize time and resources for strategic work.
- **Secure**: Audit and secure your data with a data center and network architecture with built-in certifications.
- **Machine learning**: Systematically detect potential threats and react to a system's state through machine learning, alerting, and visualization.

66.2 Backup/Restore and Disaster Recovery

The service provides backup and recovery of block volumes, cluster indexes, and application state. Automated and manual snapshots can be created via manual or scheduled API calls, or using the OpenSearch Management Console. The AWS Management Console also provides users with the option of configuring a multi-AZ domain in Amazon OpenSearch Service.

66.3 Pricing Overview

Please see the AWS UK G Cloud 14 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/opensearch-service/
- **Service quotas:** https://docs.aws.amazon.com/general/latest/gr/opensearch-service.html
- **Service FAQs:** https://aws.amazon.com/opensearch-service/faqs/

66.5 Technical Requirements

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

- Overview of Amazon OpenSearch Service: Learn about the basic concepts and features in OpenSearch 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.

67 Amazon Personalize

67.1 Service Overview

Amazon Personalize is a fully managed machine learning (ML) service that uses your data to generate product and content recommendations for your users. You provide data about your end-users (e.g., age, location, device type), items in your catalog (e.g., genre, price) and interactions between users and items (e.g., clicks, purchases). Personalize uses this data to train custom, private models that generate recommendations which can be surfaced via an API.

Personalize is powered by the same ML technology used by Amazon.com and allows any developer to easily add personalization to existing applications, websites, push notifications, marketing communications, and more – all without any ML experience required. Personalize uses real-time data insights to deliver recommendations that are personalized instantaneously depending on the user's behavior. You can get started quickly with use-case optimized recommenders for your business domain, or you can create your own configurable custom resources.

67.1.1 Features

- **Create recommendations and next best action**: Next-Best-Action (aws-next-best-action) recipe generates recommendations for actions that your users are most likely to take based on their past behavior. Use Next-Best-Action recipe to recommend high-value actions such as enrolling in loyalty programs, signing up for a newsletter, exploring a new category, or downloading an app. Amazon Personalize offers intelligent user segmentation so you can run more effective prospecting campaigns through your marketing channels. You can map your data to a recommender that applies to your business need while Amazon Personalize chooses the optimal settings for your use case and automates the work of creating and maintaining personalized recommendations, accelerating time to market.

- **User personalization**: The User-Personalization (aws-user-personalization) recipe is optimized for all personalized recommendation scenarios. It predicts the items that a user will interact with based on interactions, items, and user datasets. When recommending items, it uses automatic item exploration to improve discovery and engagement.
- **Similar item recommendations**: Improve the discoverability of your catalog by surfacing similar items to your users. The Similar-Items (aws-similar-items) generates recommendations for items that are similar to an item you specify. Use Similar-Items to help customers discover new items in your catalog based on their previous behavior and item metadata. Recommending similar items can increase user engagement, click-through rate, and conversion rate for your application.
- **Personalized rankings**: Personalized ranking is a list of recommended items that are re-ranked for a specific user. This is useful if you have a collection of ordered items, such as search results, promotions, or curated lists, and you want to provide personalized re-ranking for each of your users. Amazon Personalize enables you to highlight anything that could be relevant to your users while achieving your business priorities and ensuring the best customer experience.
- **New item recommendations**: One of the most challenging problems in building relevant recommendations is offering the right recommendations when new items are added to your catalog. Amazon Personalize allows you to generate quality recommendations for new products and fresh content by creating the appropriate balance between recommendations for new and existing items in your catalog. Learn more.
- **Trending now**: The Trending-Now recommender helps customers recommend items gaining popularity at the fastest pace among their users. Customers can also define the frequency at which it identifies trending items, with options for refreshing recommendations every 30 mins, 1 hour, 3 hours or 1 day, based on the most recent interactions data from users.
- **Personalized search**: Personalize your search results with Amazon Personalize and OpenSearch integration. Search is crucial in engaging users as it brings high-intent traffic from individuals seeking specific products. Improve the relevance of your search results by incorporating the unique interests and needs of each user through ML-based intelligence from Amazon Personalize. Using the Amazon Personalize Search Ranking plugin within OpenSearch v2.9 and above, you can boost relevant items in a specific user's search results based on their interests, context, and past interactions in real-time.
- **Business rules and filters**: Apply business rules to deliver the optimal customer experience. Amazon Personalize enables you to automatically augment recommendations.
- **Promotions**: Promote specific items or content based on rules that align with your business goals. With this feature, you can control the percentage of promoted content within your recommendations to further customize each user’s experience. Amazon Personalize automatically finds the most relevant items or content to be promoted for each user within the business rule provided and distributes it within the user's recommendations.
- **Unstructured text support**: 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 catalog, and Amazon Personalize automatically extracts key information to use when generating recommendations. Supported languages include Chinese (Simplified and Traditional), English, French, German, Japanese, Portuguese, and Spanish.

Use or disclosure of data contained on this sheet is subject to the restriction on the title page of this document.

- **Real-time or batch recommendations**: Amazon Personalize provides flexibility to use real-time or batch data based on what is most appropriate for your use case. For example, real-time data may be more appropriate for product or content recommendations on a website or app. Make your recommendations relevant by responding to the changing intent of your users in real-time. Batch data may be more appropriate for large notification campaigns. For instance, you can 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.
- **Contextual recommendations**: Providing relevant recommendations requires you to consider the context in which they are bring viewed. With contextual recommendations, you can deliver a more personalized experience for customers and improve the relevance of recommendations by generating them within a context, for instance device type, time of day, and more.
- **Metric aggregation**: Amazon Personalize enables customers to automatically understand the impact of Personalize on their business objectives such as cart-adds, page views, and clicks. Customers can measure the business outcome of any Personalize recommendation by calculating the impact of any event sent to the system. When a user completes an action (i.e. event), that data is sent to Personalize and the total impact is calculated.
- **Generative AI capabilities**: Amazon Personalize Content Generator uses generative AI to make recommendations more compelling with text generated by foundation models.
- **LangChain integration**: Builders can use a custom chain on LangChain to seamlessly integrate Amazon Personalize with generative AI solutions. With pre-configured LangChain code, you can invoke Amazon Personalize, retrieve recommendations for a campaign or a recommender, and seamlessly feed it into your generative AI applications within the LangChain ecosystem.
- **Return metadata in inference response**: Amazon Personalize improves your generative AI workflow by enabling return item metadata as part of the inference output. Select up to 10 fields, such as genre, rating, and product description, and use our LangChain integration capability to seamlessly feed these enriched recommendations into the foundation models. This richer context helps models generate highly personalized content to boost your engagement with end users.

67.1.2 Benefits

- **Unified data for seamless customer experiences**: Unify your data to create meaningful customer experiences across the entire user journey.
- **Tailor experiences to customer preferences**: Increase revenue and brand loyalty, and stand out from competitors by catering to individual customer preferences.
- **Rapid customized personalization**: Quickly implement a customized personalization engine in days—not months—with no ML expertise required.
- **Dynamic recommendations for personalized experiences**: Adapt recommendations in real time for relevant customer experiences, new users, or new catalog items.

67.2 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/.

67.3 Pricing Overview

Please see the AWS UK G Cloud 14 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/personalize/
- **Service quotas:** https://docs.aws.amazon.com/personalize/latest/dg/limits.html
- **Service FAQs:** https://aws.amazon.com/personalize/faqs/

67.5 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.

68 Amazon Pinpoint

68.1 Service Overview

Amazon Pinpoint is self-service so you can try before you buy, without talking to sales, completing RFPs, or hiring for consultants or professional services.

68.1.1 Features

- **Messaging channels**: Amazon Pinpoint helps you engage with your customers through SMS, push notifications, in-app messaging, and email.
- **SMS**: Amazon Pinpoint SMS is a cloud-based application-to-person (A2P) text messaging offering that provides a cost-effective, flexible, and scalable way for businesses to connect with their customers through SMS. Send text messages with Pinpoint SMS using its dedicated tools, SDK, and console.
- **Manage global scale and high throughput**: Send SMS messages to customers in over 240 countries with a variety of supported sender types.
- **Create resiliency**: Use phone pools to seamlessly manage all your sending resources in one place and send messages with failover capabilities.
- **Self-service registration**: Register and request various sender types directly in the Pinpoint SMS console including US 10DLC and toll-free numbers.
- **Integrate SMS messaging into your apps and workflows**: Stream SMS events to CloudWatch for monitoring, Kinesis for analytics, and SNS to fan out SMS event data to any system you choose.
- **Send and receive messages with two-way SMS**: With Amazon Pinpoint SMS, you can enable two-way messaging allowing your customers to send messages with an automatically triggered response.
- **Test SMS sending**: Use the SMS simulator in the Amazon Pinpoint SMS console to get setup quickly and start testing sending in a matter of minutes.
- **Push notifications**: Push notifications alert an app or website user to news, updates, or useful information through messages on a device's home screen or lock screen. Amazon Pinpoint supports the following channels: Firebase Cloud

Messaging (FCM), Apple Push Notification Service (APNs), Baidu Cloud Push, and Amazon Device Messaging (ADM).

- **In-app messaging**: Send targeted messages to users inside your applications. Amazon Pinpoint in-app messages are highly customizable and have configurable background and text colors, text positioning, and custom buttons and images. Use in-app messages to notify your users of news, updates, or useful information inside your app or website.
- **Email**: Amazon Pinpoint connects with Amazon Simple Email Service (SES) to send email.
- **Campaign orchestration**: Amazon Pinpoint's campaign orchestration tools serve as your communication hub, bringing together your channels, segments, personalization, and campaigns and journeys into a single tool. Create engaging customer experiences based on your customer data from one central place.
- **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. You can also create static segments by importing CSV or JSON files that contain existing first- or third-party customer data.
- **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**: Amazon Pinpoint campaigns allow you to create and execute omnichannel communications when your customers perform specific actions. Build campaigns to serve static or dynamic customer segments, then send personalized messages in real-time through channels like SMS, push notifications, and email when customers complete an action, like checking out, abandoning a cart, or signing up for marketing communications.
- **Journeys**: With Amazon Pinpoint journeys, you can create multi-step campaigns for your customers. Build journeys with multiple touchpoints, and target customers with specific communications based on their actions.
- **Web and mobile 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. You can use this data to create both dynamic and static segments and take action in a campaign.
- **Transactional analytics 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.

68.1.2 Benefits

- **Benefit for marketers**: The console provides marketers with campaign management tools to create, run, and manage multi-channel campaigns across their applications, user-base and devices. Campaigns can be scheduled or triggered on user changes and actions. For marketers that want to run multi-step campaigns across multiple

channels, they can design journeys to orchestrate an end-to-end experience. Marketers can also leverage the templating support to personalize end-user messaging. Marketers can also measure messaging effectiveness using Pinpoint analytics to understand the impact on user behavior.

- **Engage customers**: Personalize and segment communications to engage customers across SMS, email, push notifications, in-app messaging, and voice.
- **Expand your reach**: With support for over 240 countries and regions, automate campaigns and journeys to send customers the right message at the right time.
- **Connect with targeted customer groups**: Send real-time notifications to customer segments about breaking news, local updates, or event notices. Aggregate your communications into a single messaging hub.

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 14 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/pinpoint/
- **Service quotas:** https://docs.aws.amazon.com/personalize/latest/dg/limits.html
- **Service FAQs:** https://aws.amazon.com/pinpoint/faqs/

68.5 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.
- API Reference for Amazon Pinpoint: Describes all the Amazon Pinpoint REST API actions.

69 Amazon Polly

69.1 Service Overview

Amazon Polly is a service that turns text into lifelike speech. Amazon Polly enables existing applications to speak as a first class feature and creates the opportunity for entirely new categories of speech-enabled products, from mobile apps and cars, to devices and appliances. Amazon Polly includes dozens of lifelike voices and support for multiple languages, so you can select the ideal voice and distribute your speech-enabled applications in many geographies.

Amazon Polly is easy to use – you just 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 you can play it directly or store it in a standard audio file format, such as MP3. Amazon Polly supports Speech Synthesis Markup Language (SSML) tags like prosody so you can adjust the speech rate, pitch, or volume.

Amazon Polly is a secure service that delivers all of these benefits at high scale and at low latency. You can cache and replay Amazon Polly's generated speech at no additional cost. Amazon Polly lets you convert millions of characters per month for free during the first year, upon sign-up. Amazon Polly's pay-as-you-go pricing, low cost per request, and lack of restrictions on storage and reuse of voice output make it a cost-effective way to enable speech synthesis everywhere.

69.1.1 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.
- **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.
- **Newscaster Speaking Style**: Amazon Polly can be used to synthesize speech as if it is were spoken by a TV or Radio newscaster. This can be a great way to read news articles or deliver flash briefing updates. The Newscaster style is currently available for the US English (en-US) Matthew and Joanna voices, British English (en-GB) Amy and US Spanish (es-US) Lupe voice using Neural text-to-speech. Listen to an audio sample in US English, British English or US Spanish.
- **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. Brand Voice allows you to differentiate your products and applications with a unique vocal identity in a wide variety of use cases, including Amazon Connect and Alexa Skills integrations.
- **Contact center integration with Amazon Connect**: Amazon Polly is natively integrated with Amazon Connect, AWS' cloud-based contact center solution that you use to set up and manage a customer contact center and provide reliable customer engagement at any scale.
- **Contact center integration with Genesys Cloud CX**: Genesys Cloud CX is a cloud contact center solution that unifies customers and agent experiences across multiple channels such as phone, text and chat. You can deploy your voice bots using any of the existing Polly voices.
- **Contact center integration with Amazon Chime SDK**: The Amazon Chime SDK is a set of real-time communications components that supports native integration with Amazon Polly, making it easy for builders to create applications that turn text and numerical data into lifelike speech and automatically play the output to a phone caller.
- **Contact center integration with AWS Contact Center Intelligence (CCI)**: Amazon Polly is used by several AWS CCI partners, so you can seamlessly create self-service customer service virtual agents, informational bots or application bots. Amazon Polly partners include Genesys, Vonage, and Accenture.

69.1.2 Benefits

- **Generate speech in dozens of languages**: Add speech to applications with a global audience, such as RSS feeds, websites, or videos.
- **Engage customers with a natural-sounding voice**: Store and replay Amazon Polly speech output to prompt callers through interactive or automated voice response systems. The service delivers lifelike voices and conversational user experiences in consistently fast response times.
- **Adjust speaking style, speech rate, pitch, and loudness**: Use SSML, a W3C standard XML-based markup language for speech synthesis applications, to support common SSML tags for phrasing, emphasis, and intonation.
- **Control speech output**: Customize and control speech output that supports lexicons and Speech Synthesis Markup Language (SSML) tags. Store and redistribute speech in standard formats like MP3 and OGG.

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

69.5 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.

70 Amazon Quantum Ledger Database

70.1 Service Overview

Amazon Quantum Ledger Database (Amazon QLDB) is a fully managed ledger database that provides a transparent, immutable, and cryptographically verifiable transaction log.

70.1.1 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 (for example, list of all previous owners of a vehicle), and also specific details related to transaction history (for example, 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.
- **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. Blockchain frameworks are de-centralized and require peer nodes to validate a transaction before it can be stored in the ledger, impacting their performance.

- **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.
- **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. This allows you to react quickly to new events (such as a change in account balance for a banking ledger application) and easily integrate with downstream services (such as AWS Lambda, Amazon Redshift, and Amazon OpenSearch Service).

70.1.2 Benefits

- **Application history**: 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.
- **Event-driven**: Build correct, event-driven systems with QLDB ACID transactions and support for real-time streaming to Amazon Kinesis.
- **Serverless architecture**: Start small and pay only for what you use with serverless architecture that provides automatic storage and resource scaling.

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

70.5 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.
- Amazon QLDB in the AWS CLI Reference: Describes the command line interface for Amazon QLDB in detail. Provides basic syntax, options, and usage examples for each command.

71 Amazon QuickSight

71.1 Service Overview

Amazon QuickSight powers data-driven organizations with unified business intelligence (BI) at hyperscale. With QuickSight, all users can meet varying analytic needs from the same source of truth through modern interactive dashboards, paginated reports, natural language queries and embedded analytics. With Amazon Q in QuickSight, business analysts and business users can use natural language to build, discover, and share meaningful insights in seconds, turning insights into impact faster.

71.1.1 Features

- Serverless architecture and simplified management: In addition to the QuickSight serverless architecture, with QuickSight APIs, you can programmatically manage your analytics workflows, such as setting up multitenant architecture for isolating data between multiple end users, moving dashboards across accounts, automating deployments, and managing access for users with single sign-on (SSO).
- Deeper data insights: QuickSight empowers your users to gain deeper insights through as-needed analysis and machine learning (ML) capabilities such as anomaly detection, forecasting, and natural language querying (NLQ). In addition, users can author, publish, and share their own dashboards and reports to get specific answers to their customized questions. These capabilities provide you with new ways of monetizing and differentiating your applications.
- Custom-branded email reports: QuickSight Enterprise Edition now supports the ability to send custom-branded email reports. You can customize the email sender domain for email reports sent from QuickSight, along with the logo, header color, and footer text of the email. If you have your dashboard embedded in your own application, you can also customize the URL to open the dashboard from the email to the URL of your application.
- One-step embedding: QuickSight Enterprise Edition now supports one-step embedding, a feature that allows you to embed rich, interactive dashboards in applications, wikis, and portals without needing to call embedding APIs. Authorized end users can start accessing these dashboards instantly without any server deployments or infrastructure licensing needed. One-step embedding allows you to give your users insights in minutes.
- Built-in security and compliance: QuickSight provides built-in security features so that you can distribute dashboards and insights more securely to tens of thousands of users. In addition to the multi-Region availability and built-in redundancy, you can

use QuickSight to manage your users and content through a comprehensive set of security features, including role-based access controls (RBACs), Active Directory integration, AWS CloudTrail auditing, single sign-on (such as AWS Identity and Access Management [IAM] and third parties), VPC subnets, and data backup.

- Mobile app support: Amazon QuickSight Mobile for iOS and Android helps you securely get insights from your data from anywhere. Favorite, browse, and interact with all your dashboards in a straightforward mobile-optimized experience. You can explore your data by drilling down and filtering, stay ahead of the curve through 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.
- Extensive API capabilities: QuickSight now offers expanded API capabilities that make it easier to create and manage users, datasets, and assets. QuickSight APIs provide programmatic access to BI assets such as dashboards, analysis, and reports. This API access to the underlying data models allows these BI assets to be treated as software source code and managed in your DevOps pipelines. You can use these APIs to programmatically migrate from legacy BI solutions to QuickSight, thereby accelerating the BI transition to the cloud.
- SPICE (Super-fast, Parallel, In-memory Calculation Engine): With SPICE, the QuickSight 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, which saves you time and resources.
- 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. You no longer need to apply software patches or upgrade hardware to meet increasing demand.
- Amazon Q in QuickSight: The new Amazon Q dashboard authoring capabilities in QuickSight make it easier than ever before to analyze data and build interactive dashboards. Amazon Q in QuickSight makes it easier for business users to use existing dashboards to better inform business stakeholders, distill key insights, and simplify decision-making.
- Broad data source support: QuickSight allows you to directly connect to and import data from various cloud and on-premises data sources. These include software-as-a-service (SaaS) applications such as Salesforce, Square, ServiceNow, Twitter, GitHub, and Jira; third-party databases such as Teradata, MySQL, Postgres, and SQL Server; native AWS services such as Amazon Redshift, Amazon Athena, Amazon Simple Storage Service (Amazon S3), Amazon Relational Database Service (Amazon RDS), and Amazon Aurora; and virtual private cloud (VPC) subnets. You can also upload various file types including Excel, CSV, and JSON.

71.1.2 Benefits

- **Unify business intelligence siloes**: Amazon QuickSight powers data-driven organizations with unified business intelligence (BI) at hyperscale. With QuickSight, all users can meet varying analytic needs from the same source of truth through modern interactive dashboards, paginated reports, embedded analytics, and natural language queries. With Amazon Q in QuickSight, business analysts and business users can use natural language to build, discover, and share meaningful insights in seconds, turning insights into impact faster.
- **Reinvent business intelligence**: Amazon Q in QuickSight enhances business productivity using Generative BI capabilities to accelerate decision-making. With new dashboard authoring capabilities in Amazon Q, business analysts can use natural

language prompts to build, discover, and share meaningful insights in seconds. Amazon Q makes it easier for business users to understand data with executive summaries, a new context-aware data Q&A experience, and data stories.

- **Create and deliver paginated reports**: Create, schedule, and share reports and data exports from a single fully managed cloud-based BI service. Get business-critical information to users how and when they need it with critical operational reports and dashboards in the same solution.Embed analytics to differentiate your applications.
- **Embed analytics to differentiate your applications**: Embed interactive visualizations and dashboards, sophisticated dashboard authoring, or natural language query capabilities in your applications to differentiate user experiences and unlock new monetization opportunities. Blend analytics seamlessly into your application with Amazon QuickSight Embedded Analytics themes, which help you personalize the look and feel (for example, colors and fonts) of your reports and dashboards.
- **BI for everyone**: Meet needs of all your users from the same source of truth with interactive dashboards, paginated reports, embedded analytics, and natural language queries.
- **Build faster**: Speed up development by using one unified authoring experience to build and share insights across modern dashboards, paginated reports, and embedded insights.
- **Scale efficiently**: Automatically scale to tens of thousands of users without the need to set up, configure, or manage your own servers.
- **Lower costs**: Pay only for what you need with QuickSight usage-based pricing. No need to buy thousands of end-user licenses for large-scale BI or embedded analytics deployments.

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 14 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/quicksight/
- **Service quotas:** https://docs.aws.amazon.com/general/latest/gr/quicksight.html
- **Service FAQs:** https://aws.amazon.com/quicksight/resources/faqs/

71.5 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.

- API Reference: Provides descriptions, syntax, example requests, and sample responses for all of the operations available in Amazon QuickSight, plus details on data types and errors.
- Amazon QuickSight in the AWS CLI Reference: Describes all the command line interface (CLI) commands for Amazon QuickSight in detail. Provides all syntax. Also provides examples for the most common commands.

72 Amazon Redshift

72.1 Service Overview

Amazon Redshift Serverless automatically provisions and scales data warehouse capacity to deliver high performance for demanding and unpredictable workloads. Amazon Redshift offers leading price performance for diverse analytics workloads, whether it is dashboarding, application development, data sharing, ETL (Extract, Transform, Load) jobs or several others. With tens of thousands of customers running analytics on terabytes to petabytes of data, Amazon Redshift optimizes real-world customer workload performance, based on fleet performance telemetry, and delivers performance that scales linearly to the workload, while keeping costs low. Amazon Redshift lets you get insights from running real-time and predictive analytics on all your data across your operational databases, data lake, data warehouse, streaming data, and third-party datasets.

72.1.1 Features

- **Innovate faster with secure data collaboration**: Share data securely across AWS regions, teams, and third-party data warehouses without data movement or data copying. In just a few clicks multiple teams can read and update shared data sets and collaborate on up-to-date data across regions, accounts, and even third party data warehouses. Data sharing is centrally governed by AWS Lake Formation. Have the confidence that your data is secure no matter where you operate or how highly regulated your industries are. Amazon Redshift enables fine grained access controls such as role-based access controls, row and column level security and an easy authentication experience with single sign-on for your organizational identity.
- **Achieve best price-performance at any scale**: Meet your needs for a highly scalable, performant, and reliable modern cloud data warehouse to handle growing data for any number of concurrent users. With the best price-performance for your workloads, Amazon Redshift runs on Massively Parallel Processing (MPP) architecture and RA3 instances that separate compute and storage. Run and scale any kind of analytics workloads cost effectively without managing data warehouse infrastructure using Amazon Redshift Serverless, complete with AI-driven scaling and optimizations. As you adjust to demanding analytics needs for your business, a reliable cloud data warehouse such as Amazon Redshift becomes imperative for you to minimize disruptions with its multi-Az deployments delivering 99.99% SLAs.
- **Unify all your data with zero-ETL approaches**: Break through data silos in your organization and build an end to end data strategy to analyze all your data. Amazon Redshift employs a zero-ETL approach that enable interoperability and integration between the data warehouse, your Amazon S3 data lakes, operational and NoSQL databases such as Amazon Aurora, Amazon RDS, and Amazon DynamoDB, and even your streaming data services, so that the data is easily and automatically ingested into the warehouse for you or you can access the data in place. No more spending weeks or months building cumbersome and erroneous data pipelines to move the data from one system to another.
- **Maximize value with comprehensive analytics & ML**: From executing SQL queries to building complex dashboards or near real-time and AI/Gen-AI applications,

Amazon Redshift makes it easy for you to analyze all your data and drive your business forward. You can spin up a Redshift Serverless endpoint in seconds, use Amazon Redshift's Query Editor to load, analyze, visualize, and collaborate on your data, across data sources. Submit query requests in plain English and receive custom SQL code recommendations based on your organizations schema metadata with Amazon Q generative SQL in Query Editor. Seamlessly go from data to predictive analytics with Amazon Redshift ML which uses familiar SQL to build, train, and deploy machine learning or forecasting models, right within the warehouse.

- **Efficient storage and high-performance query processing**: Columnar storage, data compression, and zone maps reduce the amount of I/O needed to perform queries. Along with the industry-standard encodings such as LZO and Zstandard, Amazon Redshift also offers purpose-built compression encoding, AZ64, for numeric and date/time types to provide both storage savings and optimized query performance.
- **Concurrency scaling**: Supports virtually unlimited concurrent users and concurrent queries with consistent service levels by adding transient capacity in seconds as concurrency increases. 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.
- **BI Integration including third party tools**: Run queries within the console or connect SQL client tools, libraries, or data science tools including Amazon QuickSight, Tableau, PowerBI, QueryBook and Jupyter Notebook.
- **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 Amazon 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.
- **Multidimensional data layouts (MDDL)**: A new powerful table sorting mechanism that improves performance of repetitive queries by automatically sorting data based on the incoming query filters (for example: Sales in a specific region). This method significantly accelerates the performance of table scans compared to traditional methods.
- **Multi-AZ**: Expand recovery capabilities by reducing recovery time and guaranteeing capacity to automatically recover with no data loss. A Redshift Multi-AZ data warehouse maximizes performance and value by delivering high availability without having to use standby resources.
- **Data sharing**: 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.
- **AWS Data Exchange Integration**: 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.
- **Unified identity with IAM identity center integration**: Integration with AWS IAM Identity Center enables organizations to support trusted identity propagation between Amazon Redshift, Amazon QuickSight, and AWS Lake Formation. Customers can use their organization identities to access Amazon Redshift in a single sign-on

experience using third party identity providers (IdP), such as Microsoft Entra ID, Okta, Ping, OneLogin, etc. from Amazon QuickSight and Amazon Redshift Query Editor. Administrators can use third-party identity provider users and groups to manage fine grained access to data across services and audit user level access in AWS CloudTrail. With trusted identity propagation, a user's identity is passed seamlessly between Amazon QuickSight, Amazon Redshift, and AWS Lake Formation reducing time to insights and enabling a friction free analytics experience.

- **Support for Data sharing write**: Write to the same databases using multiple data warehouses, without moving or copying the data.
- **Partner console integration**: 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.
- **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.

72.1.2 Benefits

- **Achieve best price-performance at any scale**: Achieve up to 6x better price performance than any other cloud data warehouse, with a fully managed, AI powered, Massively Parallel Processing (MPP) data warehouse built for performance, scale, and availability.
- **Unify all your data with zero-ETL**: Easily access or ingest data across your data lakes, databases, data warehouses, streaming data - with no code/low code zero-ETL approach for integrated analytics
- **Maximize value with comprehensive analytics & ML**: Run SQL queries and open source analytics, power dashboards and visualizations, activate near real-time analytics and AI/ML applications with analytics engines and languages of your choice.
- **Innovate faster with secure data collaboration**: Share and collaborate on data easily within and across your organizations, AWS regions, and even 3rd party data sets with no manual data movement or copying and with fine grained governance, security, and compliance.

72.2 Backup/Restore and Disaster Recovery

The service provides backup and recovery of block volumes. Users can perform backups using the AWS Redshift Console, the AWS CLI or the Amazon Redshift API. The snapshots of the Redshift cluster can be automated and manual. Amazon Redshift stores these snapshots internally in Amazon S3 using an encrypted Secure Sockets Layer (SSL) connection.

72.3 Pricing Overview

Please see the AWS UK G Cloud 14 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/redshift/
- **Service quotas:** https://docs.aws.amazon.com/redshift/latest/mgmt/amazon-redshift-limits.html
- **Service FAQs:** https://aws.amazon.com/redshift/faqs/

72.5 Technical Requirements

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

- Amazon Redshift Serverless: If you are a first-time user of Amazon Redshift Serverless, read this section to learn how to create Amazon Redshift Serverless resources, as well as load and query data.
- Querying a data lake: Query data from your Amazon S3 data lake without loading the data into Amazon Redshift tables.
- Redshift Getting Started Guide: Get started with Amazon Redshift Serverless and Amazon Redshift. Set up an account and walk through a simple example.
- Redshift Management Guide: Create and manage Amazon Redshift Serverless resources and Amazon Redshift clusters.
- Redshift Database Developer Guide: Design, build, query, and maintain the databases that make up your data warehouse. Includes syntax for Redshift SQL commands and functions.

73 Amazon Rekognition

73.1 Service Overview

Amazon Rekognition is a service that adds powerful visual analysis to your applications. Rekognition Image lets you build applications to search, verify, and organize images. Rekognition Video lets you extract motion-based context from stored or live stream videos and helps you analyze them.

73.1.1 Features

- Extract information and insights from images and videos: Amazon Rekognition offers pre-trained and customizable computer vision (CV) capabilities to extract information and insights from your images and videos.
- Face liveness: Detect real users and deter bad actors using spoofs in seconds during facial verification. Amazon Rekognition Face Liveness analyzes a short selfie video to detect spoofs presented to the camera, such as printed photos, digital photos, digital videos, or 3D masks, as well as spoofs that bypass the camera, such as pre-recorded or deepfake videos.
- Content moderation: Detect potentially unsafe, inappropriate, or unwanted content across images and videos.
- Labels: Detect objects, scenes, activities, landmarks, dominant colors, and image quality.
- Face compare and search: Determine the similarity of a face against another picture or from your private image repository.

- Custom labels: Detect custom objects such as brand logos using automated machine learning (AutoML) to train your models with as few as 10 images.
- Video segment detection: Detect key segments in videos, such as black frames, start or end credits, slates, color bars, and shots.
- Face detection and analysis: Detect faces appearing in images and videos and recognize attributes such as open eyes, glasses, and facial hair for each.

73.1.2 Benefits

- Add computer vision APIs: Add pre-trained or customizable computer vision APIs to your applications without building machine learning (ML) models and infrastructure from scratch.
- Analyze at scale: Analyze millions of images, streaming, and stored videos within seconds, and augment human review tasks with artificial intelligence (AI).
- Pay only for what you use: Scale resources based on your business needs with fully managed AI capabilities and pay only for the images and videos you analyze.

73.2 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/.

73.3 Pricing Overview

Please see the AWS UK G Cloud 14 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/rekognition/
- **Service quotas:** https://docs.aws.amazon.com/rekognition/latest/dg/limits.html
- **Service FAQs:** https://aws.amazon.com/rekognition/faqs/

73.5 Technical Requirements

Please refer to https://docs.aws.amazon.com/rekognition/ 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.
- API Reference: Describes all the API operations for Amazon Rekognition in detail.

74 Amazon Relational Database Service (RDS)

74.1 Service Overview

Amazon Relational Database Service (Amazon RDS) is a managed service that makes it easy to set up, operate, and scale a relational database in the cloud. It provides cost-efficient and resizable capacity, while managing time-consuming database administration tasks, freeing you to focus on your applications and business. Amazon RDS gives you access to the capabilities of a familiar RDS for PostgreSQL, RDS for MySQL, RDS for MariaDB, RDS for SQL Server, RDS for Oracle, or RDS for Db2 database. This means that the code, applications, and tools you already use today with your existing databases can work seamlessly with Amazon RDS.

Amazon RDS can automatically back up your database and keep your database software up to date with the latest version. You benefit from the flexibility of being able to scale the compute resources or storage capacity associated with your relational database instance.

74.1.1 Features

- **Easy to set up and deploy**: You can use the AWS Management Console, the Amazon RDS Command Line Interface, SDKs, or simple API calls to access the capabilities of a production-ready relational database in minutes. Amazon RDS lowers your administrative burden by managing the administrative tasks, such as software patching, backups, provisioning, and scheduled maintenance jobs.
- **Monitor and view database metrics**: Amazon RDS provides Amazon CloudWatch metrics for your database instances at no additional charge. You can use the Amazon RDS Console to view key operational metrics, including compute/memory/storage capacity utilization, I/O activity, and instance connections.
- **Scale infrastructure with AWS CloudFormation**: Amazon RDS is integrated with AWS CloudFormation, a service that helps you model, provision, and manage your AWS resources across your AWS accounts and Regions. When creating Amazon RDS resources with AWS CloudFormation, you can simplify the management of resources and infrastructure, by reusing your templates and allowing AWS CloudFormation to provision and configure your resources.
- **Turn on event notifications with subscriptions and Amazon SNS**: Amazon RDS can notify you through email or SMS text message of database events through Amazon Simple Notification Services (Amazon SNS). You can use the AWS Management Console or the Amazon RDS APIs to subscribe to over 40 different database events associated with your database instances.
- **Enhance security and compliance with configuration governance**: Amazon RDS integrates with AWS Config to support compliance and enhance security by recording and auditing changes to the configuration of your DB instance, including parameter groups, subnet groups, snapshots, security groups, and event subscriptions.
- **Database cloning**: Amazon Aurora supports quick, efficient cloning operations, where entire multi-terabyte database clusters can be cloned in minutes. Database cloning is useful for a number of purposes including application development, testing, database updates, and running analytical queries. The immediate availability of data can significantly accelerate your software development and upgrade projects, and make analytics more accurate. You can clone an Amazon Aurora database in only a few steps, and you don't incur any storage charges, except if you use additional space to store data changes.

- **Safer, simpler, faster upgrades with Amazon RDS Blue/Green Deployments**: Amazon RDS Blue/Green Deployments allow you to make safer, simpler, and faster database updates with zero data loss. In a few steps, Blue/Green Deployments create a staging environment that mirrors the production environment and keeps the two environments in sync using logical replication. You can make changes—such as major/minor version upgrades, schema modifications, and parameter setting changes—without impacting your production workload.
- **Analyze Amazon RDS data using zero-ETL integrations with Amazon Redshift**: Zero-ETL integrations with Amazon Redshift provide access to analytics and machine learning (ML) capabilities on petabytes of your transactional data from your Aurora MySQL-Compatible Edition, Aurora PostgreSQL-Compatible Edition (Preview), and RDS for MySQL (Preview) databases. Using zero-ETL integrations with Amazon Redshift, you can remove the need to build and manage your own complex data pipelines that perform extract, transform, and load (ETL) operations.
- **Choice of General Purpose and Provisioned IOPS SSD storage options**: Amazon RDS offers a choice of two different storage types. For most use cases, Amazon RDS General Purpose storage will work for your database workloads. This is an SSD-backed storage option that delivers a consistent baseline of 3 IOPS per provisioned GB and provides the ability to burst up to 3,000 IOPS above the baseline. For workloads that require higher performance, Amazon RDS offers Provisioned IOPS storage, which is an SSD-backed storage option designed to deliver fast, predictable, and consistent I/O performance.
- **Add an Amazon ElastiCache cache to Amazon RDS databases**: To accelerate the performance of your Amazon RDS database and applications, you can add a cache from Amazon ElastiCache, a managed, in-memory caching service, to your database directly from the Amazon RDS Console. Together, they can effectively deliver real-time performance of your applications and database needs.
- **Diagnose and resolve performance bottlenecks with Amazon DevOps Guru for RDS**: Amazon DevOps Guru is a cloud operations service powered by machine learning (ML) that helps improve application availability. Developers and DevOps engineers can use DevOps Guru for RDS to automatically identify the root cause of performance issues and get intelligent recommendations to help address the issue, without needing help from database experts.
- **Improve write throughput with Amazon RDS Optimized Writes**: Amazon RDS Optimized Writes, built on top of the new AWS Nitro System Torn Write Prevention Feature, allow you to improve write transaction throughput by up to 2x at no additional cost.
- **Improve query processing with Amazon RDS Optimized Reads**: Amazon RDS Optimized Reads are designed to provide you with faster database performance with up to 2X faster query processing at no additional cost.
- **Achieve high availability with Amazon RDS Multi-AZ deployments**: Amazon RDS Multi-AZ deployments provide enhanced availability and durability for database instances with a SLA of up to 99.95%, making them a natural fit for production database workloads. When you provision a Multi-AZ database instance, Amazon RDS synchronously replicates your data to a standby instance in a different Availability Zone (AZ).
- **Protect data with 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 the use of SSL to secure data in transit, as well Transparent Data Encryption in SQL Server and Oracle.

Use or disclosure of data contained on this sheet is subject to the restriction on the title page of this document.

- **Run database instances in Amazon VPC for network isolation**: AWS recommends that you run your database instances in Amazon Virtual Private Cloud (VPC), which allows you to isolate your database in your own virtual network and connect to your on-premises IT infrastructure using industry-standard encrypted IPsec VPNs.
- **Control actions of AWS users and groups with resource-level permissions**: Amazon RDS supports several ways to authenticate users, including AWS Identity and Access Management (IAM) and Kerberos.Amazon RDS is integrated with IAM and provides you the ability to control the actions that your AWS IAM users and groups can take on specific Amazon RDS resources, from database instances through snapshots, parameter groups, and option groups.
- **Power your database with push-button compute scalability**: You can scale the compute and memory resources powering your deployment up or down, up to a maximum of 128 vCPUs and 4,096 GiB of RAM. Compute scaling operations typically complete in a few minutes. The RDS for MySQL, RDS for MariaDB, RDS for Oracle, and RDS for PostgreSQL engines allow you to scale up to 64 TiB of storage and RDS for SQL Server supports up to 16 TiB. Storage scaling is on the fly with zero downtime.
- **Elastically scale out with Amazon RDS Read Replicas**: Amazon RDS Read Replicas make it easier 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.
- **Increase resiliency and scalability with Amazon RDS Proxy**: With Amazon RDS Proxy, you can improve database efficiency and make your applications more scalable, secure, and resilient to database failures by allowing applications to pool and share connections with the database.
- **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.
- **Receive significant discounts using Amazon RDS Reserved Instances**: Amazon RDS Reserved Instances give you the option to reserve a DB instance for a one- or three-year term and in turn receive a significant discount compared to the On-Demand Instance pricing for the DB instance.
- **Stop and start database instances**: Amazon RDS allows you to easily stop and start your database instances for up to seven days at a time. This makes it easier and affordable to use databases for development and test purposes, where the database is not required to be running all of the time.
- **Safely write extensions with Trusted Language Extensions for PostgreSQL**: Trusted Language Extensions (TLE) for PostgreSQL is a development kit and open-source project that allows you to quickly build high performance extensions and safely run them on Amazon RDS without needing AWS to certify code.
- **Perform vector similarity searches and store embeddings**: With Amazon RDS for PostgreSQL, you can use pgvector, an open-source PostgreSQL extension, to perform vector similarity searches. You can also store embeddings from machine learning (ML) and artificial intelligence (AI) models in your database, such as those from Amazon Bedrock or Amazon SageMaker.
- **Deploy Amazon RDS on-premises with Amazon RDS on AWS Outposts**: AWS Outposts is a fully managed service that extends AWS infrastructure and services to any on-premises or co-location for a hybrid experience. With Amazon RDS on Outposts, you can run Amazon RDS on-premises and receive the same features and

benefits of Amazon RDS as you would in the cloud, including automated administrative tasks.

- **Gain privileged access with Amazon RDS Custom**: Amazon RDS Custom is a managed database service that gives you privileged access to the underlying operating system and database environment to run packaged, custom, and legacy applications for Microsoft SQL Server and Oracle. It's a shared responsibility model in exchange for more flexibility.
- **Perform homogenous and heterogeneous migrations with AWS DMS**: With Amazon RDS, you can use AWS Database Migration Service (AWS DMS), a managed migration and replication service for database and analytics workloads, to perform both homogeneous and heterogeneous migrations.
- **Assess and convert source schema with AWS Schema Conversion Tool**: Unlike homogeneous migrations, heterogeneous migrations require a schema conversion to ensure the source schema is compatible with the new target engine.
- **Migrate using familiar native tools**: For homogeneous migrations, you can also use familiar native database migration tools such as PostgreSQL, MySQL, Microsoft SQL Server, Oracle, Db2 to migrate workloads to AWS. You can also access the latest major and minor versions and the associated community capabilities and features.

74.1.2 Benefits

- **Easy to manage**: Remove inefficient and time-consuming database administrative tasks without needing to provision infrastructure or maintain software. Amazon RDS offers the flexibility and scalability to quickly adjust to your database needs.
- **Choice of engines**: Deploy and scale the relational database engines of your choice in the cloud or on-premises.
- **High availability**: Achieve high availability with Amazon RDS Multi-AZ deployments.
- **Operational expertise**: Benefit from over a decade of proven operational expertise, security best practices, and innovation in databases born in the cloud.
- **Choice of commercial licensing options**: For commercial engines, Amazon RDS offers several licensing models – Bring Your Own License (BYOL), Bring Your Own Media (BYOM), or License Included (LI). With BYOL and BYOM, you can take advantage of your existing commercial software license to use with Amazon RDS. With License Included, the cost of the software license is included in the Amazon RDS instance pricing. You do not need to purchase a license separately.

74.2 Backup/Restore and Disaster Recovery

By default, Amazon RDS creates and saves automated backups of your DB instance securely in Amazon S3 for a user-specified retention period. Users can use the AWS RDS Console or AWS API tools to create automated or manual backups, snapshots, of their database or cluster instance.

74.3 Pricing Overview

Please see the AWS UK G Cloud 14 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/rds/

- **Service quotas:** https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/CHAP_Limits.html
- **Service FAQs:** https://aws.amazon.com/rds/faqs/

74.5 Technical Requirements

Please refer to https://docs.aws.amazon.com/rds/ 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.

75 Amazon Route 53

75.1 Service Overview

Amazon Route 53 provides highly available and scalable Domain Name System (DNS), domain name registration, and health-checking web services. 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 example.com into the numeric IP addresses, such as 192.0.2.1, that computers use to connect to each other. You can combine your DNS with health-checking services to route traffic to healthy endpoints or to independently monitor and/or alarm on endpoints. You can also purchase and manage domain names such as example.com and automatically configure DNS settings for your domains. 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.

75.1.1 Features

- **Route 53 Resolver**: Get recursive DNS for your Amazon VPCs in AWS Regions, VPCs in AWS Outposts racks, or any other on-premises networks. Create conditional forwarding rules and Route 53 endpoints to resolve custom names mastered in Amazon Route 53 private hosted zones or in your on-premises DNS servers.
- **Route 53 Resolver on Outposts endpoints**: Connect Route 53 Resolvers on Outpost racks with DNS servers in your on-premises data centers through Route 53 Resolver endpoints. This enables resolution of DNS queries between the Outposts racks and your other on-premises resources.
- **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.
- **Geoproximity routing**: Improves application responsiveness for your end users and helps apply data residency preferences by routing traffic to the geographically nearest resource.
- **Latency based routing**: Route end users to the AWS region that provides the lowest possible latency.
- **IP-based routing**: Fine-tune your DNS routing approach based on the Classless Inter-Domain Routing (CIDR) block that the query-originating IP address belongs to.
- **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.
- **Health Checks and Monitoring**: Amazon Route 53 can monitor the health and performance of your application as well as your web servers and other resources.
- **Domain Registration**: Amazon Route 53 offers domain name registration services, where you can search for and register available domain names or transfer in existing domain names to be managed by Route 53. View a full list of supported top-level domains (TLDs) and current pricing.
- **DNSSEC**: Enable DNSSEC signing for all existing and new public hosted zones, as well as DNSSEC validation for Amazon Route 53 Resolver.
- **CloudFront Zone Apex Support**: When using Amazon CloudFront to deliver your website content, visitors to your website can now access your site at the zone apex (or "root domain"). For example, your site can be accessed as example.com instead of www.example.com.
- **S3 Zone Apex Support**: Visitors to your website hosted on Amazon S3 can now access your site at the zone apex (or "root domain").
- **Amazon ELB Integration**: Amazon Route 53 is integrated with Elastic Load Balancing (ELB).
- **Management Console**: Amazon Route 53 works with the AWS Management Console. This web-based, point-and-click, graphical user interface lets you manage Amazon Route 53 without writing any code at all.
- **Weighted Round Robin**: Amazon Route 53 offers Weighted Round Robin (WRR) functionality.

75.1.2 Benefits

- **Route end users to your site reliably with globally-dispersed Domain Name System (DNS) servers and automatic scaling**: Amazon Route 53 ensures reliable and efficient routing of end users to your website by leveraging globally-dispersed Domain Name System (DNS) servers. With automatic scaling, the service dynamically adjusts to varying workloads, optimizing performance and maintaining a seamless user experience.

Use or disclosure of data contained on this sheet is subject to the restriction on the title page of this document.

- **Set up your DNS routing in minutes with domain name registration and straightforward visual traffic flow tools**: Amazon Route 53 streamlines the setup of DNS routing by providing quick and easy domain name registration, complemented by straightforward visual traffic flow tools. This enables users to configure their DNS settings within minutes, simplifying the process of managing and directing web traffic efficiently.
- **Customize your DNS routing policies to reduce latency, improve application availability, and maintain compliance**: Amazon Route 53 allows users to tailor DNS routing policies to specific needs, such as reducing latency, enhancing application availability, and ensuring compliance. This customization empowers users to optimize their DNS configurations for performance, resilience, and adherence to regulatory requirements.

75.2 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/.

75.3 Pricing Overview

Please see the AWS UK G Cloud 14 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/route53/
- **Service quotas:** https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/DNSLimitations.html
- **Service FAQs:** https://aws.amazon.com/route53/faqs/

75.5 Technical Requirements

Please refer to https://docs.aws.amazon.com/route53/ 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.

76 Amazon SageMaker

76.1 Service Overview

Amazon SageMaker is a fully managed service that brings together a broad set of tools to enable high-performance, low-cost machine learning (ML) for any use case. With SageMaker, you can build, train and deploy ML models at scale using tools like notebooks, debuggers, profilers, pipelines, MLOps, and more – all in one integrated development environment (IDE). SageMaker supports governance requirements with simplified access control and transparency over your ML projects. In addition, you can build your own FMs, large models that were trained on massive datasets, with purpose-built tools to fine-tune, experiment, retrain, and deploy FMs. SageMaker offers access to hundreds of pretrained models, including publicly available FMs, that you can deploy with just a few clicks.

76.1.1 Features

- **Store, share, and manage features for machine learning models**: Amazon SageMaker Feature Store is a fully managed, purpose-built repository to store, share, and manage features for machine learning (ML) models. Features are inputs to ML models used during training and inference.
- **Aggregate and prepare tabular and image data for machine learning in minutes**: Amazon SageMaker Data Wrangler reduces the time it takes to aggregate and prepare tabular and image data for ML from weeks to minutes. With SageMaker Data Wrangler, you can simplify the process of data preparation and feature engineering, and complete each step of the data preparation workflow (including data selection, cleansing, exploration, visualization, and processing at scale) from a single visual interface.
- **Access geospatial, purpose-built processing operations, pretrained models, and built-in visualization tools for geospatial ML**: Amazon SageMaker supports geospatial machine learning (ML) capabilities, allowing data scientists and ML engineers to build, train, and deploy ML models using geospatial data. Access geospatial data sources, purpose-built processing operations, pretrained ML models, and built-in visualization tools to run geospatial ML faster and at scale.
- **Explore and build machine learning models with fully managed notebooks**: Launch fully managed JupyterLab from Amazon SageMaker Studio in seconds. Use the integrated development environment (IDE) for notebooks, code, and data. You can use the quick start, collaborative notebooks in the IDE to access purpose-built machine learning (ML) tools in SageMaker and other AWS services for your complete ML development, from preparing data at petabyte scale using Spark on Amazon EMR to training and debugging models, deploying and monitoring models, and managing pipelines—all in one web-based visual interface.
- **Machine learning hub with foundation models, algorithms, and solutions deployable with just a few clicks**: Amazon SageMaker JumpStart is a machine learning (ML) hub that can help you accelerate your ML journey. With SageMaker JumpStart, you can evaluate, compare, and select FMs quickly based on pre-defined quality and responsibility metrics to perform tasks like article summarization and image generation. Pretrained models are fully customizable for your use case with your data, and you can easily deploy them into production with the user interface or SDK.
- **Learn and experiment with machine learning quickly with notebooks in your browser**: The Amazon SageMaker Studio Lab is based on the open-source and extensible JupyterLab IDE. Quickly create data analytics, scientific computing, and machine learning projects with notebooks in your browser.
- **Reduce the time and cost to train and tune models at scale without managing infrastructure**: Amazon SageMaker Model Training reduces the time and cost to train and tune machine learning (ML) models at scale without the need to manage infrastructure. You can take advantage of the highest-performing ML compute infrastructure currently available, and SageMaker can automatically scale infrastructure up or down, from one to thousands of GPUs.
- **Efficiently manage machine learning experiments**: With SageMaker Experiments, you can start tracking your experiments centrally from any environment or IDE using only a few lines of data scientist friendly Python code.
- **Reduce time to train foundation models by up to 40% with a purpose-built infrastructure for distributed training at scale**: SageMaker HyperPod is pre-configured with SageMaker's distributed training libraries that enable customers to automatically split training workloads across thousands of accelerators, so workloads

can be processed in parallel for improved model performance. SageMaker HyperPod also ensures customers can continue FM training uninterrupted by periodically saving checkpoints. When a hardware failure occurs during training, SageMaker HyperPod automatically detects the failure, repairs or replaces the faulty instance, and resumes the training from the last saved checkpoint, removing the need for customers to manually manage this process and helping them train for week or months in a distributed setting without disruption.

- **Deploy machine learning models to make predictions at the best price-performance for any use case**: Amazon SageMaker makes it easy to deploy ML models including foundation models (FMs) to make inference requests at the best price-performance for any use case. From low latency (a few milliseconds) and high throughput (millions of transactions per second) to long-running inference for use cases such as natural language processing and computer vision, you can use SageMaker for all your inference needs. SageMaker is a fully managed service and integrates with MLOps tools, so you can scale your model deployment, reduce inference cost, manage models more effectively in production, and reduce operational burden.
- **Create end-to-end machine learning workflows with purpose-built CI/CD tools**: With an intuitive UI and Python SDK you can manage repeatable end-to-end ML pipelines at scale. The native integration with multiple AWS services allows you to customize the ML lifecycle based on your MLOps requirements.
- **Easily operate machine learning models running on edge devices**: Amazon SageMaker Edge Compiler automatically optimizes ML models for deployment on a wide variety of edge devices. SageMaker Edge Compiler compiles your trained model into an executable format that applies performance optimizations that can make your model run up to 25x faster on the target hardware. Using SageMaker MLOps tools, you can easily train, test, troubleshoot, deploy, and govern ML models at scale to boost productivity of data scientists and ML engineers while maintaining model performance in production.
- **ML Governance with Amazon SageMaker**: Amazon SageMaker Model Cards makes it easier to capture, retrieve, and share essential model information, such as intended uses, risk ratings, and training details, from conception to deployment.
- **Deliver high-performance ML models at scale with tools to automate processes across the ML lifecycle**: Detect bias in data and machine learning models, and explain model predictions: With the rise of generative AI, data scientists and ML engineers can leverage publicly available foundation models (FMs) to accelerate speed-to-market. Amazon SageMaker Clarify supports FM evaluation to help you quickly evaluate, compare, and select the best FM for your use case based on a variety of criteria across different tasks within minutes. It allows you to adopt FMs faster and with confidence.
- **Build highly accurate machine learning models using a visual interface, no code required**: SageMaker Canvas provides access to ready-to-use models including foundation models from Amazon Bedrock or Amazon SageMaker JumpStart or you can build your own custom ML model. With SageMaker Canvas, you can easily access and import data from 50+ sources, prepare data using natural language and 300+ built-in transforms, build and train highly accurate models, generate predictions, and deploy models to production.
- **Perform ML development using a fully integrated development environment (IDE) for machine learning**: Amazon SageMaker Studio offers a wide choice of purpose-built tools to perform all machine learning (ML) development steps, from preparing data to building, training, deploying, and managing your ML models. You can quickly upload data and build models using your preferred IDE. Streamline ML

team collaboration, code efficiently using the AI-powered coding companion, tune and debug models, deploy and manage models in production, and automate workflows—all within a single, unified web-based interface.

- **Apply human feedback across the machine learning lifecycle to create high-quality models**: Amazon SageMaker Ground Truth offers the most comprehensive set of human-in-the-loop capabilities, allowing you to harness the power of human feedback across the ML lifecycle to improve the accuracy and relevancy of models. You can complete a variety of human-in-the-loop tasks with SageMaker Ground Truth, from data generation and annotation to model review, customization, and evaluation, either through a self-service or an AWS-managed offering.

76.1.2 Benefits

- **Choice of ML tools**: Enable more people to innovate with ML through a choice of tools—IDEs for data scientists and no-code interface for business analysts.
- **Fully managed, scalable infrastructure**: Build your own ML models, including FMs to power generative AI applications, with integrated purpose-built tools and high performance, cost effective infrastructure.
- **Repeatable and responsible ML workflows**: Automate and standardize MLOps practices and governance across your organization to support transparency and auditability.
- **Human-in-the-loop**: Harness the power of human feedback across the ML lifecycle to improve the accuracy and relevancy of FMs with human-in-the-loop capabilities.

76.2 Backup/Restore and Disaster Recovery

The service provides backup and recovery of cluster data. Users can use the backup script provided by the SageMaker HyperPod. The script provides an easy way to back up critical cluster data before patching or updating clusters. This is located at 1.architectures/5.sagemaker-hyperpod/patching-backup.sh in the Awsome Distributed Training GitHub repository at https://github.com/aws-samples/awsome-distributed-training.

76.3 Pricing Overview

Please see the AWS UK G Cloud 14 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/sagemaker/
- **Service quotas:** https://docs.aws.amazon.com/general/latest/gr/sagemaker.html
- **Service FAQs:** https://aws.amazon.com/sagemaker/faqs/

76.5 Technical Requirements

Please refer to https://docs.aws.amazon.com/sagemaker/ 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.

77 Amazon Security Lake

Amazon Security Lake automatically centralizes security data from AWS environments, SaaS providers, on premises, and cloud sources into a purpose-built data lake stored in your account. With Security Lake, you can get a more complete understanding of your security data across your entire organization. You can also improve the protection of your workloads, applications, and data. Security Lake has adopted the Open Cybersecurity Schema Framework (OCSF), an open standard. With OCSF support, the service normalizes and combines security data from AWS and a broad range of enterprise security data sources.

77.1.1 Features

- **Data aggregation in your account:** Security Lake creates a purpose-built security data lake in your account. Security Lake collects log and event data from cloud, on-premises, and custom data sources across accounts and Regions. The service stores the gathered logs in your Amazon Simple Storage Service (S3) buckets, so you retain complete control and ownership over your data.
- **Various supported AWS and third-party log and event sources:** Security Lake automatically collects logs and security findings for a wide range of AWS services.
- **Data normalization and support for OCSF:** Security Lake automatically normalizes AWS log and security findings to OCSF. You can add data from third-party security solutions, other cloud sources, and your custom data such as logs from internal applications or network infrastructure that you have converted into the OCSF format. With support for OCSF, Security Lake centralizes, transforms, and makes your security data available to your preferred analytics tools.
- **Multi-account and multi-Region support:** You can activate Security Lake across multiple Regions where the service is available and across multiple AWS accounts. You can aggregate security data across accounts on a per-Region basis or consolidate security data from multiple Regions into rollup Regions. Security Lake rollup Regions can help you comply with regional compliance requirements.
- **Security data lake access management:** Security Lake helps you streamline setting up access to your data lake for your security and analytics tools. For example, you might choose to only grant access to datasets from specified sources, such as CloudTrail. There are two modes of access available: data access, which issues a notification when new objects are written to the data lake, and query access, which allows tools to query the data stored in your security data lake.
- **Data lifecycle management and optimization:** Security Lake manages the lifecycle of your data with customizable retention settings and storage costs with automated storage tiering. Security Lake automatically partitions and converts incoming security data to a storage- and query-efficient Apache Parquet format. Security Lake supports Apache Iceberg tables in AWS Glue catalog to help you easily transition your analytics tools to run queries with increased performance.
- **Security observability with AWS AppFabric:** AWS AppFabric automatically normalizes SaaS application audit logs into the OCSF format and delivers normalized OCSF data to Security Lake. With the combination of Security Lake and AppFabric, you can easily aggregate, normalize, and visualize security data across key data sources. There are no fees associated with data normalization or data ingestion for the AppFabric integration with Security Lake. Standard AppFabric charges apply.

77.1.2 Benefits

- **Analyze multiple years of security data quickly:** Centralize petabytes of data from cloud, on-premises, and custom sources in your Amazon S3 buckets and use your preferred tools for security analytics.
- **Simplify your compliance monitoring and reporting:** Effortlessly centralize security data into one or more rollup Regions, making it easier to monitor and report on compliance.
- **Facilitate your security investigations with elevated visibility:** Give your security teams broader visibility to initiate thorough security investigations and rapid response to security incidents.
- **Unify security data management across hybrid environments:** Optimize data accessibility across your organization to facilitate a more comprehensive approach to security operations.

77.2 Backup/Restore and Disaster Recovery

This service allows for backup and recovery. Using the AWS Management Console, users can manage the lifecycle of data with customizable retention and replication settings. Data is stored in Amazon S3 and can be accessed if required. Users control this via a manual or scheduled API call. Users schedule and recover backups through a web interface.

77.3 Pricing Overview

Please see the AWS UK G Cloud 14 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/security-lake/
- **Service quotas:** https://docs.aws.amazon.com/general/latest/gr/securitylake.html
- **Service FAQs:** https://aws.amazon.com/security-lake/faqs/

77.5 Technical Requirements

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

- Amazon Security Lake User Guide: Provides detailed information and instructions for working with Amazon Security Lake.
- Amazon Security Lake API Reference: Describes the Amazon Security Lake API operations. Also provides example requests and responses for each operation.
- Amazon Security Lake in the AWS CLI Reference: Provides syntax and examples for the AWS CLI commands for use with Amazon Security Lake.

Please refer to and the following links for comprehensive technical documentation regarding this service.

78 Amazon Simple Email Service (SES)

78.1 Service Overview

Amazon Simple Email Service (Amazon SES)is a cloud-based email service provider that can integrate into any application for high volume email automation. Whether you use an

email software to send transactional emails, marketing emails, or newsletter emails, you pay only for what you use. Amazon SES is an email tool that also supports a variety of deployments including dedicated, shared, or owned IP addresses. Reports on sender statistics and email deliverability tools help businesses make every email count.

78.1.1 Features

- **Configure Sender 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).
- **Improve Inbox Deliverability using Virtual Deliverability Manager**: Deliverability, or ensuring emails reach inboxes instead of spam or junk, is a core element of a successful email strategy. Amazon Simple Email Service (SES) offers deliverability features through Virtual Deliverability Manager (VDM) that make it easier and more automated to improve inbox placement.
- **Analyze email deliverability**: SES's deliverability insights help email senders understand their deliverability performance. View at-a-glance reports on sending and delivery data in a single interface in the SES console, such as bounce rates, opens, and clicks. You can see breakdowns of these metrics per ISP, sender identity, and configuration set, as well as the delivery and engagement status of each email you send.
- **Review email deliverability recommendations**: SES can notify senders of email deliverability issues and provide actionable recommendations to help improve inbox placement, like reviewing email authentication configuration such as SPF, DKIM, and DMARC. Senders have ongoing visibility if a downgrade in email reputation or improvement opportunity is identified.
- **Automatically improve email deliverability**: Senders have the option to allow SES to automatically implement email deliverability recommendations, like optimizing email delivery patterns. Once SES detects an opportunity for improvement, the system will implement the change without the need for monitoring or manual adjustment.
- **Flexible Deployment Options through 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**: Customers can manage their own IP reputations by leasing dedicated IP addresses for use with their Amazon SES account. Customers can use the dedicated IP pools feature in SES console to create pools of those IP addresses, and either send all traffic from these dedicated IPs or use configuration sets to align specific use cases to specific IPs. Managed dedicated IP addresses are automatically set up and managed by SES on your behalf to help optimize your cost and usage. Standard dedicated IP addresses are those that you manually setup and manage.
- **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.

Use or disclosure of data contained on this sheet is subject to the restriction on the title page of this document.

- **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). 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 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.
- **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 email bounces 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.
- **Create alarms that notify**: Amazon SES automatically publishes the email bounce and email sender complaint metrics from this dashboard to Amazon CloudWatch. You can use CloudWatch to create alarms that notify you when your email bounce or complaint rates reach certain thresholds. With this information, you can take immediate action on issues that could impact your sender reputation.
- **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

- **Deliver marketing emails globally**: Deliver high-volume email campaigns with the service that sends hundreds of billions of emails per year.
- **Automate transactional messages**: Keep your customers up to date by sending automated emails, such as purchase or shipping notifications, order status updates, and policy change notices.Reach more customers’ inboxes as a trusted sender with email deliverability tools and IP management.
- **Send bulk email communications**: Deliver messages—including notifications and announcements—to large groups, and track results using configuration sets.Improve your bottom line with transparent pricing designed for bulk email.
- **Compliant**: Stay compliant from day one with GDPR and ISO certified options.

78.2 Backup/Restore and Disaster Recovery

This service allows for backup and recovery. The service can back up emails to Amazon S3, including SES logs such as delivery status and bounce reports (S3). Users control this via a manual or scheduled API call (Create Amazon SES receipt rule). Users schedule and recover backups through a web interface.

78.3 Pricing Overview

Please see the AWS UK G Cloud 14 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/
- **Service quotas:** https://docs.aws.amazon.com/ses/latest/dg/quotas.html
- **Service FAQs:** https://aws.amazon.com/ses/faqs/

78.5 Technical Requirements

- Developer Guide: Learn about Amazon SES and how to use its features through the redesigned new console experience.
- 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.

79 Amazon Simple Notification Service (SNS)

79.1 Service Overview

Amazon Simple Notification Service (Amazon SNS) is a web service that you can set up, operate, and use for sending notifications from the cloud. It provides developers with a highly scalable, flexible, and cost-effective capability to publish messages from an application and immediately deliver them to subscribers or other applications. It is designed to make web-scale computing easier for developers. Amazon SNS follows the "publish-subscribe" (pub-sub) messaging paradigm, with notifications being delivered to clients using a "push" mechanism that eliminates the need to periodically check or "poll" for new information and updates. With simple APIs requiring minimal up-front development effort, no maintenance or management overhead and pay-as-you-go pricing, Amazon SNS gives developers an easy mechanism to incorporate a powerful notification system with their applications.

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.
- **Maximum throughput**: Standard topics support a nearly unlimited number of messages per second.
- **Best-effort ordering**: Occasionally, messages might be delivered in an order different from which they were published.
- **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.

- **High throughput**: FIFO topics support up to 300 messages per second or 10 MB per second per FIFO topic (whichever comes first).
- **Strict ordering**: The order in which messages are published and delivered is strictly preserved such as first-in-first-out.
- **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. 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.
- **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 data protection**: Amazon SNS message data protection empowers topic owners to define data protection policies that can discover and protect sensitive data that is transmitted via their topics. This can help you to simplify your architecture by offloading data protection logic from your applications, while helping support your compliance objectives, for example, with regulations such as the General Data Privacy Regulation (GDPR). You can audit messages that are inbound to a topic to determine how much sensitive data they contain, prevent them from being delivered to downstream subscribers via blocking, or de-identify specific data in the payload via redaction or masking.
- **Message archiving, replay, and analytics**: For Standard topics, 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. For FIFO topics, Amazon SNS offers an in-place option to store and replay your messages, without the need to provision a separate archival resource. This improves the durability of your event-driven applications, and can help you recover from downstream failure scenarios.
- **SMS text messages**: Amazon SNS supports the ability to send SMS text messages at scale to 200+ countries, using a highly available and durable service, with redundancy across multiple SMS providers. With Amazon SNS, you are be able to control your originating identity by using a sender ID, long codes, or short codes. Moreover, you can use the Amazon SNS sandbox to validate your SMS workloads before moving them to production.
- **Mobile push notifications**: Amazon SNS mobile notifications make it simple and cost effective to fan out mobile push notifications to iOS, Android, Fire, Windows, and Baidu devices. Mobile notifications can be triggered from user-driven actions or business logic. Amazon SNS delivers mobile push notifications through Amazon Device Messaging (ADM), Apple Push Notification Service (APNs), Baidu Cloud Push (Baidu), Firebase Cloud Messaging (FCM), Microsoft Push Notification Service for Windows Phone (MPNS), and Windows Push Notification Services (WNS).
- **Email notifications**: Amazon SNS supports the delivery of notifications to email addresses subscribed to topics. This feature supports a variety of use cases. For example, you can use Amazon SNS to receive application alerts, as email notifications, to bring visibility into your DevOps workflows. Thus, you can be notified immediately when an event occurs, such as a specific change to your Amazon EC2 Auto Scaling group, or a new file uploaded to your Amazon S3 bucket, or a metric threshold breached in Amazon CloudWatch.

79.1.2 Benefits

- **Integrate your applications with FIFO messaging**: Deliver messages in a strictly ordered, first in, first out (FIFO) manner to maintain accuracy and consistency across independent applications.
- **Securely encrypt notification message delivery**: Encrypt messages with AWS Key Management Service (KMS), ensure traffic privacy with AWS PrivateLink, and control access with resource policies and tags.
- **Capture and fan out events from over 60 AWS services**: Fan out events across AWS categories, such as analytics, compute, containers, databases, IoT, machine learning (ML), security, and storage.

- **Send SMS texts to customers across over 240 countries**: Use worldwide SMS, with redundancy across providers. Set SMS origination identity with a sender ID, long code, short code, TFN, or 10DLC.
- **Save costs**: Simplify your application architecture and reduce costs with message filtering, batching, ordering, and deduplication.
- **Durability**: Increase message durability with archiving, replay, delivery retries, and dead-letter queues (DLQs).

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 14 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.
- Amazon SNS in the AWS CLI Reference: Describes the command line interface for Amazon SNS in detail. Provides basic syntax, options, and usage examples for each command.

80 Amazon Simple Queue Service (SQS)

80.1 Service Overview

80.1.1 Features

- **Standard queues**: Offer unlimited throughput and support a nearly unlimited number of transactions per second (TPS) per API action. A message is delivered at least once, but occasionally more than one copy of a message is delivered. Occasionally, messages might be delivered in an order different from which they were sent.
- **IFO queues**: Offer high throughput. By default, FIFO queues support up to 3,000 messages per second with batching or up to 300 messages per second (300 send, receive, or delete operations per second) without batching. A message is delivered once and remains available until a consumer processes and deletes it. Duplicates

aren't introduced into the queue. The order in which messages are sent and received is strictly preserved (i.e. First-In-First-Out).

- **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.
- **Messages**: 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.
- **Using Amazon SQS with other AWS infrastructure web services**: Amazon SQS message queuing can be used with other AWS services such as Amazon Redshift, Amazon DynamoDB, Amazon Relational Database Service (RDS), Amazon Elastic Compute Cloud (EC2), Amazon Elastic Container Service (ECS), AWS Lambda, and Amazon S3, to make distributed applications more scalable and reliable.

80.1.2 Benefits

- **Overhead made simple**: Eliminate overhead with no upfront costs and without needing to manage software or maintain infrastructure.
- **Reliability at scale**: Reliably deliver large volumes of data, at any level of throughput, without losing messages or needing other services to be available.
- **Security**: Securely send sensitive data between applications and centrally manage your keys using AWS Key Management.

Use or disclosure of data contained on this sheet is subject to the restriction on the title page of this document.

- **Cost-effective scalability**: Scale elastically and cost-effectively based on usage so you don't have to worry about capacity planning and pre-provisioning.

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 14 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/
- **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/ 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. Amazon S3 provides a simple web service interface that you can use to store and retrieve any amount of data, at any time, for use cases such as data lakes, cloud-native applications, and mobile applications.

Data is stored as objects within resources called "buckets", and a single object can be up to 5 terabytes in size. Objects can be accessed through S3 Access Points or directly through the bucket hostname. 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

Access Grants: Amazon S3 Access Grants map identities in directories such as Active Directory, or AWS Identity and Access Management (IAM) Principals, to datasets in S3. This

helps you manage data permissions at scale by automatically granting S3 access to end-users based on their corporate identity.

Access Points: You can create unique access control policies for each S3 access point to control access to shared datasets. For shared datasets including data lakes, media archives, and user-generated content you can easily scale access for hundreds of applications by creating individualized access points with names and permissions customized for each application. You can restrict any access point to a Virtual Private Cloud (VPC) to firewall S3 data access within your private networks. You can use AWS Service Control Policies to ensure all access points are VPC restricted.

Scale Objects with S3 Batch Operations: S3 Batch Operations is an Amazon S3 data management feature that lets you manage billions of objects at scale. You can make changes to object metadata and properties, or perform other storage management tasks, such as copying or replicating objects between buckets, replacing object tag sets, modifying access controls—instead of taking months to develop custom applications to perform these tasks.

Block Public Access: Store your data in Amazon S3 and secure it from unauthorized access with S3 Block Public Access. Apply S3 Block Public Access to every bucket in your account – both existing and any new buckets created in the future – and make sure that there is no public access to any object. S3 Block Public Access is enabled by default for all new buckets.

Optimize storage costs: Amazon S3 storage classes provide low cost storage based on different access patterns. These include S3 Standard for general-purpose storage of frequently accessed data; S3 Intelligent-Tiering to automatically optimize costs for data with unknown or changing access patterns; S3 Standard-Infrequent Access (S3 Standard-IA) and S3 One Zone-Infrequent Access (S3 One Zone-IA) for long-lived, but less frequently accessed data; S3 Glacier Instant Retrieval for archive data that needs immediate access, and S3 Glacier Flexible Retrieval (formerly S3 Glacier) for archive data that does not require immediate access but needs the flexibility to retrieve large sets of data at no cost, and Amazon S3 Glacier Deep Archive (S3 Glacier Deep Archive) for long-term archive and digital preservation at the lowest storage costs in the cloud.

Mountpoint: Mountpoint for Amazon S3 is an open source file client that you can use to mount an S3 bucket on your compute instance and access it as a local file system. It automatically translates local file system API calls to REST API calls on S3 objects. Mountpoint for Amazon S3 is optimized for high-throughput performance.

Multi-Region Access Points: S3 Multi-Region Access Points provide a global endpoint for routing Amazon S3 request traffic between AWS Regions. Each global endpoint routes Amazon S3 data request traffic from multiple sources, including traffic originating in Amazon Virtual Private Clouds (VPCs), from on-premises data centers over AWS PrivateLink, and from the public internet without building complex networking configurations with separate endpoints.

Object Lambda: Add your own code to S3 GET, HEAD, and LIST requests to modify and process data as it is returned to an application with S3 Object Lambda. In the AWS Management Console, you can configure a Lambda function and attach it to a S3 Object Lambda Access Point. From that point forward, S3 will automatically call your Lambda function to process any data retrieved through the S3 Object Lambda Access Point, returning a transformed result back to the application.

Object Lock: S3 Object Lock blocks permanent object deletion during a customer-defined retention period so that you can enforce retention policies as an added layer of data protection or for regulatory compliance. With S3 Object Lock, S3 Versioning is automatically enabled, and these features work together to prevent locked object versions from being

permanently deleted (accidental or intentional) or overwritten using a write-once-read-many (WORM) model.

Replication: Configure Amazon S3 to automatically replicate S3 objects across different AWS Regions by using S3 Cross-Region Replication (CRR) or between buckets in the same AWS Region by using S3 Same-Region Replication (SRR). S3 Replication offers the flexibility of replicating to multiple destination buckets in the same, or different AWS Regions.

Storage Lens: Amazon S3 is used to store large shared datasets across tens to hundreds of accounts and buckets, multiple Regions, and thousands of prefixes. S3 Storage Lens analyzes organization-wide metrics to deliver contextual recommendations that make it easy to take immediate action. S3 Storage Lens delivers more than 60 metrics (free metrics and advanced metrics) on S3 storage usage and activity to an interactive dashboard in the S3 Console.

S3 on Outposts: Amazon S3 on Outposts delivers object storage to your on-premises AWS Outposts environment to meet local data processing and data residency needs. Using the S3 APIs and features, S3 on Outposts makes it easy to store, secure, tag, retrieve, report on, and control access to the data on your Outpost. AWS Outposts is a fully managed service that extends AWS infrastructure, services, and tools to virtually any data center, co-location space, or on-premises facility for a truly consistent hybrid experience.

IAM Access Analyzer for S3 is a feature that helps you simplify permissions management as you set, verify, and refine policies for your S3 buckets and access points. Access Analyzer for S3 monitors your existing bucket access policies to verify that they provide only the required access to your S3 resources.

Data Exchange: AWS Data Exchange for Amazon S3 accelerates time to insight with direct access to data providers' Amazon S3 data. AWS Data Exchange for Amazon S3 helps you easily find, subscribe to, and use third-party data files for storage cost optimization, simplified data licensing management, and more.

Hybrid cloud storage: AWS Storage Gateway is a hybrid cloud storage service that lets you seamlessly connect and extend your on-premises applications to AWS Storage. Customers use Storage Gateway to seamlessly replace tape libraries with cloud storage, provide cloud storage-backed file shares, or create a low-latency cache to access data in AWS for on-premises applications.

Online data transfer: AWS DataSync makes it easy and efficient to transfer hundreds of terabytes and millions of files into Amazon S3, up to 10x faster than open-source tools. DataSync automatically handles or eliminates many manual tasks, including scripting copy jobs, scheduling and monitoring transfers, validating data, and optimizing network utilization.

81.1.2 Benefits

Performance: Amazon S3 provides industry leading performance for cloud object storage. 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.

Support for data lakes: Run big data analytics, artificial intelligence (AI), machine learning (ML), and high performance computing (HPC) applications to unlock data insights.

Robustness: Meet Recovery Time Objectives (RTO), Recovery Point Objectives (RPO), and compliance requirements with S3's robust replication features.

Operational effectiveness: Move data archives to the Amazon S3 Glacier storage classes to lower costs, eliminate operational complexities, and gain new insights.

Mobile and web-based cloud-native apps: Build fast, powerful mobile and web-based cloud apps that scale automatically in a highly available configuration.

Consistency: Amazon S3 delivers strong read-after-write consistency automatically for all applications, without changes to performance or availability, without sacrificing regional isolation for applications.

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. The service supports two-way replication and S3 batch replication. Users control this via a manual or scheduled API call. Users schedule and recover backups through a web interface.

With Amazon S3 Multi-Region Access Points failover controls, you can route all S3 data request traffic through a single global endpoint and directly control the shift of S3 data request traffic between AWS Regions at any time. During a planned or unplanned regional traffic disruption, failover controls let you control failover between buckets in different AWS Regions and accounts within minutes.

81.3 Pricing Overview

Please see the AWS UK G Cloud 14 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/

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 Simple Workflow Service (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. Amazon SWF acts as the coordination hub for all of the different components of your application:

- Maintaining application state
- Tracking workflow executions and logging their progress

- Holding and dispatching tasks
- Controlling which tasks each of your application hosts will be assigned to execute

82.1.1 Features

- **Names of workflows**: Use the AWS Management Console or the Amazon SWF APIs to specify the names of workflows.
- **Amazon SWF APIs for Workflow execution**: Use the Amazon SWF APIs to “start” a new workflow, which results in a particular sequence of workflow tasks, called a “workflow execution,” being kicked into action.
- **Amazon SWF APIs for task order**: Use the Amazon SWF APIs from your worker machines (a "worker" is a component of your application which handles specific tasks) to establish the task order, manage conditional flows, and execute loops for a workflow execution. Use the Amazon SWF APIs from your worker machines to request and execute workflow tasks in the cloud or on premises.
- **Status monitoring**: Monitor the status and progress of workflow executions and their associated tasks in the AWS Management Console.

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 process flow solutions and process automation software with a fully managed cloud process flow 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 process flow service is required as you add more cloud process flows to your application or increase the complexity of your process flows.
- **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 14 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 document analysis service that detects and extracts printed text, handwriting, structured data (such as fields of interest and their values) and tables from images and scans of documents. Amazon Textract's machine learning models have been trained on millions of documents so that virtually any document type you upload is automatically recognized and processed for text extraction. When information is extracted from documents, the service returns a confidence score for each element it identifies so that you can make informed decisions about how you want to use the results. For instance, if you are extracting information from tax documents you can set custom rules to flag any extracted information with a confidence score lower than 95%. Also, all extracted data are returned with bounding box coordinates, which is a rectangular frame that fully encompasses each piece of data identified, so that you can quickly identify where a word or number appears on a document. You can access these features with the Amazon Textract API, in the AWS Management Console, or using the AWS command-line interface (CLI).

83.1.1 Features

- Custom Queries: Amazon Textract provides you the ability to customize the pretrained Queries feature and improve extraction accuracy on your business specific document types while you maintain control and ownership of your data. Through the AWS Console you can upload as few as ten sample documents, annotate the data, and customize the pretrained Queries feature within a few hours.
- Layout: Amazon Textract provides you with the ability to extract layout elements such as paragraphs, titles, lists, headers, footers, and more from documents. Layout is a feature type in the Analyze Document API. Customers can use Layout as a stand-alone feature or in combination with other Analyze Document feature types.

- Optical character recognition: Amazon Textract OCR automatically detects printed and handwritten text from documents and images. Textract's ML powered OCR can recognize text in various fonts and styles, and it can also handle noisy or distorted text.
- 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.
- Signature Detection: Amazon Textract provides the ability to detect signatures on any document or image. This makes it easy to automatically detect signatures on documents such as checks, loan application forms, and claims forms. The location of the signatures and associated confidence scores are included in the API response.
- Query based extraction: Amazon Textract provides you with the flexibility to specify the data you need to extract from documents using queries. You can specify the information you need in the form of natural language questions (e.g., "What is the customer name") and receive the exact information (e.g., "John Doe") as part of the API response. You do not need to know the data structure in the document (table, form, implied field, nested data) or worry about variations across document versions and formats. Textract Queries are pre-trained on a large variety of documents including paystubs, bank statements, W-2s, loan application forms, mortgage notes, claims documents, and insurance cards.
- Analyze Lending: Analyze Lending API is a managed, preconfigured intelligent document processing API that fully automates the extraction of information from loan packages. Customers simply upload their mortgage loan documents to the Analyze Lending API and its prebuilt machine learning models will classify and split the document package by document type.
- Invoices and receipts: 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.
- Amazon Textract pricing: 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. 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, form data, queries or process invoices and identity documents. See the FAQ for additional details about pages and acceptable use of Textract.

83.1.2 Benefits

- Efficiency: Textract can extract the data in minutes instead of hours or days. Drive higher business efficiency and faster decision making while reducing costs.
- Extract key insights with high accuracy from virtually any document: Easily extract relevant data from government-related forms such as small business loans, federal tax forms, and business applications with a high degree of accuracy.
- Scalability: Scale up or scale down the document processing pipeline to quickly adapt to market demands.
- Automate data processing securely with data privacy, encryption, and compliance standards. Any content processed by Amazon Textract is encrypted and stored at rest in the AWS region where you are using Amazon Textract. Textract is compliant with PCI, ISO, and SOC.

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 14 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 analyze 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 adaptive query engine lets you access and analyze recent and historical data together, without having to specify its location.

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

- **Performance and scalability**: Amazon Timestream is a serverless analytics database for ingesting more than tens of gigabytes of timeseries data per minute and running SQL queries on terabytes of timeseries data in seconds. It is ideal for uses cases such as security analytics and monitoring the quality of streaming video.
- **Security**: All data in Amazon Timestream is automatically encrypted by default, so you don't need to manually encrypt data at rest or in transit. With native integrations for AWS IAM and AWS KMS services, you can securely manage access to your resources and data, including specifying an AWS KMS customer managed key (CMK) for encrypting data in the magnetic store.
- **Integrations with AWS services**: Amazon Timestream integrates with commonly used services for importing and exporting data, boosting your application with Machine Learning, or visualizing your data. You can send data to Amazon Timestream using AWS IoT Core, Amazon Kinesis, Amazon MSK, and open source Telegraf. You can use Amazon SageMaker with Amazon Timestream for machine learning. You can also visualize data using Amazon QuickSight, Grafana, and business intelligence tools through JDBC.
- **Cost effective**: Amazon Timestream supports high performance and cost-effective 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.
- **Developer productivity**: You can access Amazon Timestream using the AWS SDKs. Timestream supports two SDKs per language. Supported languages include: Java, Java v2, Go, Python, Node.js, .NET.

84.1.2 Benefits

- **Scale with ease**: Analyze time-series data using SQL, with built-in analytic functions for smoothing, approximation, and interpolation
- **Fully managed**: Serverless database processes millions of queries per day and automatically scales as needed
- **Trusted security**: Simplify data lifecycle management with storage tiers, including a memory store for recent data and a magnetic store for historical data.
- **Familiarity of open source**: Derive faster insights from your data and make business decisions at a fraction of the cost of existing time-series solutions

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 14 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 AWS Artificial Intelligence (AI) service that makes it easy for you to convert speech to text. Using Automatic Speech Recognition (ASR) technology, you can use Amazon Transcribe for a variety of business applications, including transcription of voice-based customer service calls, generation of subtitles on audio/video content, and conduct (text-based) content analysis on audio/video content.

85.1.1 Features

- **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.
- **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.

- **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.
- **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.
- **Toxic audio content detection**: Amazon Transcribe Toxicity Detection uses Machine Learning to keep audio conversations civil and constructive to encourage a safe and inclusive online environment. Toxic audio content is flagged into one of several categories for human moderators to easily identify and take appropriate action.
- **Improve productivity in contact centers with generative AI-powered call summarization**: Automatically generate call summaries to help agents focus on providing excellent customer experiences and increase productivity by reducing after-call manual summarization. This capability summarizes an interaction with a caller, capturing key components such as the reason for the call, steps taken to resolve issue, and next steps. Contact center supervisors can review call summaries to quickly understand the context of an interaction without reading the whole transcript while investigating caller issues.
- **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.

- **Contact center integration with Genesys Cloud CX**: Genesys Cloud CX is a cloud contact center solution that unifies customers and agent experiences across multiple channels such as phone, text and chat. You can stream your call audio to Amazon Transcribe from the Genesys Cloud environment to improve agent productivity and extract customer interaction insights.
- **Contact center integration with Amazon Chime SDK**: The Amazon Chime SDK is a set of real-time communications components that developers can use to quickly add audio calling, video calling, and screen sharing capabilities to their own web, mobile or telephony applications.
- **Contact center integration with Amazon Chime Voice Connector**: Amazon Chime Voice Connector lets you easily integrate with SIP-based contact centers to generate live, user attributed transcripts with Amazon Transcribe.

85.1.2 Benefits

- **High quality**: 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).
- **Easy to read transcripts**: Enhance accuracy with custom models that understand your domain specific vocabulary. 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. You can extract key business insights from customer calls, video files, clinical conversations, and more.
- **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.
- **Improve business outcomes**: With state of the art speech recognition models that are fully managed and continuously trained, Transcribe Call Analytics allows to extract actionable insights from customer conversations.

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 14 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. 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 Health Insurance Portability and Accountability Act of 1996 (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-gynecology, pediatrics, oncology, radiology and urology.

86.1.1 Features

- **Medical specialties**: 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. 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

- **Improve clinical documentation**: Medical doctors and practitioners can use Amazon Transcribe Medical to quickly and efficiently document clinical conversations into electronic health record (EHR) systems for analysis. The service helps to provide privacy and safety of clinical data by masking sensitive information.

- **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 14 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.

- Amazon Transcribe Medical Developer Guide: Provides a conceptual overview of Amazon Transcribe Medical, which has transcription features specifically designed for medical professionals.

87 Amazon Translate

87.1 Service Overview

Amazon Translate is a Neural Machine Translation (MT) service for translating text between supported languages. Powered by deep learning methods, the service provides high-quality, affordable, and customizable language translation, enabling developers to translate

company and user-authored content, or build applications requiring support across multiple languages. The service can be used via an API, enabling either real-time or batch translation of text from the source language to the target language.

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.
- **Real-Time Document Translation**: Amazon Translate supports real-time document translation of Text, HTML and Docx files. Customers can submit a document from the AWS Console, CLI, or SDK API and receive the translated document in real-time while maintaining the format of the original document.
- **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

- **High quality translations**: Deliver highly accurate and continually improving translations for a wide range of use cases.

- **Batch and real-time translations**: Build batch and real-time translation capabilities into your applications with a single API call.
- **Customization**: Customize your machine learning (ML)–translated output to define brand names, model names, and other unique terms.

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 14 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.
- API Reference: Describes the Amazon Translate API operations and parameters in detail.

88 Amazon Verified Permissions

88.1 Service Overview

Amazon Verified Permissions is a scalable, fine-grained permissions management and authorization service for custom applications built by you. Verified Permissions enables your developers to build secure applications faster by externalizing authorization and centralizing policy management and administration. Verified Permissions uses the Cedar policy language to define fine-grained permissions for application users.

88.1.1 Features

- **Schema**: Verified Permissions uses the schema to validate that a static policy or policy template is consistent with the application's authorization model. You can use JSON to define a schema in Verified Permissions. You can define action groups in your schema, which are policies that permit or forbid groups of actions.
- **Authorization requests**: Connect your application to the service through the API to authorize user access requests. For each authorization request, the service retrieves the relevant policies and evaluates those policies to determine whether a user is permitted to take an action on a resource given context inputs such as users, roles, group membership, and attributes.

- **Policy Store**: A policy store is a container of policies in Verified Permissions that is logically isolated from other containers. You can have a separate policy store for each tenant use of a Verified Permissions application; you can delete one tenant's policy store without affecting the resources, schemas, policies, and policy templates of any other policy store.
- **Test Bench Feature**: The test bench feature is a tool for testing and troubleshooting Verified Permissions policies by running a simulated authorization request against all the policies in your policy store. The test bench uses the parameters that you specify to determine whether the policies in your policy store would authorize the request.
- **Policy Templates**: You can use a policy template, which is a policy statement with placeholders in the scope that are to be filled in with specific values. A policy template can have placeholders for the principal, the resource, or both. Updates to the policy template are reflected across all principals and resources that use the template, also known as a template-linked policy.
- **Query policies**: Using Verified Permissions APIs, you can run specific queries against the policies stored in Verified Permissions. You can query your policies to determine which are applied to specific principals, specific resources, or both.
- **Auditing and logging**: You can configure and connect Verified Permissions to send your policy management and authorization logs to AWS CloudTrail.
- **Integration with Amazon Cognito**: You can pass your authentication token from Amazon Cognito into an authorization request running through Verified Permissions. This allows you to pass through identity provider attributes directly into a policy evaluation and thereby an authorization decision generated by Verified Permissions.
- **Integration with AWS CloudFormation**: Verified Permissions is integrated with CloudFormation, a service that helps you model and set up your AWS resources so that you can spend less time creating and managing your resources and infrastructure. You create a template that describes all the AWS resources that you want, and CloudFormation provisions and configures those resources for you.
- **Extensibility**: The Verified Permissions SDK is available using C++, Go, Java, JavaScript, Kotlin, .NET, Node.js, PHP, Python, Ruby, Rust, and Swift.

88.1.2 Benefits

- **Decouple authorization from application logic**: Accelerate application development by decoupling authorization from business logic.
- **Protect your application resources**: Protect application resources and manage user access to the principle of least privilege. You can define your policy model, create and store policies in a central location, and evaluate access requests in milliseconds. It also highlights permissions that are overly privileged and not valid.
- **Simplify application and resource access audits**: Verified Permissions supports governance and compliance. It provides auditing tools to configure, maintain, and analyze permissions across multiple diverse applications to help answer questions, such as who has access to what. You can simplify compliance audits at scale using automated analysis to confirm that permissions work as intended.
- **Continual, real-time authorization decisions**: Build applications aligned with Zero Trust principles of continual real-time authorization decisions. As a policy engine, Verified Permissions can help your application verify user action in real time, as required for Zero Trust.
- **Cost savings**: Verified Permissions helps reduce operational costs by providing you with a fine-grained authorization system to enforce access based on roles and attributes of your identities and resources.

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 14 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/verifiedpermissions/
- **Service quotas:** https://docs.aws.amazon.com/verifiedpermissions/latest/userguide/quotas.html
- **Service FAQs:** https://aws.amazon.com/verified-permissions/faqs/

88.5 Technical Requirements

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

- Amazon Verified Permissions User Guide: This guide provides a conceptual overview of Verified Permissions and explains how to use it to define policies as permissions within your application.
- API Reference Guide: This guide provides details about using the actions, data types, and API requests that can be made with Verified Permissions.
- Cedar policy**:** This documentation provides a detailed overview of the Cedar policy language for access control. This guide helps you understand how to define principals, actions, resource types, and conditions for your schema and policies.
- Designing an authorization model: This guide shares key considerations and criteria to consider when designing an authorization model for your application.

89 Amazon Virtual Private Cloud (VPC)

89.1 Service Overview

Amazon VPC lets you provision a logically isolated section of the Amazon Web Services (AWS) cloud where you can launch AWS resources in a virtual network that you define. You have complete control over your virtual networking environment, including selection of your own IP address ranges, creation of subnets, and configuration of route tables and network gateways. You can also create a hardware Virtual Private Network (VPN) connection between your corporate datacenter and your VPC and leverage the AWS cloud as an extension of your corporate datacenter.

89.1.1 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.
- **Ingress Routing**: 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.
- **Network Access Analyzer**: You can use Network Access Analyzer to understand network access to your resources, helping you identify improvements to your cloud security posture and 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.
- **Network Manager**: Network Manager provides tools and features to help you manage and monitor your network on AWS. Network Manager makes it easier to perform connectivity management, network monitoring and troubleshooting, IP management, and network security and governance.
- **Reachability Analyzer**: This static configuration analysis tool enables you to analyze 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**: Security groups can act as a firewall for associated Amazon EC2 instances, controlling inbound and outbound traffic at the instance level.
- **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.
- **Lattice**: This service helps you consistently connect, monitor, and secure communications between your applications. You can define policies for network traffic management, access, and monitoring to connect compute services in a simplified and consistent way across instances, containers, and serverless applications.

89.1.2 Benefits

- **Increase security**: Secure and monitor connections, screen traffic, and restrict instance access inside your virtual network.
- **Save time**: Spend less time setting up, managing, and validating your virtual network.

- **Manage and control your environment**: Customize your virtual network by choosing your own IP address range, creating subnets, and configuring route tables.

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 14 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/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/

89.5 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.

90 Amazon WorkMail (Lot 2)

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

- **Amazon WorkMail now offers no cost options to migrate**: Amazon has collaborated with email migration specialists audriga and Transend to make it easier

for you to move your email to Amazon WorkMail. You can use these solutions to migrate email, meetings, contacts, and notes from Microsoft Office 365, Microsoft Exchange Server, Google Apps for Work, IMAP and POP. See audriga and Transend websites for specific migration capabilities, and which usage fees and basic support costs are covered.

- **Microsoft Outlook Compatible**: Amazon WorkMail provides 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.
- **Enterprise-Grade Security**: Amazon WorkMail automatically encrypts all of your data at rest with encryption keys you control, using the AWS Key Management Service.
- **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.
- **Journaling**: Journaling allows you to record all email communication sent or received by your Amazon WorkMail organization. By enabling journaling, you can meet your data security and information protection policies using third party compliance solutions for email archiving and e-discovery. Because Amazon WorkMail uses industry standards for journaling, you can integrate with any third party solution that supports SMTP journaling.
- **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. You can also programmatically receive Exchange Web Services (EWS) push notifications. This means your existing IT service management tools, workflows, and third party applications can automate WorkMail migration and management, and respond to events in your mailboxes in real-time.
- **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. You can require device encryption, require that devices lock, specify requirements for device password strength, and wipe devices remotely using the AWS Management Console.
- **Email Flow Rules**: Amazon WorkMail allows you to use email flow rules to filter inbound email traffic for your Amazon WorkMail organizations, which helps you

reduce email from unwanted senders, route suspicious mail to junk folders, and make sure important messages are successfully delivered. Email flow rules can be applied based on specific email addresses, or entire email domains.

- **Spam and Virus Protection**: Amazon WorkMail scans all incoming and outgoing email for spam, malware, and viruses to help protect your users from malicious email.
- **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 14 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.
- API Reference: Describes the API operations for Amazon WorkMail.

91 Amazon WorkSpaces (Lot 2)

91.1 Service Overview

Amazon WorkSpaces is a fully managed desktop virtualization service for Windows, Linux, and Ubuntu that allows you to access resources from any supported device, including Amazon WorkSpaces Thin Client.

91.1.1 Features

- **Operating System support**: Choose your operating system (Windows, Amazon Linux, Ubuntu Linux) and select from a range of hardware configurations, software configurations, and AWS Regions.
- **Choose your protocol**: PCoIP or WorkSpaces Streaming Protocol (WSP).
- **Persistent desktop experience**: Connect to your WorkSpace and pick up from right where you left off. WorkSpaces provides a persistent desktop experience.
- **Billing flexibility**: WorkSpaces provides the flexibility of either monthly or hourly billing for WorkSpaces.
- **BYOL**: For Windows desktops, you can bring your own licenses and applications, or purchase them from the AWS Marketplace for Desktop Apps.
- **Managed directory**: Create a standalone managed directory for your users, or connect your WorkSpaces to your on-premises directory so that your users can use their existing credentials to obtain seamless access to corporate resources.
- **Security and encryption**: Use multi-factor authentication (MFA) for additional security and AWS Key Management Service (AWS KMS) to encrypt data at rest, disk I/O, and volume snapshots. Control the IP addresses from which users are allowed to access their WorkSpaces.

91.1.2 Benefits

- **Security**: Provide secure access to apps and data stored on AWS rather than on vulnerable endpoint devices.

- **Productivity**: Maximize productivity and business continuity with a financially backed uptime SLA.
- **Scalability**: Scale on demand with fixed-rate hourly billing, no overprovisioning, and no upfront costs.
- **Flexibility**: Use the same tools to manage WorkSpaces that you use to manage on-premises desktops. Bring a Microsoft 365 Apps for enterprise license to run on WorkSpaces or purchase Microsoft Office bundles.

91.2 Backup/Restore and Disaster Recovery

This service allows for backup and recovery. The service provides cross-Region redirection, a feature that works with your Domain Name System (DNS) failover routing policies to redirect your WorkSpaces users to alternative WorkSpaces in another AWS Region when their primary WorkSpaces aren't available. For more information, see Cross-Region redirection for Amazon WorkSpaces.

91.3 Pricing Overview

Please see the AWS UK G Cloud 14 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/workspaces/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 WorkSpaces.
- API Reference: Describes the API operations for Amazon WorkSpaces.
- WorkSpaces section of the AWS CLI Reference: Describes the AWS CLI commands that you can use to administer Amazon WorkSpaces. Provides syntax, options, and usage examples for each command.

92 AWS Amplify

92.1 Service Overview

AWS Amplify allows you to build full-stack web and mobile apps on AWS. AWS Amplify consists of a set of tools (open source framework, visual development environment, console) and services (web app and static website hosting) to accelerate the development of mobile and web applications on AWS. Build and host your frontend, add features like auth and storage, connect to real-time data sources, deploy, and scale to millions of users.

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/desktop/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 behavior 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 analyze 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. Powered by Amazon Lex.
- **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. Use 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. Powered by Amazon S3.
- **Amplify Hosting**: AWS Amplify offers a fully managed service for deploying and hosting full stack web applications, with built-in CI/CD workflows that accelerate your application release cycle. A full stack 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.
- **User management**: Use Amplify Studio to manage app users and content. Amplify Studio offers a way to manage authenticated users. Create/edit users and groups, update user properties, auto-verify signups, and more, without going through verification flows.
- **Content management**: Developers can give content editors and testers access to update the app data via Amplify Studio. Admins can save content as markdown to render rich text.
- **Override generated resources**: Modify fine-grained backend resource settings and override via CDK. Amplify does the heavy lifting for you. For example, use Amplify to add new Cognito resources with default settings to your backend. Override just the settings you want with amplify override auth.
- **Custom AWS resources**: The Amplify CLI provides escape hatches to add custom AWS resources with CDK or CloudFormation. Running the “amplify add custom” command in your Amplify project provides CDK or CloudFormation placeholders along with mechanisms to reference other Amplify-generated resources.
- **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 artifacts 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 artifacts (including CloudFormation templates, API resolver code, and client-side code generation).
- **Amplify Libraries**: AWS Amplify offers use case-centric open source libraries 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 modeling 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.
- **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.
- **Amplify Hosting**: Host your front-end web app, create/delete backend environments, setup CI/CD on the front end and backend, and access Amplify Studio to configure your backend and manage app content and users.

92.1.2 Benefits

- **Everything you need to build apps in the cloud**: Simply configure a backend and connect it to your frontend in just a few lines of code. No cloud expertise required.
- **Build cloud-connected UI**: Accelerate frontend UI development with dozens of pre-built, themeable, and responsive cloud-connected React components.
- **Connect your app to data**: Build secure, real-time GraphQL APIs backed by AWS databases quickly and easily.
- **Add Authentication and authorization**: Enable secure authentication flows with email/social sign-up and multi-factor auth. Control data and file access with fine-grained authorization rules.
- **Store and manage app content and data**: Store user-created content and application data in Amazon S3 buckets. Enable file uploading and display with pre-built components.
- **Deploy your frontend, fast**: Easily deploy and host fast, secure, and reliable websites and server-side rendered apps in a few clicks.
- **Full-stack CI/CD workflows**: Connect your git branches to automatically deploy frontend and backend changes in a single workflow.

- **Scale to zero or millions**: Serverless pay-as-you-go pricing means you only pay for what you use.
- **Wide reaching:** Access over 200 AWS services to support new use cases, DevOps practices, and user growth.

92.2 Backup/Restore and Disaster Recovery

The service provides backup of code and application data. 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 14 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 Runner

93.1 Service Overview

AWS App Runner is a fully managed application service that lets you build, deploy, and run web applications and API services without prior infrastructure or container experience. AWS App Runner builds and deploys web applications automatically, load balances traffic with encryption, scales to meet your traffic needs, and allows for the configuration of how services are accessed and communicate with other AWS applications in a private Amazon VPC.

93.1.1 Features

- **Accelerate application development and deployment**: AWS App Runner is a fully managed service that removes the operation burden involved in running your application, by automatically managing load balancing and providing a secure and scalable environment. This allows developers to focus on writing code and creating applications, rather than managing infrastructure, and drastically reduce the time-to-production.
- **Automatic scaling**: AWS App Runner automatically scales applications based on traffic needs, ensuring high availability and low latency. For businesses running microservices at scale, it allows for flexible scaling of each component of the application, ensuring optimal performance even under heavy load.

- **Built-in security and compliance**: AWS App Runner helps meet security and compliance requirements for your application, by providing built-in security features and adhering to AWS's compliance standards. This includes fully managed TLS encryption for secure connections, integration with AWS Secrets Manager and System Manager Parameter Store for secure handling of sensitive information, and private service endpoints that enables you to run internal applications securely within your Virtual Private Cloud (VPC).
- **Seamless integration with AWS services**: AWS App Runner provides seamless integration with other AWS services, such as databases, caches, and message queues. This facilitates building and running of complex applications that rely on these services, making it easier for you to leverage the full capabilities of AWS ecosystem.
- **Optimize costs with pay-as-you-go pricing**: With AWS App Runner, you only pay for the compute and memory resources used by your applications. This makes it a cost-effective solution for businesses of all sizes, allowing them to scale their applications based on demand without incurring unnecessary costs.
- **Maintain brand consistency with custom domains**: With AWS App Runner, you can associate any domain name you own to your App Runner service. AWS App Runner automatically provisions and automatically renews custom domain certificates for your application before their expiration date.
- **Gain insights with metrics and logs**: AWS 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.
- **Deployment flexibility**: AWS App Runner supports deployment from both source code and container image, allowing you to choose the deployment method that best suits your needs and workflows.

93.1.2 Benefits

- **Build frontend and backend applications**: Use App Runner to build and run API services, backend web services, websites, and more.
- **Run thousands of microservices simultaneously**: Flexibly scale each component of your application, decreasing risks to application resiliency and fostering increased agility and innovation.
- **Improve your time to market**: Use AWS best practices and technologies to deploy and run web applications at scale, drastically decreasing time to market.

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

Please see the AWS UK G Cloud 14 Pricing Document accompanying this service in the Digital Marketplace.

93.4 Service Constraints

Information on service features and constraints is available online here:

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

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

93.5 Technical Requirements

Please refer to https://docs.aws.amazon.com/apprunner/ 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.

94 AWS Application Discovery Service

94.1 Service Overview

AWS Application Discovery Service collects and presents data to enable enterprise customers to understand the configuration, usage, and behavior of servers in their IT environments. Server data is retained in the Application Discovery Service where it can be tagged and grouped into applications to help organize AWS migration planning. Collected data can be exported for analysis in Excel or other cloud migration analysis tools.

94.1.1 Features

- **Discover on-premises infrastructure**: AWS Application Discovery Service collects both server and database configuration information. Server information includes hostnames, IP addresses, MAC addresses, as well as the resource allocation and utilization details of key resources such as CPU, network, memory, and disk. Collected database information includes the database engine identity, version, and edition. Once collected, this information can be used to size AWS resources as part of migration planning.
- **Identify server dependencies**: 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**: 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 allows you to 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 analyze the time-series system performance for each server, the type of processes that are running on them, and the network dependencies between different servers.

94.1.2 Benefits

- **Discover on-premises server and database inventory**: Conduct inventory discovery to help accelerate migration by gathering server and database configuration information, operating system version history, and key resource allocation details. Gain a comprehensive snapshot of on-premises inventory to work with AWS experts and accelerate the migration journey.
- **Map network communication patterns**: Examine the connections between applications and servers to uncover unknown servers, better understand dependencies, and establish move groups. Protect collected data by encrypting it in transit to AWS and at rest within the service's data store.
- **Mobilize for migration**: Gather configuration, performance and Transport Control Protocol (TCP) connection data for both servers and databases to create a detailed migration plan. Integrate discovery data with other AWS services, such as AWS Migration Hub and AWS DMS Fleet Advisor, to support multiple migration tasks. Generate rightsized Amazon Elastic Compute Cloud (EC2) instances and plan for migration in Migration Hub or in AWS Database Migration Service.

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 14 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/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/

94.5 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.
- API Reference: Describes all the API operations for the AWS Application Discovery Service in detail. Also provides sample requests, responses, and errors for the supported web services protocols.

95 AWS Application Migration Service

95.1 Service Overview

AWS Application Migration Service (AWS MGN) is the recommended service for migrations to AWS. AWS Application Migration Service simplifies and expedites your migration to AWS by automatically converting your source servers from physical, virtual, or cloud infrastructure to run natively on AWS. It further simplifies your migration and reduces costs by allowing you

to use the same automated process for a wide range of applications, without changes to applications, their architecture, or the migrated servers.

95.1.1 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 APIs that are 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.
- Automate modernizations to your migrated applications.
- Plan and manage your application migrations.
- Plan and manage your migration waves.

95.1.2 Benefits

- **Migrate**: Migrate applications from any source infrastructure that runs supported operating systems.
- **Modernize**: Modernize your applications during migration with options such as disaster recovery and operating system or license conversion.
- **Maintain**: Maintain normal business operations throughout the application replication process.
- **Reduce costs**: Use one tool for a wide range of applications, with no need to invest in application-specific skills.

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

Please see the AWS UK G Cloud 14 Pricing Document accompanying this service in the Digital Marketplace.

95.4 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/

95.5 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.

96 AWS AppSync

96.1 Service Overview

AWS AppSync connects apps to data and events with secure, serverless, and performant GraphQL and Pub/Sub APIs. The GraphQL and Pub/Sub APIs help to simplify application development through a single endpoint to securely query, update, or publish data.

GraphQL APIs built with AWS AppSync give frontend developers the ability to query multiple databases, microservices, and APIs from a single GraphQL endpoint. Pub/Sub APIs built with AWS AppSync give frontend developers the ability to publish real-time data updates to subscribed API clients through serverless WebSocket connections.

96.1.1 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.
- **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.
- **Offline data synchronization**: Interact with and update your data, even when offline, with the Amplify DataStore. The Amplify DataStore is a queryable on-device DataStore for web, mobile and IoT developers that enables a local-first and familiar programming model to interact with data seamlessly whether you’re online or offline. Powered by AWS AppSync, the DataStore provides versioning, conflict detection and resolution in the cloud to automatically merge data from different clients as well as provide data consistency and integrity.
- **Generate a unified GraphQL API based on multiple independent source APIs from different teams**: A Merged API composes multiple GraphQL schemas, imports AppSync data sources and resolvers from multiple APIs, and consolidates all resources, combining all the APIs into a single Merged API endpoint that can be exposed to clients. This improves multi-service isolation and minimizes the need for service teams to interact with each other by automating the integration and composition of GraphQL schemas.
- **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.

- **Support for complex real-time use case**: With the subscriptions filtering capabilities in available with AWS AppSync, you can create a broad range of real-time experiences in your applications by leveraging filter operators, server-side filtering, and the ability to trigger subscription invalidations.
- **Get started without GraphQL**: If all you want is a an ephemeral API that receives messages published to a channel and delivers those messages in real-time to subscribed clients, you can set up your API without any GraphQL knowledge or experience using AWS AppSyncs PubSub API wizard.
- **Control your data**: 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 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.
- **Monitoring, analytics, logging and tracing**: With AWS AppSync you can easily configure AWS CloudWatch and AWS X-Ray to provide comprehensive logging and tracing for your GraphQL API.
- **Secure your APIs**: AWS AppSync offers support for private APIs that can be used to limit API exposure to only within a customer's VPC. Traffic to a private API uses connections that are designed to be secure and does not leave the Amazon network.

96.1.2 Benefits

- **Combine multiple data sources**: Access data from multiple sources with a single request. Instantly create APIs for your databases. Combine APIs into a single Merged API.
- **Create engaging real-time experiences**: Publish data from any event source to subscribed clients through serverless WebSockets. Subscriptions are supported with AWS AppSync against any of the data sources, so that when a mutation occurs, the results can be passed down to clients subscribing to the event stream immediately using over WebSockets.
- **Analyze multiple years of security data quickly**: Built-in security, monitoring, logging, and tracing. Optional caching for low latency. Centralize petabytes of data from cloud, on premises, and custom sources in your Amazon S3 buckets and use your preferred tools for security analytics.
- **Cost effective**: Pay only for requests to your API and any real-time messages delivered to connected clients.

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

96.5 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.

97 AWS Artifact

97.1 Service Overview

AWS Artifact, available in the console, is a self-service audit artifact retrieval portal that provides our customers with on-demand access to AWS' compliance documentation and AWS agreements.

97.1.1 Features

- **AWS Artifact Reports**: You can use AWS Artifact Reports to download AWS security and compliance documents, such as AWS ISO certifications, Payment Card Industry (PCI), and System and Organization Control (SOC) reports.
- **AWS Artifact Agreements**: You can use AWS Artifact Agreements to review, accept, and track the status of AWS agreements such as the Business Associate Addendum (BAA).
- **Central resource for compliance-related information**: AWS Artifact provides on-demand access to security and compliance reports from AWS and ISVs who sell their products on AWS Marketplace.
- **AWS Artifact notifications**: AWS Artifact notifications leverage AWS User Notifications service in order to configure notifications and to send emails. You can configure notifications using the AWS Artifact console. You can also view and configure notifications using the AWS User Notifications console.
- **AWS security and compliance posture**: Auditor-issued reports, certifications, accreditations, and other third-party attestations of AWS in a comprehensive resource help you to understand AWS's security and compliance posture.

97.1.2 Benefits

- **Save time**: Save time with on-demand access to AWS and Independent Software Vendor (ISV) compliance reports in a self-service portal. You can provide the AWS audit artifacts to your auditors or regulators as evidence of AWS security controls.
- **Manage at scale**: On-demand accept, terminate, and download compliance agreements with AWS.
- **Deploy with more confidence**: Improve your workload deploying confidence by understanding the compliance and security posture of AWS.
- **Manage select online agreements**: Review, accept, and manage your agreements with AWS and apply them to current and future accounts within your organization.
- **Assess third-party security and compliance**: Perform due-diligence of ISVs that sell products on AWS Marketplace, with on-demand access to their security and compliance reports.

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

This is a no cost, self-service portal for on-demand access to AWS' compliance reports.

97.4 Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/artifact/
- **Service FAQs:** https://aws.amazon.com/artifact/faq/

97.5 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.
- API Reference: Describes all the API operations for AWS Artifact.

98 AWS Audit Manager

98.1 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 make it easier 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.

98.1.1 Features

- **Prebuilt frameworks**: AWS Audit Manager offers prebuilt frameworks that cover a range of compliance standards. These frameworks help map your AWS resources to the requirements for industry standards and regulations. Examples of prebuilt

frameworks in AWS Audit Manager include AWS Control Tower, AWS License Manager, the General Data Protection Regulation (GDPR), the Service Organization Control 2 (SOC 2), and many more. AWS Audit Manager also provides a prebuilt framework to help gain visibility into how your generative AI implementation on Amazon Bedrock is working against AWS recommended best practices.

- **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. You can define custom controls to collect evidence from specific data sources to help show you are meeting internal audit and compliance 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 such as AWS CloudTrail, AWS Config, AWS Security Hub, and AWS License Manager.
- **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.
- **Search evidence**: Audit Manager more easily allows you to sift through thousands of pieces of collected evidence from multiple disparate sources, using search filters and groupings to identify trends and cross-reference issues. This will help you deep dive into issues identified via flagged compliance checks in the service, either in assessments – an automated data collection process against a specific set of controls – or on the Audit Manager dashboard.
- **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. The final assessment report contains a summary file on your assessment and provides links to an organized set of folders containing related evidence, which are named and organized as defined by the control set in each framework. The Audit Manager assessment report uses cryptographic verification to help you ensure the integrity of the assessment report.
- **Third-party risk assessment**: AWS Audit Manager provides features that help reduce the manual effort of third-party risk assessment. One example is the framework-sharing feature that allows you to share custom frameworks with your vendors in accordance with your organization's compliance requirements. Vendors can then gain access to these customized frameworks and use them to create assessments. In Audit Manager, an assessment is used to collect evidence for controls within the scope of your audit. Using the shared framework as a starting point, vendors can create an assessment that collects evidence for the controls in that framework.

98.1.2 Benefits

- **Map frameworks**: Map your AWS usage and controls with prebuilt and custom frameworks.

- **Automation**: Save time with automated evidence collection, and focus on confirming that your controls work properly.
- **Streamline operations**: Streamline collaboration across teams, and ensure the integrity of your audits with read-only permissions.

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

Please see the AWS UK G Cloud 14 Pricing Document accompanying this service in the Digital Marketplace.

98.4 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/

98.5 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.
- API Reference: Describes all the API operations for AWS Audit Manager in detail. Also provides sample requests, responses, and errors for the supported web service protocols.

99 AWS Backup

99.1 Service Overview

AWS Backup is a fully managed service that centralizes and automates data protection across AWS services like Amazon Simple Storage Service (S3), Amazon FSx, Amazon Elastic Compute Cloud (EC2), and Amazon Relational Database Service (RDS), and hybrid workloads like VMware on premises, VMware Cloud on AWS, and VMware Cloud on AWS Outposts. AWS Backup offers a cost-effective, fully managed, policy-based service that further simplifies data protection at scale. Using the AWS Backup Audit Manager, you can audit and report on the compliance of your data protection policies to help meet your business and regulatory needs. Together with AWS Organizations, use AWS Backup to centrally deploy data protection policies to configure, manage, and govern your backup activities across your AWS accounts and resources.

99.1.1 Features

- **Data protection of application resources on AWS and hybrid services overview**: AWS Backup helps protect application resources, including your AWS storage, database, and compute services as well as hybrid workloads like VMware. AWS Backup supports the following capabilities for all its supported services and third-party applications: automated backup scheduling and retention management, centralized data protection monitoring, AWS KMS-integrated backup encryption, cross-account management with AWS Organizations, data protection auditing and compliance reporting with AWS Backup Audit Manager, and write-once, read-many (WORM) with AWS Backup Vault Lock.
- **Centralized, policy-based data protection**: AWS Backup provides a backup console, public APIs, and a command line interface to centrally manage backups across the AWS storage, compute, database, and hybrid services your applications run on, including certain AWS services and hybrid applications like VMware workloads running on premises and in VMware CloudTM on AWS and AWS Outposts.

  The **AWS Backup vault** is a logical container that stores and manages your encrypted backups. When creating a backup vault, you must specify the AWS Key Management Service (AWS KMS) encryption key that encrypts the backups placed in this vault. All copied backups are encrypted with the key of the target vault. For more information about encryption, see the chart in Encryption for backups in AWS.
- **Encryption**: AWS Backup encrypts your backup data at rest and in transit, providing a comprehensive encryption solution that secures your backup data and helps meet compliance requirements. Your backup data is encrypted using encryption keys managed by the AWS Key Management Service (KMS), reducing the need to build and maintain a key management infrastructure.
- **Data protection policies**: You can create data protection policies and use AWS Organizations to enforce the protection policies throughout all the accounts in that organization. This provides multi-account backups and gives an additional layer of protection should the source account experience disruption from accidental or malicious deletion, disasters, or ransomware.
- **Backup role-based access control**: You can set resource-based access policies on backup vaults. With resource-based access policies, you can control access to backups in a backup vault across all users, rather than having to define permissions for each user. You can delegate backup policy management in AWS Organizations and cross account monitoring in AWS Backup.
- **Backup activity monitoring**: AWS Backup console includes an Amazon CloudWatch dashboard to see metrics on completed or failed backup, copy, and restore jobs. Within this dashboard, you can view job status by time period, customized to the schedule you desire. AWS Backup integrates with AWS CloudTrail, which provides a consolidated view of backup activity logs and simplifies the audit process for protected resources. AWS Backup integrates with Amazon Simple Notification Service (Amazon SNS), which can automatically alert you on backup activity such as when a backup succeeds or a restore is initiated. For a fully managed experience, you can use AWS Backup Audit Manager to monitor your backup activity across your accounts and Regions.
- **Application stack protection and recovery**: AWS Backup provides capabilities that help protect and recover critical data from ransomware events and account compromise. You can back up your AWS CloudFormation stack along with its resources like AWS IAM roles and Amazon VPC security groups. This means you can more easily recover your entire application stack, and manage compliance of your data protection policies across the entire application stack. You can import

application definitions and create application-wide protection plans managed on a recurring schedule and cross-account or cross-Region copy for additional protection from ransomware events. You can perform automated and periodic evaluation of restore viability as well as monitor restore job duration times with the restore testing feature. You can conduct recovery readiness test drills to prepare for possible data downtime or data loss events, satisfying compliance or regulatory requirements.

- **Data immutability (SEC-17a, WORM) and vaults**: AWS Backup Vault Lock allows you to protect your backups from deletion or changes to their lifecycle (making data immutable) 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 with backup policies such as retention periods, cold storage transitioning, cross-account, and cross-Region copy. This provides an additional layer of protection and helps meet your compliance requirements.
- **Data protection compliance with real-time analytics and insights**: AWS Backup Audit Manager is a capability that monitors and generates audit reports of your data protection activity, such as backup frequency or backup retention period. AWS Backup Audit Manager is a fully managed experience that can generate daily reports with insights on the compliance status of your data protection frameworks.
- **Backup audits for compliance posture**: You can audit and report on the compliance of your data protection policies to help you meet your business and regulatory needs with AWS Backup Audit Manager. You can continuously evaluate backup activity and generate audit reports that can help you demonstrate compliance with regulatory requirements.
- **Compliant retention for legal hold and e-discovery**: AWS Backup supports legal hold, which is used when an organization must retain certain data either for preservation, auditing, or as evidence in legal proceedings and e-Discovery. You can use legal holds to prevent backups from being deleted even if their retention period is over, and remain in place until they are explicitly released.
- **Reporting for compliance and governance**: You can use compliance report templates to generate you daily reports on the compliance of your backup activity and resources against the controls you defined in one or more frameworks.

99.1.2 Benefits

- **Simplify data protection of application resources on AWS and hybrid services**: Centralize backup and recovery with a fully managed, policy-based service.
- **Build a foundation for disaster recovery and business continuity**: Back up across different AWS Regions and accounts. Perform automatic restore testing and validation of backups.
- **Protect and recover from ransomware and account compromise**: Create immutable backups and use logically air-gapped vaults as part of a defense in depth strategy.
- **Govern data protection compliance**: Monitor, audit, and report on data protection policies and operations using real-time analytics and insights.

99.2 Backup/Restore and Disaster Recovery

AWS Backup is the backup service storing data resiliently on Amazon 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.

99.3 Pricing Overview

Please see the AWS UK G Cloud 14 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/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/

99.5 Technical Requirements

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/

100 AWS Batch

100.1 Service Overview

AWS Batch is a set of batch management capabilities that dynamically provisions the optimal quantity and type of compute resources such as CPU or memory optimized compute resources, based on the volume and specific resource requirements of the batch jobs. With AWS Batch, there is no need to install and manage batch computing software or server clusters, allowing you to instead focus on analyzing results and solving problems. AWS Batch plans, schedules, and executes your batch computing workloads using Amazon ECS, Amazon EKS, and AWS Fargate with an option to utilize spot instances.

100.1.1 Features

- **Multi-container jobs**: Multi-container jobs is a feature that makes it easier and faster to run simulations at scale when testing complex systems like those used in autonomous vehicles and robotics. With the ability to run multiple containers in a job, you can benefit from the advanced scaling, scheduling, and cost optimization capabilities of AWS Batch without rebuilding your system into a complex monolithic container. Instead, you can maintain the use of multiple smaller, modular containers that each represent different system components like a 3D virtual environment, robot perception sensors, or a data logging sidecar.
- **Support for tightly-coupled HPC workloads**: AWS Batch supports multi-node parallel jobs, so you can run single jobs that span multiple EC2 instances. With this feature, you can use AWS Batch to efficiently run workloads such as large-scale, tightly-coupled, high performance computing (HPC) applications or distributed GPU model training. AWS Batch also supports Elastic Fabric Adapter, a network interface where you can run applications that require high levels of internode communication at scale on AWS.
- **Granular job definitions and straightforward job dependency modeling**: With AWS Batch, you can 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 run. AWS Batch runs your jobs as containerized applications on Amazon ECS. You can also define dependencies between different jobs. For example, your batch job can be composed of three stages of processing with different resource needs. With dependencies, you can create three jobs with different res

- **Support for popular workflow engines**: AWS Batch can be integrated with commercial and open-source workflow engines such as Pegasus WMS, Luigi, Nextflow, Metaflow, Apache Airflow, and AWS Step Functions, so you can use familiar workflow languages to model your batch computing pipelines.
- **Integration with EC2 launch templates**: AWS Batch supports EC2 launch templates, so you can build customized templates for your compute resources while Batch scales instances to meet your requirements. You can specify your EC2 launch template to add storage volumes, select network interfaces, or configure permissions, among other capabilities.
- **Integrated monitoring and logging**: AWS Batch displays key operational metrics for your batch jobs in the AWS Management Console. You can view metrics related to compute capacity, as well as metrics for running, pending, and completed jobs. Logs for your jobs (for example, STDERR and STDOUT) are available in the console and are also written to Amazon CloudWatch Logs.
- **Fine-grained access control**: AWS Batch uses IAM to control and monitor the AWS resources that your jobs can access, such as Amazon DynamoDB tables. Through IAM, you can also define policies for different users in your organization.
- **AWS Batch on Amazon EKS**: AWS Batch can run your batch jobs on your existing Amazon EKS clusters. AWS Batch manages both the scaling of Kubernetes nodes and the placement of pods within your nodes. Additionally, AWS Batch manages queueing, dependency tracking, job retries, prioritization, and pod submission, along with providing support for Amazon Elastic Compute Cloud (EC2) On-Demand and Spot Instances.
- **AWS Batch with Fargate**: AWS Batch with Fargate resources gives you a completely serverless architecture for your batch jobs. If you're a current AWS Batch user, Fargate allows an additional layer of separation from EC2 instances. There's no need to manage or patch Amazon Machine Images (AMIs). When submitting your Fargate compatible jobs to Batch, you don't have to worry about maintaining two different services if you have some workloads that run on Amazon EC2 and others that run on Fargate. AWS provides a cloud-based scheduler complete with a managed queue and the ability to specify priorities, job retries, dependencies, timeouts, and more. AWS Batch manages submission to Fargate and the lifecycle of your jobs.
- **Priority-based job scheduling**: With AWS Batch, you can set up multiple queues with different priority levels. Batch jobs are stored in the queues until compute resources are available to run 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 (for example, memory-optimized compared to 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 then 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.

- **Dynamic compute resource provisioning and scaling**: When using Fargate or AWS Fargate Spot with AWS Batch, you only need to set up a few concepts in Batch (a compute environment, a job queue, and a job definition), and you have a complete queue, scheduler, and compute architecture, without managing a single piece of compute infrastructure.

  For those wanting EC2 instances, AWS Batch provides managed compute environments that dynamically provision and scale compute resources based on the volume and resource requirements of your submitted jobs. You can configure your managed compute environments with requirements such as the type of EC2 instances, VPC subnet configurations, the min/max/desired vCPUs across all instances, and the amount you are willing to pay for Spot Instances as a percent of the On-Demand Instance price.
- **Flexible allocation strategies**: With AWS Batch, you can choose three methods to allocate compute resources. These strategies allow you to factor in throughput and price when deciding how AWS Batch should scale instances for you.
- **Best fit**: AWS Batch selects an instance type that best fits the needs of the jobs with a preference for the lowest-cost instance type. If additional instances of the selected instance type are not available, AWS Batch will wait for the additional instances to become available. If there are not enough instances available, or if you are hitting Amazon EC2 service limits, then additional jobs will not run until currently running jobs have completed. This allocation strategy keeps costs lower but can limit scaling.
- **Best fit progressive**: AWS Batch will select additional instance types that are large enough to meet the requirements of the jobs in the queue, with a preference for instance types with a lower-cost-per-unit vCPU. If additional instances of the previously selected instance types are not available, AWS Batch will select new instance types.
- **Spot capacity–optimized**: AWS Batch will select one or more instance types that are large enough to meet the requirements of the jobs in the queue, with a preference for instance types that are less likely to be interrupted. This allocation strategy is only available for Spot Instance compute resources.

100.1.2 Benefits

- **Focus on analyzing results, not managing infrastructure**: Run hundreds of thousands of batch machine learning (ML), simulation, and analytics computing jobs without installing software or servers.
- **Native integration with AWS**: Natively integrate with AWS to implement scaling, networking, and management capabilities.
- **Optimize compute costs**: Reduce costs by optimizing computing job distribution based on volume and resource requirements.
- **Scale compute resources automatically**: Scale your compute resources automatically with fully managed infrastructure that supports large-scale processing and simulations.

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 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.

100.4 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/

100.5 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).

101 AWS Budgets

101.1 Service Overview

Using AWS Budgets, you can set a budget that alerts you when you exceed (or are forecasted to exceed) your budgeted cost or usage amount. You can also set alerts based on your RI or Savings Plans Utilization and Coverage using AWS Budgets.

101.1.1 Features

- **Monitor**: With custom budget threshold and variance tracking, you can monitor how your cost, usage, or your Reserved Instances and Savings Plans' utilization and coverage progresses towards your target.
- **Maintain**: With Budget Reports, you can keep your stakeholders up to date with your budget status.
- **Stay informed**: With Budget Alerts, you can stay informed when your budget limits are (or are about to be) exceeded.
- **Configure**: With Budget Actions, you can authorize AWS to take actions to stop the activities that caused your cost overage. You can use AWS Budgets to run an action on your behalf when a budget exceeds a certain cost or usage threshold. To do this, after you set a threshold, configure a budget action to run either automatically or after your manual approval.

101.1.2 Benefits

- Track your costs, usage, and coverage with custom budgets.
- Stay informed on forecasted spend and resource use.
- Create custom actions to prevent overages, inefficient resource use, or lack of coverage.

101.2 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/.

101.3 Pricing Overview

Please see the AWS UK G Cloud 14 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/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/

101.5 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.

102 AWS Certificate Manager

102.1 Service Overview

AWS Certificate Manager (ACM) is a service that lets you 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. ACM removes the time-consuming manual process of purchasing, uploading, and renewing SSL/TLS certificates.

With AWS Certificate Manager, you can request a certificate, deploy it on AWS resources such as Elastic Load Balancers, Amazon CloudFront distributions, and APIs on API Gateway, and let ACM handle certificate renewals. You pay for the AWS resources you create to run your application.

102.1.1 Features

- **Centrally manage certificates**: AWS Certificate Manager (ACM) makes it easy for you to centrally manage your SSL/TLS certificates from the AWS Management Console, AWS CLI, or ACM APIs. You can also audit the use of each certificate by reviewing your AWS CloudTrail logs.

- **Secure key management**: ACM 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.
- **AWS service integration**: ACM 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 and Amazon API Gateway. ACM 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.
- **Import third-party certificates**: ACM 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.

102.1.2 Benefits

- **Protect and secure your website**: Provision and manage certificates so you can securely terminate traffic to your website or application. Use key management for your certificates.
- **Protect your internal resources**: Secure communication between connected resources on private networks, such as servers, mobile and IoT devices, and applications. Use no-cost certificates with ACM integrated services.
- **Improve uptime**: Maintain SSL/TLS certificates, including certificate renewals, with automated certificate management. Get managed certificate renewal.

102.2 Backup/Restore and Disaster Recovery

The service provides backup and recovery of certificates.

102.3 Pricing Overview

Public SSL/TLS certificates provisioned through AWS Certificate Manager are free. You pay only for the AWS resources you create to run your application. If you manage AWS Private Certificate Authority (CA) through ACM, refer to the AWS Private CA Pricing page for more details and examples.

Please see the AWS UK G Cloud 14 Pricing Document accompanying this service in the Digital Marketplace.

102.4 Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/acm/
- **Service quotas:** https://docs.aws.amazon.com/acm/latest/userguide/acm-limits.html
- **Service FAQs:** https://aws.amazon.com/certificate-manager/faqs/

102.5 Technical Requirements

Please refer to https://docs.aws.amazon.com/acm/ 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.

103 AWS Chatbot

103.1 Service Overview

AWS Chatbot is an interactive agent that makes it easy to set up ChatOps for AWS in your Amazon Chime, Microsoft Teams, or Slack channels—and to securely interact with multiple AWS services. You can receive notifications about operational events, security findings, or budget alerts right in your chatroom, where your entire team can see and discuss them.

You can issue AWS Command Line Interface (CLI) commands from Microsoft Teams and Slack channels to retrieve diagnostic information, invoke AWS Lambda functions, configure Amazon Simple Storage Service (S3) buckets, change Kinesis shards, restart Amazon Elastic Compute Cloud (EC2) instances, and resolve AWS System Manager incidents.

103.1.1 Features

- **Supports Microsoft Teams, Slack and Amazon Chime**: You can add AWS Chatbot to your Amazon Chime, Microsoft Teams, or Slack channel with just a few clicks**.**
- **Receive notifications**: AWS Chatbot can set up ChatOps for AWS in your Amazon Chime, Microsoft Teams, or Slack channels. Choose these channels to receive notifications and Amazon Simple Notification Service (SNS) topics to trigger notifications.
- **Retrieve diagnostic information and configure AWS resources**: AWS Chatbot supports AWS Command Line Interface (CLI) commands for most AWS services from Microsoft Teams and 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.
- **Operationalize DevOps workflows and streamline communication in chat channels**: Operationalize frequently used DevOps runbook processes and incident response tasks in chat channels with custom notifications, customizable actions, and command aliases.
- **Receive expert guidance from Amazon Q - generative AI assistant**: With Amazon Q in AWS Chatbot, get expert guidance on AWS-related questions in a conversation experience to accelerate understanding of the AWS services, troubleshoot issues, and identify remediation paths.
- **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.

103.1.2 Benefits

- **Receive notifications in real time**: Establish integrations and security permissions between AWS resources and chat channels to receive preselected or event-driven notifications in real time.
- **Respond quickly to incidents**: Collaborate, retrieve observability telemetry, and respond quickly to incidents, security findings, and other alerts for applications in your AWS environment.
- **Prevent context switching**: Safely configure AWS resources, resolve incidents, and run tasks from Microsoft Teams and Slack without context switching to other AWS management tools.
- **Generative AI**: Get assistance from Amazon Q, the generative AI assistant to diagnose issues and identify next steps.

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

103.5 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.

104 AWS Cloud9

104.1 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. The service is pre-packaged with essential tools for popular programming languages, including JavaScript, Python, and PHP. You can quickly share your development environment with your team, pair program and track inputs in real time.

104.1.1 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 allows you to customize your view. You can adjust your panels in any direction with a drag-and-drop action.

- **Broad selection of run configurations**: AWS Cloud9 supports more than 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, and more.
- **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.
- **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.
- **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 color schemes that control syntax highlighting and the UI. You can 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.

104.1.2 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**: 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 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 has 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 run commands and directly access AWS services.
- **Start new projects quickly**: AWS Cloud9 is prepackaged with tooling for more than 40 programming languages, including Node.js, JavaScript, Python, PHP, Ruby, Go, and C++. This allows 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.

104.2 Backup/Restore and Disaster Recovery

The service provides backup and recovery of machines images and volumes. Users schedule backups through manual or schedules API calls. Users recover backups through a web interface.

104.3 Pricing Overview

Please see the AWS UK G Cloud 14 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/cloud9/
- **Service quotas:** https://docs.aws.amazon.com/cloud9/latest/user-guide/limits.html
- **Service FAQs:** https://aws.amazon.com/cloud9/faqs/

104.5 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.

- API Reference: Describes all the API operations for AWS Cloud9. Also provides sample requests, responses, and errors for the supported web services protocols.
- AWS CLI Reference for AWS Cloud9: Describes the AWS CLI commands that you can use to create and manage AWS Cloud9 development environments.

105 AWS Cloud Map

105.1 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.

105.1.1 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.

105.1.2 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.

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

Please see the AWS UK G Cloud 14 Pricing Document accompanying this service in the Digital Marketplace.

105.4 Service Constraints

Information on service features and constraints is available online here:

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

105.5 Technical Requirements

Please refer to https://docs.aws.amazon.com/cloud-map/ 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.

106 AWS CloudFormation

106.1 Service Overview

AWS CloudFormation is a service that helps you model and set up your AWS resources so that you can spend less time managing those resources and more time focusing on your applications that run in AWS.

You create a template that describes all the AWS resources that you want (like Amazon EC2 instances or Amazon RDS DB instances), and CloudFormation takes care of provisioning and configuring those resources for you. You don't need to individually create and configure AWS resources and figure out what's dependent on what; CloudFormation handles that. The following scenarios demonstrate how CloudFormation can help.

106.1.1 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 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.
- **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.
- **Git management of stacks**: AWS CloudFormation supports updating a stack from a CloudFormation template stored in a remote Git repository. With this feature, you will know if a stack deployment succeeded or failed, without having to navigate away from your remote repository, saving time by avoiding context switching.

106.1.2 Benefits

- **Scalability**: Scale your infrastructure worldwide and manage resources across all AWS accounts and regions through a single operation.
- **Extensibility**: Extend and manage your infrastructure to include cloud resources published in the CloudFormation Registry, the developer community, and your library.
- **Automation**: Automate resource management across your organization with AWS service integrations offering turnkey application distribution and governance controls.

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 14 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/cloudformation/
- **Service quotas:** https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/cloudformation-limits.html
- **Service FAQs:** https://aws.amazon.com/cloudformation/faqs/

106.5 Technical Requirements

Please refer to https://docs.aws.amazon.com/cloudformation/ 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.

107 AWS CloudHSM

107.1 Service Overview

AWS CloudHSM combines the benefits of the AWS cloud with the security of hardware security modules (HSMs). A hardware security module (HSM) is a computing device that processes cryptographic operations and provides secure storage for cryptographic keys. With AWS CloudHSM, you have complete control over high availability HSMs that are in the AWS Cloud, have low-latency access, and a secure root of trust that automates HSM management (including backups, provisioning, configuration, and maintenance).

107.1.1 Features

- **Security**: Manage your keys on FIPS validated hardware, protected with customer-owned, single-tenant HSM instances running in your own Virtual Private Cloud

(VPC). Separation of duties and role-based access control is inherent in the design of the AWS CloudHSM. AWS monitors the health and network availability of your HSMs; you control the HSMs and the generation and use of your encryption keys.

- **Performance**: Meet performance requirements of your applications through elasticity, adding or removing HSM instances while achieving latency and reliability goals. For more details, refer to the performance page in the AWS CloudHSM user guide.
- **Compliance**: Demonstrate compliance with security and privacy regulations, such as PCI, GDPR, HIPAA, and FedRAMP.
- **Availability**: Automatically load balance requests and securely replicate keys across HSM instances in your cluster, improving key durability and increasing capacity.
- **Flexibility**: Deploy an open solution that supports a wide range of cryptographic algorithms through industry-standards such as PKCS#11, JCE, OpenSSL, and CNG/KSP.

107.1.2 Benefits

- **HSMs are FIPS 140-2 level-3 validated**: AWS CloudHSM uses general purpose HSMs that are standards-compliant, single-tenant, and FIPS 140-2 level-3 validated. They provide more flexibility when compared to the fully-managed AWS services that have predetermined algorithms and key lengths for your application.
- **E2E encryption is not visible to AWS**: Because your data plane is end-to-end (E2E) encrypted and not visible to AWS, you control your own user management (outside of IAM roles).
- **Full control of your keys, algorithms, and application development**: AWS CloudHSM gives you full control of the algorithms and keys you use. You can generate, store, import, export, manage, and use cryptographic keys (including, session keys, token keys, symmetric keys and asymmetric key pairs). Additionally, AWS CloudHSM SDKs give you full control over application development, application language, threading, and where your applications physically exist.
- **Migrate your cryptographic workloads to the cloud**: Customers migrating public key infrastructure that use Public Key Cryptography Standards #11 (PKCS #11), Java Cryptographic Extension (JCE), Cryptography API: Next Generation (CNG), or key storage provider (KSP) can migrate to AWS CloudHSM with fewer changes to their application.

107.2 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 can invoke defined actions with periodic (24 hour) backup coverage. Users recover backups through a web interface.

107.3 Pricing Overview

Please see the AWS UK G Cloud 14 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/cloudhsm/index.html
- **Service quotas:** https://docs.aws.amazon.com/cloudhsm/latest/userguide/limits.html

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

107.5 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.

108 AWS CloudTrail

108.1 Service Overview

CloudTrail enables auditing, security monitoring, and operational troubleshooting by tracking user activity and API usage. CloudTrail logs, continuously monitors, and retains account activity related to actions across your AWS infrastructure, giving you control over storage, analysis, and remediation actions.

108.1.1 Features

- **AWS CloudTrail Event History**: Event history provides a viewable, searchable, downloadable, and immutable record of the past 90 days of management events in an AWS Region. There are no CloudTrail charges for viewing Event history.
- **Always on**: CloudTrail Event history is enabled on all AWS accounts and records management events across AWS services without the need for any manual setup. With AWS Free Tier, you can view, search, and download the most recent 90-day history of your account's management events at no charge using the CloudTrail console or by using the CloudTrail lookup-events API.
- **AWS CloudTrail trails**: Trails capture a record of AWS account activities, delivering, and storing these events in Amazon S3, with optional delivery to Amazon CloudWatch Logs and Amazon EventBridge. These events can be fed into your security monitoring solutions. You can use your own third-party solutions or solutions such as Amazon Athena for searching and analyzing logs captured by CloudTrail. You can create trails for a single AWS account or for multiple AWS accounts by using AWS Organizations.
- **Storage and monitoring**: You can deliver your ongoing management and data events to S3 and optionally to CloudWatch Logs by creating trails. By doing this, you get the complete event details, and you can export and store events as you like.
- **Encrypted activity logs**: You can validate the integrity of CloudTrail log files stored in your S3 bucket and detect whether the log files were unchanged, modified, or deleted since CloudTrail delivered them to your S3 bucket. You can use log file integrity validation in your IT security and auditing processes. By default, CloudTrail encrypts all log files delivered to your specified S3 bucket by using S3 server-side encryption (SSE). If necessary, you can also add a layer of security to your CloudTrail log files by encrypting the log files with your AWS Key Management Service (KMS) key. If you have decrypt permissions, S3 automatically decrypts your log files.
- **Multi-Region**: You can configure CloudTrail to capture and store events from multiple AWS Regions in a single location. This configuration certifies that all settings apply consistently across existing and newly launched Regions.

- **Multi-account**: You can configure CloudTrail to capture and store events from multiple AWS accounts in a single location. This configuration verifies that all settings apply consistently across all existing and newly created accounts.

108.1.2 Benefits

- **Aggregate and consolidate multisource events**: With CloudTrail Lake, you can ingest activity events from AWS and sources outside AWS, including other cloud providers, in-house applications, and SaaS applications running in the cloud or on premises.
- **Immutably store audit-worthy events**: In AWS CloudTrail Lake, you can immutably store audit-worthy events. Easily generate audit reports required by internal policies and external regulations.
- **Derive insights and analyze unusual activity**: Detect unauthorized access and analyze activity logs using SQL-based queries or Amazon Athena. Respond with rules-based EventBridge alerts and automated workflows.

108.2 Backup/Restore and Disaster Recovery

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

108.3 Pricing Overview

Please see the AWS UK G Cloud 14 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/cloudtrail/index.html
- **Service quotas:** https://docs.aws.amazon.com/awscloudtrail/latest/userguide/WhatIsCloudTrail-Limits.html
- **Service FAQs:** https://aws.amazon.com/cloudtrail/faqs/

108.5 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.
- Partner Onboarding Guide: This guide is for AWS Partners in the AWS Partner Network who want to onboard their applications and solutions to CloudTrail Lake, and let their customers integrate application events into CloudTrail Lake.
- CloudTrail-Data API Reference: Describes all API operations for AWS CloudTrail event data in detail. Also provides sample requests, responses, and errors for supported web service protocols.

109 AWS CodeArtifact

109.1 Service Overview

AWS CodeArtifact is a fully managed artifact 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 works with commonly used package managers and build tools like Maven and Gradle (Java), npm and yarn (JavaScript), or pip and twine (Python), or NuGet (.NET).

109.1.1 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, NuGet and SwiftPM 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.

109.1.2 Benefits

- **Share artifacts**: Store and share artifacts across accounts, with appropriate levels of access granted to your teams and build systems.
- **Fully managed**: Reduce overhead from setup and maintenance of an artifact server or infrastructure with a fully managed service.
- **Pay-as-you-go**: Only pay for software packages stored, number of requests made, and data transferred out of Region with pay-as-you-go pricing..

109.2 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/.

109.3 Pricing Overview

Please see the AWS UK G Cloud 14 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/codeartifact/index.html
- **Service quotas:** https://docs.aws.amazon.com/codeartifact/latest/ug/service-limits.html
- **Service FAQs:** https://aws.amazon.com/codeartifact/faq/

109.5 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.

110 AWS CodeBuild

110.1 Service Overview

AWS CodeBuild is a fully managed continuous integration service in the cloud. CodeBuild compiles source code, runs tests, and produces packages that are ready to deploy. CodeBuild eliminates the need to provision, manage, and scale your own build servers. CodeBuild automatically scales up and down and processes multiple builds concurrently, so your builds don’t have to wait in a queue. You can get started quickly by using CodeBuild prepackaged build environments, or you can use custom build environments to use your own build tools. With CodeBuild, you only pay by the minute.

110.1.1 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 GitHub, GitHub Enterprise, Bitbucket, or Amazon S3.
- **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.
- **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.

110.1.2 Benefits

- Avoid having to set up, manage, or patch your own build servers.

- Scale capacity automatically so builds aren't waiting in queue to run.
- Pay only for the build minutes that you use.
- Use pre-packaged build environments or your own, and encrypt artifacts with your own keys.

110.2 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/.

110.3 Pricing Overview

Please see the AWS UK G Cloud 14 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/codebuild/index.html
- **Service quotas:** https://docs.aws.amazon.com/codebuild/latest/userguide/limits.html
- **Service FAQs:** https://aws.amazon.com/codebuild/faqs/

110.5 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.

111 AWS CodePipeline

111.1 Service Overview

AWS CodePipeline is a continuous delivery service that enables you to model, visualize, and automate the steps required to release your software. With AWS CodePipeline, you model the full release process for building your code, deploying to pre-production environments, testing your application and releasing it to production. AWS CodePipeline then builds, tests, and deploys your application according to the defined workflow every time there is a code change. You can integrate partner tools and your own custom tools into any stage of the release process to form an end-to-end continuous delivery solution.

111.1.1 Features

- **Workflow modeling**: 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 GitHub, Amazon ECR, or Amazon S3. It can run builds and unit tests in AWS CodeBuild. CodePipeline can deploy your changes using AWS Elastic Beanstalk, Amazon Elastic Container Service (Amazon ECS), or AWS Fargate.
- **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.
- **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.

111.1.2 Benefits

- **Automate your release processes**: CodePipeline fully automates your release process from end to end, starting from your source repository through build, test, and deployment. You can prevent changes from moving through a pipeline by including a manual approval action in any stage except a Source stage. You can release when you want, in the way you want, on the systems of your choice, across one instance or multiple instances.
- **Establish a consistent release process**: Define a consistent set of steps for every code change. CodePipeline runs each stage of your release according to your criteria.
- **Speed up delivery while improving quality**: You can automate your release process to allow your developers to test and release code incrementally and speed up the release of new features to your customers.
- **Use your favorite tools**: You can incorporate your existing source, build, and deployment tools into your pipeline. For a full list of AWS services and third-party tools currently supported by CodePipeline, see Product and service integrations with CodePipeline.
- **View progress at a glance**: You can review real-time status of your pipelines, check the details of any alerts, retry failed stages or actions, view details about the source revisions used in the latest pipeline execution in each stage, and manually rerun any pipeline.
- **View pipeline history details**: You can view details about executions of a pipeline, including start and end times, run duration, and execution IDs.

111.2Backup/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/.

111.3Pricing Overview

Please see the AWS UK G Cloud 14 Pricing Document accompanying this service in the Digital Marketplace.

111.4Service 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/

111.5Technical 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.

112 AWS Compute Optimizer

112.1Service Overview

AWS Compute Optimizer helps you identify the optimal AWS resource configurations, such as Amazon Elastic Compute Cloud (EC2) instance types, Amazon Elastic Block Store (EBS) volume configurations, task sizes of Amazon Elastic Container Service (ECS) services on AWS Fargate, commercial software licenses, and AWS Lambda function memory sizes, using machine learning to analyze historical utilization metrics. Compute Optimizer provides a set of APIs and a console experience to help you reduce costs and increase workload performance by recommending the optimal AWS resources for your AWS workloads.

112.1.1 Features

- **Compute Optimizer console**: The Compute Optimizer console provides you with a global, cross-account view of all resources analyzed by Compute Optimizer and recommendations so that you can quickly identify the most impactful optimization opportunities.
- **Recommendations for optimal EC2 instance types**: Compute Optimizer delivers recommendations to help customers identify optimal EC2 instance types for their workloads. You can use the external metrics ingestion feature to configure AWS Compute Optimizer to ingest EC2 memory utilization metrics from one of the four observability products: Datadog, Dynatrace, Instana, and New Relic. When enabled, Compute Optimizer analyzes your external EC2 memory utilization metrics in addition to your CPU, disk, network, IO, and throughput data to generate EC2 rightsizing recommendations. AWS Compute Optimizer provides EC2 instance type and size recommendations for EC2 Auto Scaling groups with a fixed group size, meaning desired, minimum, and maximum are all set to the same value and have no

scaling policy attached. Additionally, all Auto Scaling group member instances must be of type M, C, R, T, X, I, D, H, Z, G, and P instance families.

- **Cost Savings**: Recommendations with memory data can help you identify additional savings opportunities and ways to improve application performance. The Cost Explorer console and API surfaces a subset of these recommendations that may lead to cost savings, and augments them with customer-specific cost and savings information (such as billing information, available credits, RI and Savings Plans) to help cost-management owners quickly identify savings opportunities through infrastructure rightsizing.
- **Integration with Cost Optimization Hub**: AWS Compute Optimizer is integrated with Cost Optimization Hub. After enabling Cost Optimization Hub, you can view estimated monthly savings that incorporate your discounts in Compute Optimizer.
- **Analyze SQL server licensing options**: AWS Compute Optimizer analyzes your current configurations such as SQL Server edition, licensing options, and specific database level features you are using. Based on the analysis, it determines whether your SQL Server instances are optimized. Lastly, it generates recommendations based the pre-defined optimization criteria and then makes recommendations available in console or through APIs. When you opt in AWS Compute Optimizer, you authorize the service to use AWS service configuration data and CloudWatch Application Insights metrics. For SQL Server, the data includes edition, licensing options, and SQL Server feature configurations tracked by CloudWatch Application Insights.

112.1.2 Benefits

- **Rightsizing**: Rightsize workloads according to your workload preferences through artificial intelligence and machine learning-based analytics to reduce costs by up to 25%.
- **Performance**: Resolve performance issues by implementing recommendations that identify under provisioned resources.
- **Enhance memory utilization**: Increase recommendation savings and visibility into memory utilization by enabling Amazon CloudWatch metrics.
- **Cost optimization**: Optimize licensing cost through automated license optimization recommendations after authentication.

112.2 Backup/Restore and Disaster Recovery

Please see the AWS UK G Cloud 14 Pricing Document accompanying this service in the Digital Marketplace.

112.3 Pricing Overview

Please see the AWS UK G Cloud 14 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/compute-optimizer/
- **Service quotas:** https://docs.aws.amazon.com/general/latest/gr/compute-optimizer.html
- **Service FAQs:** https://aws.amazon.com/compute-optimizer/faqs/

112.5Technical 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.

113 AWS Config

113.1Service Overview

AWS Config is a fully managed service that provides you with resource inventory, configuration history, and configuration change notifications for security and governance. With AWS Config, you can discover existing AWS resources, record configurations for third-party resources, export a complete inventory of your resources with all configuration details, and determine how a resource was configured at any point in time. These capabilities use compliance auditing, security analysis, resource change tracking, and troubleshooting.

113.1.1 Features

- **Configuration history of resources**: AWS Config records details of changes to your 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 Simple Storage Service (S3) bucket that you specify.
- **Configuration history of software**: AWS Config helps you record software configuration changes within your Amazon Elastic Compute Cloud (EC2) instances and servers running on-premises, 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.
- **Resource relationships tracking**: AWS Config discovers, maps, and tracks AWS resource relationships in your account. For example, if a new EC2 security group is associated with an EC2 instance, AWS Config records the updated configurations of both the EC2 security group and the EC2 instance.
- **Configurable and customizable rules**: AWS Config provides you with pre-built rules evaluating the configurations of your cloud resources, and software within managed instances, including EC2 instances, and resources outside AWS like servers running on premises, before and after provisioning. You can customize pre-built rules to evaluate your AWS resource configurations and configuration changes, or create custom rules on 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.

- **Multi-account, multi-Region data aggregation**: AWS Config aggregators enable centralized auditing and governance that provides an enterprise-wide view of your AWS resource configuration metadata and AWS Config rule compliance status. You can set up AWS Config aggregators for "single account and multi-Regions," "multi-account and multi-Regions," or "all the accounts in an AWS Organization that have AWS Config enabled."
- **Querying configuration state**: AWS Config advanced queries enable you to search the current configuration state of AWS resources based on configuration properties. With advanced queries you can search within a single account and AWS Region or query against an AWS Config aggregator to search from a central account across accounts, AWS Regions, or an AWS Organization.
- **Extensibility**: AWS Config supports extensibility by helping you publish the configuration of third-party resources into AWS Config using our public API operations. Examples of third-party resources include version control systems such as GitHub, Microsoft Active Directory resources, or any on-premises server. AWS Config helps you view and monitor the resource inventory and configuration history of these third-party resources using the AWS Config console and API operations, like you do for AWS resources.
- **Configuration snapshots**: AWS Config can provide you with a configuration snapshot, which is a point-in-time capture of all your resources and their configurations. Configuration snapshots are generated on demand by using the AWS CLI or API and delivered to the Amazon S3 bucket that you specify.
- **Cloud governance dashboard**: AWS Config provides three types of dashboards. First, an AWS account- and AWS Region-specific dashboard that displays your resources' compliance posture, to spot noncompliance resources and take appropriate action. Second, a high-level dashboard per aggregator that shows insights such as total count of non-compliant rules across your AWS Organization, the top five non-compliant rules by number of resources, and the top five AWS accounts that have the highest number of non-compliant rules. Third, detailed dashboards for each aggregator that dive deep into inventory and compliance details, such as compliance summary by resources, top 10 accounts with non-compliant resources, top 10 accounts by noncompliant rules across conformance packs, comparison of running vs. stopped EC2 instances by type, and EBS volumes by volume type and size.
- **Partner solutions**: You can choose from numerous AWS Partner Network (APN) partners who provide solutions that integrate with AWS Config for resource discovery, change management, compliance, or security.

113.1.2 Benefits

- **Manage resource configuration changes**: Continually assess, monitor, and record resource configuration changes to simplify change management.
- **Evaluate configurations against desired state**: Audit and evaluate compliance of your resource configurations with your organization’s policies on a continual basis.
- **Simplify troubleshooting and remediation**: Simplify operational troubleshooting by correlating configuration changes to particular events in your account.

113.2 Backup/Restore and Disaster Recovery

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

113.3 Pricing Overview

Please see the AWS UK G Cloud 14 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/config/
- **Service quotas:** https://docs.aws.amazon.com/general/latest/gr/awsconfig.html
- **Service FAQs:** https://aws.amazon.com/config/faq/

113.5 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.

114 AWS Control Tower

114.1 Service Overview

AWS Control Tower helps to set up and govern a secure, multi-account AWS environment. It establishes a landing zone that is based on best-practices blueprints, and enables governance using controls you can choose from a pre-packaged list. The landing zone is a well-architected, multi-account baseline that follows AWS best practices. Controls implement governance rules for security, compliance, and operations.

114.1.1 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 the following: Create a multi-account environment using AWS Organizations; Provide identity management using the default directory found within AWS IAM Identity Center; Provide federated access to accounts using IAM Identity Center; Centralize logging from AWS CloudTrail and AWS Config stored in Amazon Simple Storage Service (Amazon S3); Enable cross-account security audits using IAM Identity Center.
- **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 by using the AWS Control Tower predefined account

blueprint with default resources, configurations, or VPC settings. You can also define and implement your own custom account resources and requirements in addition to the preapproved account configurations. By configuring your account factory with preapproved network configuration and AWS Region selections, you enable self-service for your builders to configure and provision new accounts.

- **Comprehensive Controls Management**: Comprehensive controls management in AWS Control Tower helps you reduce the time it takes to define, map, and manage the controls required to meet your most common control objectives such as enforcing least privilege, restricting network access, and enforcing data encryption.
- **Dashboard**: The AWS Control Tower dashboard gives you continuous visibility into your AWS environment. You can view the number of organisational units and accounts provisioned and the number of controls enabled and check the status of your organisational units and accounts against those controls. You can also see a list of noncompliant resources with respect to enabled controls.
- **Solutions for AWS Control Tower in AWS Marketplace**: AWS Marketplace 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).

114.1.2 Benefits

- **Set up multi-account environment quickly**: Set up a well-architected, multi-account environment in under 30 minutes.
- **Automation and built-in governance**: Automate the creation of AWS accounts with built-in governance.
- **Preconfigured Controls**: Enforce best practices, standards, and regulatory requirements with preconfigured controls.
- **Integrate third-party software at scale**: Seamlessly integrate third-party software at scale to enhance your AWS environment.

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 14 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/controltower/index.html
- **Service quotas:** https://docs.aws.amazon.com/controltower/latest/userguide/limits.html
- **Service FAQs:** https://aws.amazon.com/controltower/faqs/

114.5 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.
- API Reference: Describes all the API operations for AWS Control Tower in detail. Also provides details of request and response syntax errors for the supported web services protocols.

115 AWS Cost Explorer

115.1 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 analyze 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.

115.1.1 Features

- **Cost and Usage Visualization**: AWS Cost Explorer provides users with visual representation of their AWS cost and usage over time, using graphs and tables. Visualization simplifies analysis of complex cost and usage data, facilitating quick insights and informed decisions.
- **Filter and Group**: You can filter and group your data by a variety of dimensions. For example, you can view your daily/monthly cost grouped by member accounts to understand who is driving your most cost.
- **Cost and usage Forecasting**: AWS Cost Explorer can predict future costs and usage based on your historical and current cost and usage pattern. This helps organizations plan and budget for upcoming expenses.
- **Saved Reports**: Cost Explorer provides preconfigured reports to monitor monthly cost by service, monthly cost by linked account, monthly EC2 running hours costs and usage, Daily costs, and AWS Marketplace costs.
- **Multi-year history**: Users can perform multi-year historical data analysis at monthly granularity using up to 38 months of history and understand patterns and year-over-year changes in their AWS cost and usage.
- **Hourly granularity**: Cost Explorer allows customers to understand their spend for the past 14 days at the most granular hourly level. You can monitor your hourly cost and usage to determine appropriate SP/RI commitment or understand peak hour usage.
- **Resource level granularity**: Cost Explorer allows customers to monitor their AWS cost and usage at a resource level to identify cost drivers. Cost Explorer supports resource level data for all services at daily granularity and EC2 resource level data at hourly granularity. Customers can enable this additional data for the past 14 days.
- **Download your CE view in csv format**: You can download your filtered CE view in CSV file format for offline access as well as easy sharing with other team members.
- **API access**: AWS Cost Explorer provides programmatic access to all features through API, enabling automation of Cost Management tasks.

115.1.2 Benefits

- **Cost analysis**: AWS Cost Explorer lets you explore your AWS costs and usage at both a high level and at a detailed level of analysis, and empowering you to dive

deeper using a number of filtering dimensions (e.g., AWS Service, Region, Member Account, etc.) AWS Cost Explorer also gives you access to a set of default reports to help you get started, while also allowing you to create custom reports from scratch.

- **Default reports**: Cost Explorer supports five default reports through Reports page – monthly cost by service, monthly cost by linked account, Monthly EC2 running hours costs and usage, Daily costs, and AWS Marketplace costs. A Customer can click on any of these reports to open them in Cost Explorer console.
- **Cost forecasts**: Cost Explorer forecasts your AWS cost and usage by analyzing your historical pattern. Cost Explorer estimates your cost and usage for the next 12 months at monthly granularity and for the next 3 months at daily granularity.
- **Easy to set up**: The AWS Cost Explorer API is the low-latency, ad-hoc query service that powers AWS Cost Explorer, and is accessible via a command-line interface and supported AWS SDKs. Using the AWS Cost Explorer API, you can build custom, interactive cost management applications without having to set up and maintain any additional infrastructure.

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 14 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/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-and-usage-reporting/faqs/

115.5 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.

116 AWS Data Exchange (BYOS)

116.1 Service Overview

AWS Data Exchange makes it easy for AWS customers to securely exchange and use third-party data on AWS. AWS Data Exchange has a single global product catalog offered by providers available from any supported AWS Region.

AWS Data Exchange is the only data marketplace with more than 3,500 products from over 300 providers delivered —through files, APIs, or Amazon Redshift queries— directly to the data lakes, applications, analytics, and machine learning models that use it.

With AWS Data Exchange, you can streamline all third-party data consumption, from existing subscriptions —which you can migrate at no additional cost to you or the provider— to future data subscriptions, in one place.

As an AWS service, AWS Data Exchange is secure and compliant, integrated with AWS and third-party tools and services, and offers consolidated billing and subscription management.

116.1.1 Features

- **Data grant in AWS Data Exchange**: A data grant is the unit of exchange in AWS Data Exchange that is created by a data sender in order to grant a data receiver access to a data set. When a data sender creates a data grant, a grant request is sent to the data receiver's AWS account. As a data receiver, you can view all of your current, pending, and expired data grants from the AWS Data Exchange console. You can also discover and subscribe to new third-party data sets available through AWS Data Exchange from the AWS Marketplace catalog.
- **Data Exchange APIs**: For data senders, AWS Data Exchange eliminates the need to build and maintain any data delivery and entitlement infrastructure. Anyone with an AWS account can create and send data grants to data receivers. As a data sender or provider, you can access AWS Data Exchange through the following options: Directly through the AWS Data Exchange console. Data providers with data products available in AWS Marketplace can access programmatically using 1) AWS Data Exchange API—Use the API operations to create, view, update, and delete data sets and revisions. 2) AWS Marketplace Catalog API—Use the API operations to view and update data products published to AWS Marketplace. Anyone, with or without an AWS account, can find and use publicly available data sets that are part of the Open Data on AWS program.
- **Supported data sets**: AWS Marketplace takes a responsible approach to facilitating data transactions by promoting transparency through use of the service. AWS Marketplace reviews permitted data types, restricting products that are not permitted. Providers are limited to distributing data sets that meet the legal eligibility requirements set forth in the Terms and Conditions for AWS Marketplace Sellers.

116.1.2 Benefits

- **Extensive Data Set Selection**: Offers 3,500+ data sets from 300+ data providers with over 1,000 free data products and custom data products. Provides automatic access to new data
- **Streamlined Data Procurement & Governance**: Offers one place to exchange data publicly or privately; You can migrate existing subscriptions at no additional cost with simplified contracts and consolidated secure billing.
- **Better Data Technology**: Provides integration of data into AWS where data is encrypted at rest and in transit.

- **Ease of Use for Data Analytics & Machine Learning**: Supports data files, tables, and APIs with self-service options.

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

Please see the AWS UK G Cloud 14 Pricing Document accompanying this service in the Digital Marketplace.

116.4 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/

116.5 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.

117 AWS Database Migration Service

117.1 Service Overview

AWS Database Migration Service (AWS DMS) is a managed migration and replication service that helps you move your databases and analytics workloads 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.

117.1.1 Features

- **Migration planning**: AWS DMS Fleet Advisor is a free, fully managed capability of AWS Database Migration Service (AWS DMS). It automates migration planning and helps you migrate database and analytics fleets to the cloud at scale with minimal

effort. DMS helps you confidently migrate your databases and analytics systems to AWS with virtually no downtime.

- **AWS DMS Fleet Advisor**: To accelerate migrations, AWS DMS Fleet Advisor automatically inventories and assesses your on-premises database and analytics server fleet and identifies potential migration paths. Using historical performance and usage patterns collected from self-managed databases, Fleet Advisor can recommend target database engines and instance options as well as estimate costs. AWS DMS Fleet Advisor discovers and analyzes the same source databases supported in AWS DMS, including Oracle, Microsoft SQL Server, MySQL, PostgreSQL, and more. DMS Fleet Advisor delivers results in a few hours, instead of weeks or even months, without using third-party tools or hiring migration experts.
- **Schema assessment and conversion**: AWS DMS Schema Conversion (DMS SC) is a fully managed feature of AWS DMS that allows you to automatically assess and convert database schemas and code objects at scale with zero downtime. AWS DMS Schema Conversion supports popular heterogeneous migrations, such as Oracle to Amazon RDS for PostgreSQL, SQL Server to Amazon RDS for MySQL, SQL Server to Amazon Aurora PostgreSQL-Compatible Edition, and Oracle to Amazon Aurora MySQL-Compatible Edition. You can save weeks or months of manual time and resources with a few clicks in the DMS console.
- **Database and analytics migration**: AWS DMS migrates your data—to and from—widely used commercial and open-source databases. It supports homogeneous migrations such as Oracle to Oracle, as well as heterogeneous migrations between different database infrastructures, such as Oracle to Amazon Aurora. AWS DMS supports over 20-plus databases and data warehouses, as source and target endpoints.
- **Serverless**: AWS DMS Serverless offers the flexibility and ease to migrate data without the need to provision replication instances or manually monitor use and adjust capacity. It automatically provisions, monitors, and scales migration resources to the optimal capacity needed to meet demand. DMS Serverless supports popular DMS use cases including continuous data replication, database consolidation, and migrations, even if the source and target database engines differ.
- **Continuous data replication**: With AWS DMS, you can continuously replicate data with low latency from a supported source to a supported target. Continuous data replication has a multitude of use cases, including Disaster Recovery instance synchronization, geographic database distribution, and Dev/Test environment synchronization.

117.1.2 Benefits

- **Simple to use**: AWS DMS 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 from the AWS Management Console.
- **Minimal downtime**: AWS DMS 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 switch over the database at a convenient time.
- **Cost effectiveness**: AWS DMS is a low-cost service. You pay only for the compute resources used during the migration process and any additional log storage. Migrating a terabyte-size database can be done at a low cost. This applies to both homogeneous and heterogeneous migrations of supported databases. This is in

stark contrast to conventional database migration methods that can be very expensive.

- **Reliable**: AWS DMS 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.
- **On-going replication**: You can set up a DMS task for either one-time migration or ongoing replication. An ongoing replication task keeps your source and target databases in sync. Once set up, the ongoing 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 a replication task.
- **Developer productivity**: AWS DMS can be used to migrate data both into and out of the cloud for development purposes. There are two common scenarios. The first is deploying development, test, or staging systems on AWS to take advantage of the cloud's scalability and rapid provisioning. This lets developers and testers use copies of real production data, then copy updates back to the on-premises production system. The second scenario is when development systems are on-premises (often on personal laptops) and you migrate a current copy of an AWS production database to these on-premises systems—either once or continuously. You can avoid a disruption to existing DevOps processes while ensuring the up-to-date representation of your production system.
- **Database Consolidation**: You can use AWS DMS to consolidate multiple source databases into a single target database. This can be done for homogeneous and heterogeneous migrations, and you can use this feature with supported database engines. The source databases can be located in your own premises outside of AWS, running on an Amazon EC2 instance, or it can be an Amazon RDS database. The source databases can also be spread across different locations.

117.2 Backup/Restore and Disaster Recovery

The service provides point-in-time backup and recovery of target database from replication logs stored in Amazon S3. Users can configure backup windows and retention settings.

117.3 Pricing Overview

Please see the AWS UK G Cloud 14 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/dms/
- **Service quotas:** https://docs.aws.amazon.com/general/latest/gr/dms.html
- **Service FAQs:** https://aws.amazon.com/dms/faqs/

117.5 Technical Requirements

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

- Data Migration Guide: Provides a high-level overview of AWS DMS concepts. Includes instructions on migrating data and replicating ongoing changes from

relational databases, data warehouses, NoSQL databases, and other types of data stores into the AWS Cloud.

- 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.

118 AWS DataSync

118.1 Service Overview

AWS DataSync is an online data movement and discovery service that simplifies and accelerates data migrations to AWS as well as moving data to and from on-premises storage, edge locations, other cloud providers, and AWS Storage services.

AWS DataSync Discovery helps you simplify migration planning and accelerate your data migration to AWS by giving you visibility into your on-premises storage performance and utilization, and providing recommendations for migrating your data to AWS Storage services. DataSync Discovery enables you to better understand your on-premises storage performance and capacity usage through automated data collection and analysis, enabling you to quickly identify data to be migrated and use generated recommendations to select AWS Storage services that align to your performance and capacity needs.

118.1.1 Features

- **Discovery**: AWS DataSync Discovery helps you simplify migration planning and accelerate data migration to AWS by giving you visibility into your on-premises storage performance and utilization, and providing recommendations for migrating your data to AWS Storage services, such as Amazon FSx for NetApp ONTAP, Amazon FSx for Windows File Server, and Amazon Elastic File System (EFS). DataSync Discovery enables you to better understand your on-premises storage performance and capacity usage through automated data collection and analysis, enabling you to quickly identify data to be migrated and use generated recommendations to select AWS Storage services that align to your performance and capacity needs.
- **Data Movement**: For online data transfers, AWS DataSync simplifies, automates, and accelerates copying large amounts of data between on-premises storage, edge locations, or other cloud providers, and AWS Storage services.
- **Multicloud Data Movement**: AWS DataSync helps you move data between AWS, on-premises file systems, and other cloud storage services. AWS has continued to extend its cloud services to help customers streamline, manage, and govern their hybrid and multicloud infrastructure and applications. For customers who operate in multicloud environments, AWS DataSync can now move data to and from storage on various clouds.
- **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.

- **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 between the DataSync agent and the DataSync service using Transport Layer Security (TLS). DataSync supports using default at-rest encryption for Amazon S3 buckets. DataSync also supports encryption of data at rest and in transit for Amazon EFS and Amazon FSx.
- **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, to your self-managed object storage or Google Cloud Storage using the Amazon S3 application programming interface (API), or to Azure Blob Storage using the Blob 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 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, Amazon FSx for Lustre, Amazon FSx for OpenZFS, or Amazon FSx for NetApp ONTAP.
- **Integration with AWS Infrastructure and Management Services**: DataSync works 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, and Amazon FSx, DataSync supports AWS Virtual Private Cloud (VPC) endpoints (powered by AWS PrivateLink) to move files directly into your Amazon VPC.
- **Monitoring and Auditing**: DataSync task reports provide JSON-formatted output files that include a summary and detailed reports for all files transferred, skipped, verified, and deleted, enabling you to easily verify and audit the data transfer operations for each task execution. Task reports are generated after the completion of your transfer tasks and they are stored in your Amazon S3 bucket. This allows you to easily use AWS services such as AWS Glue, Amazon Athena, and Amazon QuickSight to automatically catalog, analyze, and visualize task report output to check the progress of your data transfers across all task executions.
- **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.

118.1.2 Benefits

- **Simplify migration planning**: With automated data collection and recommendations, DataSync Discovery can minimize the time, effort, and costs associated with planning your data migrations to AWS. You can use recommendations to inform your budget planning and re-run discovery jobs to validate your assumptions as you approach your migration.
- **Automate data movement**: DataSync makes it easier to move data over the network between storage systems and services. DataSync automates both the management of data-transfer processes and the infrastructure required for high performance and secure data transfer.
- **Transfer data securely**: DataSync provides end-to-end security, including encryption and integrity validation, to help ensure that your data arrives securely, intact, and ready to use. DataSync accesses your AWS storage through built-in AWS security mechanisms, such as AWS Identity and Access Management (IAM) roles. It also supports virtual private cloud (VPC) endpoints, giving you the option to transfer data without traversing the public internet and further increasing the security of data copied online.
- **Move data faster**: DataSync uses a purpose-built network protocol and a parallel, multi-threaded architecture to accelerate your transfers. This approach speeds up migrations, recurring data-processing workflows for analytics and machine learning, and data-protection processes.
- **Reduce operational costs**: Move data cost-effectively with the flat, per-gigabyte pricing of DataSync. Avoid having to write and maintain custom scripts or use costly commercial transfer tools.

118.2 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/.

118.3 Pricing Overview

Please see the AWS UK G Cloud 14 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/datasync/index.html
- **Service quotas:** https://docs.aws.amazon.com/datasync/latest/userguide/datasync-limits.html
- **Service FAQs:** https://aws.amazon.com/datasync/faqs/

118.5 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.
- AWS DataSync in the AWS CLI Reference: Describes the AWS CLI commands that you can use with DataSync transfers and storage discovery.

119 AWS DeepRacer

119.1 Service Overview

AWS DeepRacer is the fastest way to get rolling with reinforcement learning (RL), literally, with a fully autonomous 1/18th scale race car driven by reinforcement learning, 3D racing simulator, and a global racing league.

Developers can train, evaluate, and tune RL models in the online simulator, deploy their models onto AWS DeepRacer for a real-world autonomous experience and compete in the AWS DeepRacer League for a chance to win the AWS DeepRacer Championship Cup.

119.1.1 Features

- **AWS DeepRacer console**: An AWS Machine Learning service for training and evaluating reinforcement learning models in a three-dimensional simulated autonomous-driving environment.
- **AWS DeepRacer vehicle**: A 1/18th scale RC car capable of running inference on a trained AWS DeepRacer model for autonomous driving.
- **AWS DeepRacer League**: The world's first global, autonomous racing league. Race for prizes, glory, and an opportunity to advance to the World Championship Cup.

119.1.2 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.
- **Access**: The new AWS DeepRacer storefront on amazon.com provides a complete list of every recommended item needed to host an in person race with your AWS DeepRacer device. From cars to tracks, batteries to zip ties, we've taken the guess work out of planning so you can focus on racing.

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 14 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/deepracer/

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

119.5Technical 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.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.

120 AWS Device Farm

120.1Service 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.

120.1.1 Features

- **Automated app testing**: Device Farm allows you to upload your own tests or use built-in, script-free compatibility tests. Because testing is performed in parallel, tests on multiple devices begin in minutes. As tests are completed, a test report that contains high-level results, low-level logs, pixel-to-pixel screenshots, and performance data is updated.
- **Supported test types**: Device Farm supports testing of native and hybrid Android and iOS apps, including those created with PhoneGap, Titanium, Xamarin, Unity, and other frameworks. It supports remote access of Android and iOS apps for interactive testing.
- **Remote access interaction**: Remote access allows you to swipe, gesture, and interact with a device through your web browser in real time. There are a number of situations where real-time interaction with a device is useful. For example, customer service representatives can guide customers through the use or setup of their device. They can also walk customers through the use of apps running on a specific device. You can install apps on a device running in a remote access session and then reproduce customer problems or reported bugs.

- **Session logs**: During a remote access session, Device Farm collects details about actions that take place as you interact with the device. Logs with these details and a video capture of the session are produced at the end of the session.

120.1.2 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.

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 14 Pricing Document accompanying this service in the Digital Marketplace.

120.4Service 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/

120.5Technical 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.

121 AWS Digital Investigation and Forensics Storage (created for G-Cloud)

121.1Service 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.

121.1.1 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.

121.1.2 Benefits

- **Automation**: 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.
- **Cost-effective Storage**: 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.

121.2 Backup/Restore and Disaster Recovery

Amazon S3 offers a highly durable, scalable, and secure solution for backing up and archiving your critical data. This includes 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.

121.3 Pricing Overview

Please see the AWS UK G Cloud 14 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:**
    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/

121.5Technical 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.

122 AWS Direct Connect

122.1Service Overview

AWS Direct Connect is a networking service that provides an alternative to using the internet to connect to AWS. Using AWS Direct Connect, data that would have previously been transported over the internet is delivered through a private network connection between your facilities and AWS. In many circumstances, private network connections can reduce costs, increase bandwidth, and provide a more consistent network experience than internet-based connections. All AWS services, including Amazon Elastic Compute Cloud (EC2), Amazon Virtual Private Cloud (VPC), Amazon Simple Storage Service (S3), and Amazon DynamoDB can be used with AWS Direct Connect.

122.1.1 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, AWS 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.

122.1.2 Benefits

- **Improve Application Performance**: Improve application performance by connecting directly to AWS and bypassing the public internet.
- **Reduce networking costs**: Low rates for transferring data out of AWS may reduce your networking costs compared to other AWS services.
- **Protect your data in transit**: Secure your data as it moves between your network and AWS with multiple encryption options.

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 14 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/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

122.5 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.

123 AWS Directory Service

123.1 Service Overview

AWS Directory Service is a managed service offering, providing directories that contain information about your organization, including users, groups, computers, and other resources. As a managed offering, AWS Directory Service is designed to reduce management tasks, thereby allowing you to focus more of your time and resources on your business. There is no need to build out your own complex, highly-available directory topology because 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. There is no software to install and AWS handles all of the patching and software updates.

123.1.1 Features

- **Multiple Availability Zones**: Since directories are mission-critical infrastructure, AWS Managed Microsoft AD is deployed in highly available AWS infrastructure and across multiple Availability Zones. Domain controllers are deployed across two Availability Zones in a region by default and connected to your Amazon Virtual Private Cloud (VPC). Backups are automatically taken once per day, and the Amazon Elastic Block Store (EBS) volumes are encrypted so that data is secured at rest. Domain controllers that fail are automatically replaced in the same Availability Zone using the same IP address, and disaster recovery can be performed using the latest backup.

- **Scale out with additional domain controllers**: When you first create your directory, AWS Managed Microsoft AD deploys two domain controllers across multiple Availability Zones, which is required for highly availability purposes. Later, you can deploy additional domain controllers via the AWS Directory Service console by specifying the total number of domain controllers that you want. AWS Managed Microsoft AD distributes the additional domain controllers to the Availability Zones and VPC subnets on which your directory is running.
- **Managed AD infrastructure**: AWS Managed Microsoft AD runs on AWS managed infrastructure powered by Windows Server 2019. When you select and launch this directory type, it is created as a highly available pair of domain controllers connected to your virtual private cloud (VPC). The domain controllers run in different Availability Zones in a Region of your choice. Host monitoring and recovery, data replication, snapshots, and software updates are configured and managed for you according to the Service Level Agreement (SLA) for AWS Directory Service.
- **Daily snapshots**: AWS Managed Microsoft AD provides built-in, daily, automated snapshots. You can also take additional snapshots before critical application updates to make sure you have the most recent data in case you need to roll back a change.
- **Share directory with multiple AWS accounts**: AWS Managed Microsoft AD integrates tightly with AWS Organizations to allow seamless directory sharing across multiple AWS accounts. You can share a single directory with other trusted AWS accounts within the same organization or share the directory with other AWS accounts that are outside your organization. You can also share your directory when your AWS account is not currently a member of an organization.
- **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.
- **Group-based policies**: AWS Managed Microsoft AD allows you to manage users and devices using native Microsoft Active Directory Group Policy Objects (GPOs). You can create GPOs with existing tools, such as the Group Policy Management Console (GPMC).
- **Schema extension**: You can extend your AWS Managed Microsoft AD schema by adding new object classes and attributes. You can also use schema extensions to enable support for applications that rely on specific Active Directory object classes and attributes. This can be especially useful in the case where you need to migrate corporate applications that are dependent on AWS Managed Microsoft AD, to the AWS cloud.
- **Single sign-on**: 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 can sign in with SSO to AWS applications and resources using a single set of credentials.
- **Directory security settings**: You can configure fine-grained directory settings for your AWS Managed Microsoft AD to meet your compliance and security requirements without any increase in operational workload. In directory settings, you can update secure channel configuration for protocols and ciphers used in your directory.
- **Server-side LDAPS**: Server-side LDAPS encrypts LDAP communications between your commercial or homegrown LDAP-aware applications (acting as LDAP clients) and AWS Managed Microsoft AD (acting as an LDAP server).

- **Client-side LDAPS**: Client-side LDAPS encrypts LDAP communications between AWS applications such as WorkSpaces (acting as LDAP clients) and your self-managed Active Directory (acting as LDAP server).
- **AWS Private CA Connector for Active Directory (AD)**: AWS Managed Microsoft AD and AD Connector integration with AWS Private Certificate Authority (AWS Private CA) Connector for AD allows you to enroll AD domain-joined objects, including users, groups and machines, with certificates issued by AWS Private CA. You can use AWS Private CA as a drop-in replacement for your self-managed enterprise CAs without the need to deploy, patch, or update local agents or proxy servers. You can set up AWS Private CA integration with your directory in just a few clicks or programmatically via API.
- **Directory status monitoring**: Using Amazon Simple Notification Service (Amazon SNS), you can receive email or text (SMS) messages when the status of your directory changes. You get notified if your directory goes from an Active status to an Impaired or Inoperable status. You also receive a notification when the directory returns to an Active status.
- **Domain controller metrics**: AWS Directory Service integrates with Amazon CloudWatch to help provide you with important performance metrics for each domain controller in your directory. This means that you can monitor domain controller performance counters, such as CPU and memory utilization. You can also configure alarms and initiate automated actions to respond to periods of high utilization.
- **Domain controller metrics**: AWS Directory Service integrates with Amazon CloudWatch to help provide you with important performance metrics for each domain controller in your directory. This means that you can monitor domain controller performance counters, such as CPU and memory utilization. You can also configure alarms and initiate automated actions to respond to periods of high utilization.
- **Federated access to the AWS account and applications**: You can grant your on-premises AD users access to sign in to the AWS Management Console and AWS CLI with their existing AD credentials with AWS Identity Center (successor to AWS SSO) by selecting AWS Managed Microsoft AD as the identity source. This enables your users to assume one of their assigned roles at sign-in, and to access and take action on the resources according to the permissions defined for the role. An alternative option is using AWS Managed Microsoft AD to enable your users to assume an AWS Identity and Access Management (IAM) role.
- **Seamless directory integration with AWS applications**: 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 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.

123.1.2 Benefits

- **Simplify your Active Directory (AD) operations**: Simplify your AD operations, availability, maintenance, and on-premises workload management.
- **Streamline your time to market from on-premises to AWS**: Streamline time to market by securely migrating directory-aware workloads to AWS.
- **Provide existing AD users and groups with access to AWS**: Provide user experiences with less friction by using your existing Microsoft AD credentials to access AWS resources.

- **Improve you security posture and resiliency**: Improve your security posture and resiliency with advanced active directory monitoring, logging, and networking.

123.2 Backup/Restore and Disaster Recovery

The service provides backup and recovery of directory snapshots. Users control this via a manual or scheduled API call. Users schedule and recover backups through a web interface.

123.3 Pricing Overview

Please see the AWS UK G Cloud 14 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/directory-service/
- **Service quotas:** https://docs.aws.amazon.com/general/latest/gr/ds_region.html
- **Service FAQs:** https://aws.amazon.com/directoryservice/faqs/

123.5 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.

124 AWS Elastic Beanstalk

124.1 Service Overview

AWS Elastic Beanstalk makes it even easier for developers to quickly deploy and manage applications in the AWS Cloud. Developers simply upload their application, and Elastic Beanstalk automatically handles the deployment details of capacity provisioning, load balancing, auto-scaling, and application health monitoring.

124.1.1 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.

- **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.

124.1.2 Benefits

- **Quickly launch web applications**: Upload and deploy web applications in a simplified, fast way. Focus on writing code instead of provisioning and managing infrastructure.
- **Control**: Select and retain full control of the optimal AWS resources for powering your applications.
- **Scale**: Use adjustable settings to scale your application for handling peaks in traffic, while minimizing costs.

124.2 Backup/Restore and Disaster Recovery

The service provides backup and recovery of machine images and volumes. Users control this via a manual or scheduled API call. Users schedule and recover backups through a web interface.

124.3 Pricing Overview

Please see the AWS UK G Cloud 14 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/elastic-beanstalk/
- **Service quotas:** https://docs.aws.amazon.com/general/latest/gr/elasticbeanstalk.html
- **Service FAQs:** https://aws.amazon.com/elasticbeanstalk/faqs/

124.5 Technical Requirements

Please refer to https://docs.aws.amazon.com/elastic-beanstalk/ 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.

125 AWS Elastic Disaster Recovery

125.1 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.

125.1.1 Features

- **Launch management settings**: The launch management settings in AWS DRS let you control how recovery instances will be launched for your source servers. You can define default launch settings to apply to newly added source servers, as well as modify launch settings in bulk for existing source servers, providing a simple method for you to manage source servers at scale. For applications running on AWS, you can also modify the recovery Availability Zone (AZ) for multiple source servers, which helps simplify cross-AZ recovery.
- **Post-launch actions**: Using post-launch actions on AWS DRS, you can define actions that run automatically after launching recovery instances. You can create custom actions for any AWS Systems Manager (AWS SSM) command or automation, including documents you create or that were shared with you. You can also use a pre-defined set of actions, such as installing a CloudWatch agent, volume

free space detection, instances network connectivity verification, volume integrity validation or creating an Amazon Machine Image from the instance.

- **Automated network replication and recovery**: AWS DRS makes it easy to replicate and recover your AWS network components and configurations to maintain the readiness and security of your AWS recovery site. These components include subnet settings (such as CIDR), security groups, route tables, Internet gateways, and network ACLs. Automating the replication of your network (VPC) configuration substantially increases the chances of a smooth and rapid recovery. It also enhances security and saves time and resources.
- **Recover into existing instance**: AWS DRS allows you to recover into an original or pre-defined existing instance instead of spinning up a new Amazon EC2 instance. By being able to recover into an existing instance, you can retain the existing metadata and security parameters associated with the instance.

125.1.2 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.

125.2 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.

125.3 Pricing Overview

Please see the AWS UK G Cloud 14 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/drs/index.html
- **Service quotas:** https://docs.aws.amazon.com/general/latest/gr/drs.html
- **Service FAQs:** https://aws.amazon.com/disaster-recovery/faqs/

125.5 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.

126 AWS Elemental MediaConnect

126.1 Service Overview

AWS Elemental MediaConnect is a high-quality transport service for live video. It delivers the reliability and security of satellite and fiber-optic combined with the flexibility, agility, and economics of IP-based networks.

126.1.1 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.
- **Transmit compressed live video between on-premises multicast environments and the cloud using AWS Elemental MediaConnect Gateway**: With AWS Elemental MediaConnect Gateway, a cloud-connected software application, you can send video to and from your customer-managed multicast infrastructure. MediaConnect Gateway sends multicast feeds from your network operations center as unicast, transporting the video to MediaConnect. Once there, you can replicate your video and send it to other AWS Regions, process it using AWS Media Services or other applications, share it with partners and affiliates, and even deliver it to other on-premises MediaConnect Gateway installations.
- **Transport uncompressed and visually-lossless video with AWS Cloud Digital Interface (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 allow listing 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. You can also use Amazon EventBridge to build event-driven applications to respond to issues impacting your MediaConnect flows.

- **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.

126.1.2 Benefits

- **High-quality video**: Achieve uninterrupted, high-quality video delivery by adding a quality-of-service layer over standard IP transport.
- **Security**: Secure your live video and control distribution using industry-standard encryption.
- **Resilience**: Gain resilience at the signal and flow levels, and build workflows without any single failure points.
- **Cost effective**: Save 30% or more compared to a typical satellite primary distribution use case, which has 70 destinations. Satellite and fiber solutions have long lead times, limited agility, require up-front commitments, and lack the scale to meet demand as workloads change over time. MediaConnect is a pay-as-you-go service with predictable costs and you can get started in minutes, not months.

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 14 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/mediaconnect/index.html
- **Service quotas:** https://docs.aws.amazon.com/mediaconnect/latest/ug/quotas.html
- **Service FAQs:** https://aws.amazon.com/mediaconnect/faqs/

126.5 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.

127 AWS Elemental MediaConvert

127.1 Service Overview

AWS Elemental MediaConvert is a file-based video processing service that formats and compresses offline content for delivery to televisions or connected devices. With AWS

Elemental MediaConvert's high-quality video transcoding, you can create on-demand video assets for playback on virtually any device. 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.

127.1.1 Features

- **Broadcast Capabilities for Video-on-Demand Content**: 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**: 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**: AWS Elemental MediaConvert eliminates the burden of managing video processing infrastructure by automating the key aspects of workload provisioning and management. AWS Elemental MediaConvert handles resource provisioning and optimization, service orchestration, scaling, healing, resiliency failover, monitoring, and reporting.
- **Built-in Reliability**: 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.

127.1.2 Benefits

- **High-quality video**: Produce high-quality video output with broadcast-grade video and audio features.
- **Easy to manage**: Submit jobs with the video processing settings you want without spending resources to manage transcoding infrastructure.
- **Pay-as-you-go pricing**: Pay based only on the duration of video processed and the features you use.

127.2 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/.

127.3 Pricing Overview

Please see the AWS UK G Cloud 14 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:** 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/

127.5 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.

128 AWS Elemental MediaLive

128.1 Service Overview

AWS Elemental MediaLive is a cloud-based live video encoding service that offers fast, reliable and easy-to-use delivery of high-quality live video streams without the need to manage infrastructure. AWS Elemental MediaLive streamlines live video operations by automating the configuration and management of ingest and encoding components for highly reliable delivery of live streams. The service provides broadcast quality features, configurable capability and support for industry standard formats and technologies. Combining broadcast-grade encoding capabilities with the scale and elasticity of AWS, video providers can efficiently deliver live stream to their audiences and focus on their content and differentiated viewing experiences.

128.1.1 Features

- **Comprehensive Video Standards Support**: 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).
- **Broadcast Capabilities for Live Video Streams**: 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.

- **Statistical Multiplexing (Statmux)**: 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.

128.1.2 Benefits

- **Send HD or UHD video**: Use small, portable AWS Elemental Link HD and UHD devices to transfer video to AWS Elemental MediaLive for delivery to viewers.
- **Schedule or dynamically switch inputs**: Add multiple video inputs and switch using the schedule feature, or dynamically switch by creating input switching actions as needed.
- **Automatically deploy a channel**: AWS Elemental MediaLive is easy to use and scales with AWS resources globally. Create a media infrastructure that generates a channel and loops through an .mp4 file using the AWS Cloud Development Kit (CDK).

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 14 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/medialive/index.html
- **Service quotas:** https://docs.aws.amazon.com/medialive/latest/ug/limits.html
- **Service FAQs:** https://aws.amazon.com/medialive/faqs/

128.5 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**.**

129 AWS Elemental MediaPackage

129.1 Service Overview

AWS Elemental MediaPackage is a highly scalable, video origination and just-in-time packaging service that helps video providers securely, reliably and cost-efficiently package and deliver live video streams. Video providers can improve the viewing experience, easily integrate advanced, broadcast-grade capabilities, increase workflow resiliency, and better protect and monetize their multiscreen content. The service uses just-in-time packaging to

cost effectively output multiple standards-based streaming protocols and DRM types in different combinations to support an array of multiscreen devices. It supports consistent quality of service by elastically scaling to meet demand and manage failover within a highly available managed service. AWS Elemental MediaPackage does not limit customers' choices of video players, CDNs or ad providers, and works seamlessly with other AWS services to build a solution for high-quality, resilient live streaming for 24/7 channels or live events.

129.1.1 Features

- **Comprehensive Output Formats**: 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**: 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**: 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**: 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.

129.1.2 Benefits

- **Maximize viewership**: Deliver video content to the maximum number of viewers on multiple playback devices
- **Distribute content**: Package and distribute your content to a broad range of video playback devices.
- **Automate resource scalability**: Automatically manage resources across multiple Availability Zones, and detect and resolve issues without disrupting streams.
- **Support DRM providers**: Configure a DVR-like experience for viewers and support a range of digital rights management (DRM) providers.
- **Integrate with CDN**: Integrate with the Amazon CloudFront content delivery network (CDN) and AWS Elemental MediaLive.

129.2 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/.

129.3 Pricing Overview

Please see the AWS UK G Cloud 14 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/mediapackage/
- **Service quotas:** https://docs.aws.amazon.com/mediapackage/latest/userguide/quotas.html
- **Service FAQs:** https://aws.amazon.com/mediapackage/faqs/

129.5 Technical Requirements

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

- User Guide: Describes the components and features that AWS Elemental MediaPackage version 2 provides and how to use them.
- Live API Reference: Describes all AWS Elemental MediaPackage API operations for live content in detail. Also provides sample requests, responses, and errors for the supported web services protocols.

130 AWS Elemental MediaStore

130.1 Service Overview

AWS Elemental MediaStore is a media origin and storage service that offers the performance, predictable low latency, and consistency required for delivery and processing workloads like live streaming video. The service provides a write-behind cache, designed for performance, in front of object storage. It is an inexpensive method for pass-through and low-latency segmented content delivery, with predictable pay-as-you-go pricing.

130.1.1 Features

- **Performance when you need it**: 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**: 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.
- **Access Control Management**: Integrated security management tools let you create Identity and Access Management (IAM) roles, improving security and restricting access to content to specific roles.
- **Use Independently or with AWS Media Services**: 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.

130.1.2 Benefits

- **High quality**: Deliver content optimized for high-scale, high-quality media workloads with consistent service quality.
- **Scalability**: Scale automatically so that increases in load don't negatively impact the quality of the viewing experience.
- **Secure access**: Use AWS Identity and Access Management (IAM) to specify granular access controls.

130.2 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/.

130.3 Pricing Overview

Please see the AWS UK G Cloud 14 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/mediastore/index.html
- **Service quotas:** https://docs.aws.amazon.com/mediastore/latest/ug/quotas.html
- **Service FAQs:** https://aws.amazon.com/mediastore/faqs/

130.5 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.

131 AWS Elemental MediaTailor

131.1 Service Overview

AWS Elemental MediaTailor is a channel assembly and personalized ad insertion service that lets video providers create live OTT (internet delivered) channels using existing video content and monetize those channels, or other live streams and VOD content, with personalized advertising. Live streams maintain a TV-like experience across multiscreen video applications. With MediaTailor, virtual live channels are created without the expense, complexity, and management of real-time live encoding. Adverts are seamlessly stitched into the content and can be tailored to individual viewers, maximizing the monetization opportunity for every ad break and mitigating ad blocking schemes.

131.1.1 Features

- **Server-side Ad Insertion**: 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**: 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**: 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**: 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**: 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**: 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.

131.1.2 Benefits

- **Low costs**: Create live streams with low running cost by using existing encoded content.
- **Enhanced audience engagement**: Boost audience engagement by delivering relevant program content and personalized ads.
- **Seamless viewing**: Give viewers a seamless, TV-like viewing experience without buffering between program content and ad breaks.
- **Opportunities to monetize**: Optimize monetization opportunities with personalized ads and assembled linear channels.

131.2 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/.

131.3 Pricing Overview

Please see the AWS UK G Cloud 14 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/mediatailor/index.html
- **Service quotas:** https://docs.aws.amazon.com/mediatailor/latest/ug/quotas.html
- **Service FAQs:** https://aws.amazon.com/mediatailor/faqs/

131.5 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.

132 AWS Fargate

132.1 Service Overview

AWS Fargate is a serverless compute engine for containers that works with both Amazon Elastic Container Service (ECS) and Amazon Elastic Kubernetes Service (EKS). AWS Fargate makes it easy to focus on building your applications by eliminating the need to provision and manage servers, lets you specify and pay for resources per application, and improves security through application isolation by design.

132.1.1 Features

- **Serverless compute for containers**: AWS Fargate manages capacity needs, operating system (OS) updates, compliance requirements, resiliency, and more, freeing you to focus on applications, not servers. Deploy any OCI compliant container image to AWS Fargate and move the undifferentiated heavy lifting of managing the underlying infrastructure to AWS.
- **Serverless compute for stateful workloads**: For customers running stateful workload in containers, you can leverage Amazon EFS with AWS Fargate to externalize data outside of your application.
- **Seamless scaling with serverless compute**: Respond to changes in demand seamlessly by dynamically scaling capacity to reduce wasted resources while ensuring application availability. Utilize AWS Auto Scaling's scheduled scaling, target tracking, and step scaling techniques, and Kubernetes' Horizontal Pod Autoscaler (HPA) and Vertical Pod Autoscaler (VPA) components with AWS Fargate.
- **Security and Compliance**: For sensitive, multi-tenant, or regulated containerized workloads, AWS Fargate provides a secure virtualization boundary between every Amazon ECS task or Amazon EKS pod. Each task/pod is deployed on to a dedicated single use, single tenant piece of compute, providing secure isolation between workloads.

- **Secure and monitor network traffic with Amazon Virtual Private Networks (VPC)**: AWS Fargate connects applications to existing investments in Virtual Private Networks (VPC), providing load balancing, segmentation through VPC security groups, and traffic monitoring with VPC Flow Logs. For client-based routing and out of the box traffic observability, workloads can be configured with Amazon ECS Service Connect for inter-service communication.
- **Compliant foundation for regulated industries**: For customers that need to run containerized workloads in regulated or controlled environments, AWS Fargate has been approved by numerous industry-standard compliance programs.
- **Secure terminal access for debugging**: To help with troubleshooting and debugging Amazon ECS Exec provides secure access to running containers without logging into SSH to underlying servers. Amazon ECS Exec commands are recorded for audit purposes in AWS CloudTrail and Amazon CloudWatch.
- **Extending AWS Fargate**: Containerized workloads running on AWS Fargate workloads are deeply integrated with AWS services. Workloads are secured with AWS Identity and Access Management, monitored with Amazon CloudWatch Container Insights and scaled with AWS Application Auto Scaling.
- **Modernize .NET applications for improved performance and utility**: You can deploy .NET Framework applications with Windows Containers on Amazon ECS and AWS Fargate,
providing a seamless deployment story, improved monitoring and observability, and removing the management and licensing expenses of maintaining Windows Server instances.
- **Resource-based pricing and per-second billing**: With AWS Fargate you pay by the task size and only for the time for which resources are consumed by the task. By adding a tag to Amazon ECS tasks, you can use AWS Cost Explorer to visualize costs, identify opportunities for savings, and provide granular per workload billing.
- **Flexible pricing models**: Use Compute Savings Plans to lock in savings over one or three year increments, choose between On-Demand and Fargate Spot to blend cost and availability, and select AWS Graviton for improved cost-performance. AWS Compute Optimizer uses machine learning to analyze historical usage and provide recommendations on sizing tasks to reduce cost.

132.1.2 Benefits

- **Manage your applications, not infrastructure**: Deploy and manage your applications, not infrastructure. Remove the operational overhead to scale, patch, help secure, and manage servers.
- **Monitor your applications to gain metrics and insights**: Monitor your applications through built-in integrations with AWS services like Amazon CloudWatch Container Insights or gather metrics and logs with third-party tools.
- **Improve security through isolation**: Improve security through workload isolation by design. Amazon ECS tasks and Amazon EKS pods run in their own dedicated runtime environment.
- **Optimize for cost**: Pay only for compute resources used, with no upfront expenses. Further optimize costs with Savings Plans, Fargate Spot, or AWS Graviton processors.

132.2 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.

132.3 Pricing Overview

Please see the AWS UK G Cloud 14 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/ecs/index.html
- **Service quotas:** https://docs.aws.amazon.com/general/latest/gr/ecs-service.html
- **Service FAQs:** https://aws.amazon.com/fargate/faqs/

132.5 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.

133 AWS Fault Injection Service

133.1 Service Overview

AWS Fault Injection Service (AWS FIS) is a fully managed service for running fault injection experiments on AWS that makes it easier to improve an application's performance, observability, and resilience. 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.

133.1.1 Features

- **Simple setup**: AWS FIS makes it easy to get started building and running fault injection experiments, without needing to install any agents. 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 Service supports Amazon CloudWatch so that you can use your existing metrics to monitor Fault Injection Service experiments.
- **Run real-world scenarios**: Scenarios are provided through the Scenario Library in the FIS console, and are run using an FIS experiment template. To run an experiment using a scenario, select the scenario from the library, copy it to your experiment template, and specify your application details. Each scenario includes a detailed description and suggested metrics to measure the response of your application during the experiment, helping you improve the resilience posture of your applications over time. You can find a list of supported scenarios here.
- **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 FIS allows you to target scenarios 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”.

- **Integrated security model**: AWS FIS 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 Service experiments, and which resources and services can be affected.
- **Visibility throughout an experiment**: FIS provides visibility throughout every stage of an experiment via the console and APIs. To support accurate operational metrics and effective troubleshooting, you can also identify what resources and APIs are affected by a Fault Injection Service experiment.
- **Console and programmatic access**: You can use AWS FIS with the AWS Management Console, AWS CLI, and AWS SDKs. The FIS 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.

133.1.2 Benefits

- **Track performance**: Find performance bottlenecks or other unknown weaknesses missed by traditional software tests.
- **Specify conditions**: Define specific conditions to stop an experiment or roll back to the pre-experiment state.
- **Save time**: Run experiments in minutes using pre-built scenarios from the FIS Scenario Library.
- **Gain insights**: Get superior insights by generating real-world failure conditions, such as impaired performance of different resources.

133.2 Backup/Restore and Disaster Recovery

This service allows for backup and recovery. Experiment templates (json) can be created and exported; Experiment results can be downloaded/stored.

Users control this via a manual or scheduled API call. Users schedule and recover backups through a web interface.

133.3 Pricing Overview

Please see the AWS UK G Cloud 14 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/fis/
- **Service quotas:** https://docs.aws.amazon.com/fis/latest/userguide/fis-quotas.html
- **Service FAQs:** https://aws.amazon.com/fis/faqs/

133.5 Technical Requirements

Please refer to https://docs.aws.amazon.com/fis/ 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.

134 AWS Firewall Manager

134.1 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 Organization. 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.

134.1.1 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.
- **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.
- **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.

- **AWS Marketplace third-party firewall support**: AWS Firewall Manager allows you to centrally deploy and monitor AWS Marketplace subscribed third-party cloud firewalls across all virtual private clouds (VPCs) in your organization. The service is a single firewall management solution to deploy and manage both AWS native firewalls and AWS Marketplace subscribed third-party firewalls. You can automate cross-account deployment of firewalls, association of rules, and configuration of VPC routes, even as new accounts and VPCs are created in your organization.

134.1.2 Benefits

- **Manage rules centrally**: Use a central administrator account to manage firewall rules across multiple AWS accounts. Deploy managed rules, such as pre-configured WAF rules on your applications, across accounts.
- **Automate security policies**: Automatically enforce your defined security policies across existing and newly created resources. Centrally deploy baseline security group rules to protect your virtual private clouds (VPCs).

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 14 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/firewall-manager/
- **Service quotas:** https://docs.aws.amazon.com/waf/latest/developerguide/fms-limits.html
- **Service FAQs:** https://aws.amazon.com/firewall-manager/faqs/

134.5 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.

135 AWS Global Accelerator

135.1 Service Overview

AWS Global Accelerator is a networking service that helps you improve the availability and performance of the applications that you offer to your global users. AWS Global Accelerator is easy to set up, configure, and manage. It provides static IP addresses that provide a fixed

entry point to your applications and eliminate the complexity of managing specific IP addresses for different AWS Regions and Availability Zones.

AWS Global Accelerator always routes user traffic to the optimal endpoint based on performance, reacting instantly to changes in application health, your user's location, and policies that you configure. You can test the performance benefits from your location with a speed comparison tool. AWS Global Accelerator is a self-service, pay-per-use offering, requiring no long term commitments or minimum fees.

135.1.1 Features

- **Global Edge Network**: AWS Global Accelerator uses a global network of 111 Points of Presence in 93 cities across 50 countries.
- **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. 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. 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**: 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.
- **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.
- **Custom routing accelerator**: Custom routing accelerators are a new type of accelerator in AWS Global Accelerator. This new accelerator lets you use your own application logic to route user traffic to a specific Amazon EC2 instance destination in a single or multiple AWS Regions. Because the traffic is routed over the AWS global network, you get all the performance improvements of going through Global Accelerator.
- **Multi-Region Storage**: You can use S3 Multi-Region Access Points to get the benefits of Global Accelerator for your object storage workload. S3 Multi-Region Access Points use Global Accelerator transparently to provide a single global endpoint to access data sets that are replicated across multiple AWS Regions. This allows you to build multi-region applications with the same simple architecture used in a single region, and then to run those applications anywhere in the world.

135.1.2 Benefits

- **Performance**: Improve network performance for your applications by up to 60%.
- **Availability**: Deliver highly available applications with fast failover for multi-Region and multi-AZ architectures.
- **Routing**: Achieve deterministic routing by removing DNS cache dependencies.
- **Security**: Protect your applications from DDoS attacks closer to the source.

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 14 Pricing Document accompanying this service in the Digital Marketplace.

135.4Service Constraints

Information on service features and constraints is available online here:

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

135.5Technical Requirements

Please refer to https://docs.aws.amazon.com/global-accelerator/ 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.
- AWS Global Accelerator section of the AWS CLI Reference: Describes the AWS Global Accelerator commands in the AWS CLI. Provides basic syntax, options, and usage examples for each command.

136 AWS Glue

136.1Service Overview

AWS Glue is a serverless data integration service that makes it easier to discover, prepare, and combine data for analytics, machine learning (ML), and application development. AWS Glue provides all the capabilities needed for data integration, so you can start analyzing your data and putting it to use in minutes instead of months. AWS Glue provides both visual and code-based interfaces to make data integration easier. Users can more 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 in a few steps in AWS Glue Studio.

136.1.1 Features

- **Discover and search across all your AWS datasets**: 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, progress through a prioritized list of classifiers to determine the schema for your data, and then create 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 extract, transform, and load (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, helps you validate and control the evolution of streaming data using registered Apache Avro schemas at no additional charge.

Through Apache-licensed serializers and deserializers, AWS Glue Schema Registry integrates with Java applications developed for Apache Kafka, Amazon Managed Streaming for Apache Kafka (Amazon MSK), Amazon Kinesis Data Streams, Apache Flink, Amazon Kinesis Data Analytics for Apache Flink, and AWS Lambda. When data streaming applications are integrated with AWS Glue Schema Registry, you can improve data quality and safeguard against unexpected changes using compatibility checks that govern schema evolution.

- **Automatically scale based on workload**: Auto Scaling, a serverless feature in AWS Glue, dynamically scales resources up and down based on workload. With Auto Scaling, your job is assigned workers only when needed. As the job progresses, and it goes through advanced transforms, AWS Glue adds and removes resources depending on how much it can split up the workload. You no longer need to worry about over-provisioning resources, spending time optimizing the number of workers, or paying for idle resources.
- **Deduplicate and cleanse data with built-in machine learning (ML)**: AWS Glue helps clean and prepare your data for analysis without you having to become an ML expert. Its FindMatches feature deduplicates and finds records that are imperfect matches of each other.
- **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 integrated development environment (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.
- **Define, detect, and remediate sensitive data**: AWS Glue Sensitive Data Detection helps you define, identify, and process sensitive data in your data pipeline and data lake. Once identified, you can remediate sensitive data by redacting, replacing, or reporting on personally identifiable information (PII) data and other types of data deemed sensitive. AWS Glue Sensitive Data Detection simplifies the identification and masking of sensitive data, including PII such as name, Social Security number, address, email, and driver's license.
- **Scale existing Python code with Ray**: AWS Glue for Ray helps data engineers process large datasets using Python and popular Python libraries. AWS Glue for Ray uses Ray.io, an open-source unified compute framework that helps scale Python workloads from a single node to hundreds of nodes. AWS Glue for Ray is serverless, so there is no infrastructure to manage.
- **Create custom visual transforms**: AWS Glue helps you create custom visual transforms so you can define, reuse, and share ETL logic. With AWS Glue Custom Visual Transforms, data engineers can write and share business-specific Apache Spark logic, reducing dependence on Spark developers and making it simpler to keep ETL jobs up to date. These transforms are available to all jobs in your AWS account, whether visual or code-based.
- **Simplify data integration job development**: AWS Glue Interactive Sessions, a serverless feature of job development, simplifies the development of data integration jobs. With AWS Glue Interactive Sessions, data engineers can interactively explore and prepare data. Engineers can also explore, experiment on, and process data interactively using the IDE or notebook of their choice.
- **Built-in job notebooks**: AWS Glue Studio Job Notebooks provide serverless notebooks with minimal setup in AWS Glue Studio, so developers can get started quickly. With AWS Glue Studio Job Notebooks , you have access to a built-in

interface for AWS Glue Interactive Sessions where you can save and schedule your notebook code as AWS Glue jobs.

- **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.
- **Apply and deploy DevOps best practices with Git integration**: AWS Glue integrates with Git, the widely used open-source version-control system. You can use GitHub to maintain a history of changes to your AWS Glue jobs and apply existing DevOps practices to deploy them. Git integration in AWS Glue works for all AWS Glue job types, whether visual or code-based. It includes built-in integration GitHub and also makes it simpler to use automation tools like Jenkins to deploy AWS Glue jobs.
- **Reduce costs for nonurgent workloads with flexible job execution**: AWS Glue Flex is a flexible execution job class that allows you to reduce the cost of your nonurgent data integration workloads (for example, preproduction jobs, testing, and data loads) by up to 35%. AWS Glue has two job execution classes: standard and flexible. The standard execution class is ideal for time-sensitive workloads that require fast job startup and dedicated resources. AWS Glue Flex is appropriate for non-time-sensitive jobs whose start and completion times may vary.
- **Read, insert, update, and delete files in your data lake**: AWS Glue natively supports three open-source frameworks, including Apache Hudi, Apache Iceberg, and Linux Foundation Delta Lake. These frameworks help you manage data in a transactionally consistent manner for use in your Amazon S3 based data lake.
- **Deliver high-quality data across your data lakes and pipelines**: AWS Glue Data Quality helps you improve your data quality and confidence. It automatically measures, monitors, and manages data quality in your data lakes and pipelines.
- **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.
- **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.

136.1.2 Benefits

- **Support all workloads**: Flexible support for ETL, ELT, batch, streaming and more, with no lock-in
- **Scale on demand**: Petabyte scale, pay-as-you-go billing, any data size
- **Tailored tools**: Support all data users from developers to business users
- **All in one**: Complete data integration capabilities in one serverless service

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

136.5 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.

137 AWS Ground Station

137.1 Service Overview

AWS Ground Station is a fully managed service that enables you to control satellite communications, process satellite data, and scale your satellite operations. This means that you no longer have to build or manage your own ground station infrastructure.

AWS Ground Station enables you to focus on innovating and rapidly experimenting with new applications that ingest satellite data and dynamically scale your server and storage use, rather than spend resources on operating and maintaining your own ground stations.

137.1.1 Features

- **Fully managed global ground station network integrated with AWS Global Infrastructure**: AWS Ground Station antennas are located within fully managed AWS locations and are interconnected via Amazon's low-latency, highly reliable, scalable and secure global network backbone. As a result, downlinked data is synchronously delivered to an AWS Region and can be sent to other AWS Regions over the global network.
- **Schedule satellites and download data using AWS services**: After onboarding your satellite to AWS Ground Station, you can easily schedule contacts with your satellite. Simply access the AWS Management Console via AWS Ground Station, and just select your satellite, a start and end time, and the preferred ground station location. After scheduling your Contact, you launch Amazon EC2 instances to send satellite commands or receive telemetry and/or mission data. When sending satellite commands, an EC2 instance communicates over an elastic network interface (ENI) connection in Amazon VPC for the duration of the contact. Received telemetry and/or

mission data is delivered to EC2 or Lambda for real-time processing or stored in an Amazon S3 bucket for asynchronous processing.

- **Direct access to AWS services**: AWS Ground Station antenna locations provide direct access to AWS services for faster, simpler and more cost-effective storage and processing of downloaded data. This allows you to reduce data processing and analysis times for use cases like weather prediction or natural disaster imagery from hours to minutes or seconds. This also enables you to quickly create business rules and workflows to organize, structure, route the satellite data before it can be analyzed and incorporated into key applications such as imaging analysis and weather forecasting.
- **Support most common satellites and communication frequencies**: You can use AWS Ground Station to connect with satellites in Low Earth Orbit (LEO) and Medium Earth Orbit (MEO) operating in X-band and S-band frequencies, including: S-band uplink (2025-2120 MHz) and downlink (2200-2300 MHz), X-band narrowband and wideband downlink (7750-8400 MHz).
- **Wideband Digital Intermediate Frequency (DigIF)**: Wideband DigIF allows customers to build a global ground segment by centralizing their software-defined radios (SDRs) in a separate AWS Region. AWS Ground Station leverages Wideband Digital Intermediate Frequency (DigIF) to let customers downlink up to five channels (400 MHz total) per polarity.
- **AWS Ground Station console and APIs**: You can use the AWS Management Console or AWS Ground Station APIs to reserve contacts and antenna time for your satellite communications. With AWS Ground Station's console or APIs, you can review, cancel, and reschedule contact reservations up to 15 minutes prior to scheduled antenna times.
- **Pay-per-minute pricing**: Once you complete all licensing and onboard your satellite, you can schedule access to AWS Ground Station antennas on a per-minute basis and pay only for the scheduled time, using On-Demand or Reserved pricing. On-Demand pricing lets you pay for antenna access with no long-term commitments; whereas, Reserved pricing provides a discounted rate and improved scheduling window, compared to On-Demand, with a monthly commitment.

137.1.2 Benefits

- **Ground Station as a Service**: AWS Ground Station provides a global network of ground stations in close proximity to our global network of AWS infrastructure regions. With AWS Ground Station, you no longer need to worry about buying, leasing, building, scaling or managing your own satellite ground stations.
- **Security you can trust**: Security at AWS is the highest priority. As an AWS customer, you will benefit from facility and network architecture built to meet the requirements of the most security-sensitive organizations. AWS Ground Station provides premium data security and physical security without any additional costs.
- **Pay-as-you-go**: With AWS Ground Station, you only pay for the actual antenna time you use. There are no long-term contracts, or hidden fees. With a single price, you can use any antenna in the global AWS Ground Station network.
- **Faster data downloads**: With multiple antennas at each of our geographically diverse AWS Ground Stations you can downlink your satellite data directly into an AWS region for immediate processing. Our Ground Station scheduling interface allows you to get your data when you need it without scheduling delays or antenna conflicts.
- **Immediate data processing**: AWS Ground Station provides satellite antennas in close proximity to AWS infrastructure regions, giving you low-latency and low-cost

access to AWS services to store and process your data. This allows you to reduce data processing and analysis times for use cases like weather prediction or natural disaster imagery from hours to minutes or seconds.

- **Self-service scheduling**: You can easily schedule Contacts with your satellite using the AWS Management Console and APIs. Once reserved, Contacts can be rescheduled or cancelled up to 15 minutes prior to start.

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 14 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/ground-station/
- **Service quotas:** https://docs.aws.amazon.com/general/latest/gr/gs.html
- **Service FAQs:** https://aws.amazon.com/ground-station/faqs/

137.5 Technical Requirements

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

- User Guide: Describes the conceptual framework of AWS Ground Station and includes instructions on managing satellite operations with the AWS Ground Station console.
- API Reference: Describes all the API operations for AWS Ground Station in detail. Also provides sample requests, responses, and errors for the supported web services protocols.
- AWS Ground Station in the AWS CLI Reference: Describes the Ground Station commands in the AWS CLI that you can use to schedule contacts from a ground station location. Provides syntax, options, and usage examples for each command.

138 AWS Health Dashboard

138.1 Service Overview

AWS Health is the authoritative data source for events and changes affecting your AWS cloud resources. AWS Health notifies you about service events, planned changes, and account notifications to help you manage and take actions. Sign into your AWS Health Dashboard at https://health.aws.amazon.com/health/home, to view account-specific health information or receive AWS Health event updates using Amazon EventBridge. You can also access AWS Health programmatically using AWS Health API, available with AWS Premium Support.

AWS Health Dashboard shows information in the following ways:

- **Your account events** – This page shows events that are specific to your account. You can view open, recent, and scheduled changes. You can also view notifications and an event log that shows all events from the past 90 days.

- **Your organization events** – This page shows events that are specific to your organization in AWS Organizations. You can view open, recent, and scheduled changes for your organization. You can also view notifications, as well as an event log that shows all organization events from the past 90 days.

138.1.1 Features

- **AWS Health Dashboard**: If you have an AWS account, your AWS Health Dashboard shows both public events and account-specific events. We recommend that you use your AWS Health Dashboard to learn about events that provide general awareness, such as an upcoming maintenance issue for a service in a Region. You can also use the AWS Health Dashboard to learn about events that might affect you directly, such as a deprecated resource in your account. You can sign in to the AWS Management Console to view your AWS Health Dashboard at https://health.aws.amazon.com/health/home.
- **AWS Health Dashboard - Service health**: If you don't have an account, you can use the AWS Health Dashboard – Service health at https://health.aws.amazon.com/health/status to view public events. Public events are reported service issues for AWS that provide information about service availability. This website only shows public events, which aren't specific to any account. You don't need to sign in or have an account to view this page.
- **AWS Health API**: You can use the AWS Health API to programmatically access the information that appears in the AWS Health Dashboard, such as the following: Get information about events that might affect your AWS services and resources; Enable or disable the organizational view feature for your AWS organization; Filter your events by specific services, event type categories, and event type codes.
- **Event status**: The event status tells you if the Health event is open, closed, or upcoming. You can view Health events in the AWS Health Dashboard or the AWS Health API for up to 90 days.
- **Event type categories**: All Health events have an associated event type category. For some events, the event type category might appear in the event type code, such as the AWS_RDS_MAINTENANCE_SCHEDULED code.
- **Event type code**: The event type codes shown in a Health event include the affected service and the type of event. For example, if you receive a Health event that has the AWS_EC2_SYSTEM_MAINTENANCE_EVENT event type code, this means that the service is scheduling a maintenance event that might affect you. Use this information to plan ahead or take action for your account.
- **AWS Health events on Amazon EventBridge**: You can setup Amazon EventBridge rules for your accounts to automate actions after the appropriate AWS Health event is received by an account. These can be general actions, such as sending all planned lifecycle event messages to a chat interface. Or, they can be specific actions, such as triggering a workflow in an IT service management tool.
- **Affected entities**: Affected entities are AWS resources that might be affected by the event. For example, if you receive a scheduled event for Amazon EC2 maintenance for a specific instance type that you're using in your account, you can use the Health event to determine the ID of the affected instances. Use this information to address any potential service issue, such as creating or deprecating resources.
- **Organizational view**: You can use this feature to aggregate all health events for AWS accounts in your AWS Organizations into a single view in the AWS Health Dashboard. You can then sign in to the management account of your organization or use the AWS Health API to view all events that might affect the different accounts and resources. You can enable this feature from the AWS Health console or API.

138.1.2 Benefits

- **Authoritative data source for events and changes**: Integrates with 200+ AWS services to provide coverage for operational issues and planned changes.
- **Comprehensive, timely, and actionable**: Get timely, actionable guidance to help you troubleshoot and remediate operational events, and take actions to prepare for scheduled changes.
- **Manage service health at scale**: Provides centralized, aggregated view across your organization in AWS Organizations for teams that manage operational health across the company.
- **Integration and automation**: Receive AWS Health events using Amazon EventBridge or integrate programmatically with AWS Health API. AWS Health also offers out-of-the-box integration with other IT Service Management (ITSM) tools.

138.2 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/.

138.3 Pricing Overview

Please see the AWS UK G Cloud 14 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/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/

138.5 Technical Requirements

Please refer to https://docs.aws.amazon.com/health/latest/ug/getting-started-phd.html for comprehensive technical documentation regarding this service.

139 AWS HealthImaging

139.1 Service Overview

You can use AWS HealthImaging as the data management layer for your medical imaging applications while reducing your total cost of ownership (TCO) up to 40%. HealthImaging optimizes your medical imaging data (X-Ray, CT, MRI, Ultrasound, and so on) and stores a single, authoritative copy of each image that can be used to drive all of your clinical and research workflows. When creating de-identified copies of an image, pixel data is not copied, saving you storage costs.

139.1.1 Features

- **Developer-friendly metadata**: HealthImaging simplifies application development by returning DICOM metadata in a developer-friendly format. Individual attributes are accessible using human-readable keywords rather than unfamiliar group and element hexadecimal numbers. Patient and study attributes are normalized, removing the need for application developers to deal with inconsistencies between SOP instances.

- **SIMD-accelerated image decoding**: HealthImaging returns images encoded using High-Throughput JPEG 2000 (HTJ2K)—a state-of-the-art codec ideal for medical imaging. HTJ2K takes advantage of single instruction multiple data (SIMD) instructions on modern processors to deliver new levels of performance. HTJ2K decodes images an order of magnitude faster than JPEG 2000 and is at least twice as fast as all other DICOM transfer syntaxes. WASM SIMD can be utilized to bring this extreme speed to zero-footprint web viewers.
- **Industry-leading performance**: HealthImaging sets a new standard for image loading performance thanks to its efficient metadata encoding, lossless compression, and progressive resolution data access. Efficient metadata encoding enables image viewers and AI algorithms to understand the contents of a DICOM study without having to load the image data. Images load faster without any compromise in quality due to state-of-the-art image compression.
- **Scalable ingestion**: HealthImaging import processes use modern cloud-native technologies capable of ingesting many DICOM studies in parallel. Historical archives can be quickly ingested without impacting clinical workloads for new data.

139.1.2 Benefits

- **Lower total cost of ownership**: Reduce costs up to 40% by employing advanced compression and storing a single copy of each image in the cloud.
- **Fast image access**: Access medical imaging data with sub-second image latency from anywhere, powered by cloud-native APIs and applications.
- **Automated infrastructure management**: Reduce the burden of infrastructure management to focus time and resources on delivering high-quality patient care.
- **Robust partner network**: Supported by major medical imaging vendors with cloud-native medical imaging workflows.

139.2 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/.

139.3 Pricing Overview

Please see the AWS UK G Cloud 14 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/healthimaging/
- **Service quotas:** https://docs.aws.amazon.com/healthimaging/latest/devguide/endpoints-quotas.html#quotas
- **Service FAQs:** https://aws.amazon.com/healthimaging/faqs/

139.5 Technical Requirements

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

- Developer Guide: Provides a conceptual overview of AWS HealthImaging and detailed instructions for using its features.

- API Reference: Describes all the API actions for AWS HealthImaging in detail.

140 AWS Healthlake

140.1 Service Overview

AWS HealthLake (AHL) is a HIPAA-eligible service offering healthcare companies a complete view of individual and patient population health data using FHIR (Fast Healthcare Interoperable Resources) API based transactions to securely store and transform their data into a queryable format at petabyte scale, and further analyze this data using machine learning (ML) models. Using the HealthLake FHIR-based APIs, healthcare organizations can easily import large volumes of health data, including medical reports or patient notes, from on-premises systems to a secure, compliant, and pay-as-you-go service in the cloud. HealthLake offers built-in natural language processing (NLP) models to help customers understand and extract meaningful medical information from a single copy of raw health data, such as medications, procedures, and diagnoses.

140.1.1 Features

- **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. AWS HealthLake meets rigorous security and access controls to ensure patients' sensitive health data is protected and meets regulatory compliance. With customer-managed keys, GDPR compliance, and HIPAA eligibility, customers can be confident their patients' data remains secure and their privacy is protected.
- **Leverage transactional FHIR server**: Currently in preview, leverage FHIR APIs for standard resource validation, SMART on FHIR authorization, and Bulk data FHIR API export capabilities to support unifying and analyzing your data to reduce operational costs and improve decision making. AWS HealthLake supports customer conformity to the latest ONC and CMS regulatory standards including: HL7 FHIR API R4 v4.0.1, FHIR Bulk Data Access v2.0.0, US Core IG STU v3.1.1, HL7 SMART App Launch Framework IG v1.0.0, and integrate with OAuth 2.0 & OpenID Connect.
- **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. The NLP extracted entities are stored as native FHIR resources within the data store that can be accessed through FHIR APIs or through SQL from the Athena query engine.
- **Powerful query & search capabilities**: AWS HealthLake supports FHIR Create/Read/Update/Delete (CRUD) and FHIR Search operations. You can add new records by performing a Create Operation for adding new patients and their information, like medications. You can read the most recent version of that record by performing a Read Operation. You can update a previously created record by performing an Update Operation. On delete, records are soft deleted and hidden from search results.. You can also search with predefined filters to find all the information on a patient.

- **Identify trends & make predictions**: AWS HealthLake Analytics supports Athena SQL queries on HealthLake data that enables users to analyze without needing to export the data. With Amazon QuickSight, you can create dashboards on HealthLake data to quickly explore patient trends. Here is an example of a population health dashboard created using Amazon QuickSight. You can also build, train, and deploy your own predictive analytics using machine learning models with Amazon SageMaker. Here is how to build a number of predictive chronic or acute disease models using Amazon SageMaker with AWS HealthLake normalized data. Here is an example of building an ML-enabled cognitive search application where every clinical evidence is tagged, indexed, and structured to provide evidence-based topics on things like transmission, risk factors, therapeutics, and incubation. This particular functionality is tremendously valuable for clinicians or scientists, who can quickly ask questions to validate their clinical decisions or advance their research.

140.1.2 Benefits

- **Store clinical and claims data using interoperability standards**: Store patient medical history from multiple data sources in the normalized common data model (FHIR-based) format and leverage FHIR APIs to build transactional applications and patient 360 views
- **Build highly scalable interoperability solutions to meet regulatory needs**: Leverage HealthLake Patient Access APIs and Bulk FHIR APIs with built-in support for US CORE and CARIN BB profile validation to meet the 21st Century Cures Act for patient access and interoperability requirements
- **Automatic transformation of FHIR data to analytics format**: Empower healthcare analysts to query HealthLake FHIR data using SQL through AWS Athena query engine and create metric tracking care gap dashboards
- **Apply analytics and ML on health data**: Make predictions and identify trends with health data using Amazon SageMaker machine learning models and Amazon QuickSight analytics

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

140.5 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.

141 AWS HealthOmics

141.1 Service Overview

AWS HealthOmics is an AWS service that helps users such as bioinformaticians, researchers, and scientists to store, query, analyze, and generate insights from genomics and other biological data. It simplifies and accelerates the process of storing and analyzing genomic information for research and clinical organizations, and makes scientific discovery and insight generation faster.

HealthOmics has three primary components. HealthOmics Storage helps you store and share petabytes of genomics data efficiently and at low cost per gigabase. HealthOmics Analytics simplifies how you prepare genomics data for multiomics and multimodal analyses. HealthOmics Workflows automatically provisions and scales the underlying infrastructure for your bioinformatics computation.

141.1.1 Features

- **HealthOmics Storage**: Helps you store and share petabytes of raw genomics data efficiently and at low cost per gigabase.
- **HealthOmics Analytics**: Simplifies how you prepare genomics data for multiomics and multimodal analyses.
- **HealthOmics Workflows**: Automatically provisions and scales the underlying infrastructure for your bioinformatics workflows.

141.1.2 Benefits

- **Securely store and combine genomic data**: HealthOmics integrates with other AWS services such as AWS Lake Formation and Amazon Athena. You can securely store your genomics data and then query or combine it with medical history data for better diagnoses and personalized treatment plans.
- **Protect patient privacy**: HealthOmics is HIPAA eligible. It also integrates with IAM and Amazon CloudWatch so that you can control and log data access, as well as how the data has been used in analyses.
- **Built to scale**: Support large population data analyses with simplified billing and new collaboration tools.
- **Maximize efficiency**: Use automated workflows and integrated tools to streamline data processing and analysis.

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 14 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/omics/
- **Service quotas:**
  - HealthOmics Storage quotas: https://docs.aws.amazon.com/omics/latest/dev/storage-quotas.html
  - HealthOmics Analytics quotas: https://docs.aws.amazon.com/omics/latest/dev/analytics-quotas.html
  - HealthOmics Workflows quotas: https://docs.aws.amazon.com/omics/latest/dev/workflows-quotas.html
- **Service FAQs:** https://aws.amazon.com/healthomics/faqs/

141.5 Technical Requirements

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

- AWS HealthOmics User Guide: Provides a conceptual overview of AWS HealthOmics and detailed instructions for using the various features.
- API Reference: Describes all the API operations for AWS HealthOmics in detail. Also provides details of request and response syntax and errors for the supported web services protocols.

142 AWS Identity and Access Management (IAM)

142.1 Service Overview

AWS Identity and Access Management (IAM) is a web service that helps you securely control access to AWS resources. With IAM, you can centrally manage permissions that control which AWS resources users can access. You use IAM to control who is authenticated (signed in) and authorized (has permissions) to use resources.

142.1.1 Features

- **Shared access to your AWS account**: You can grant other people permission to administer and use resources in your AWS account without having to share your password or access key.
- **Granular permissions**: You can grant different permissions to different people for different resources. For example, you might allow some users complete access to Amazon Elastic Compute Cloud (Amazon EC2), Amazon Simple Storage Service (Amazon S3), Amazon DynamoDB, Amazon Redshift, and other AWS services. For other users, you can allow read-only access to just some S3 buckets, or permission to administer just some EC2 instances, or to access your billing information but nothing else.
- **Secure access to AWS resources for applications that run on Amazon EC2**: You can use IAM features to securely provide credentials for applications that run on EC2 instances. These credentials provide permissions for your application to access other AWS resources. Examples include S3 buckets and DynamoDB tables.

- **Multi-factor authentication (MFA)**: You can add two-factor authentication to your account and to individual users for extra security. With MFA you or your users must provide not only a password or access key to work with your account, but also a code from a specially configured device. If you already use a FIDO security key with other services, and it has an AWS supported configuration, you can use WebAuthn for MFA security.
- **Identity federation**: You can allow users who already have passwords elsewhere—for example, in your corporate network or with an internet identity provider—to get temporary access to your AWS account.
- **Identity information for assurance**: If you use AWS CloudTrail, you receive log records that include information about those who made requests for resources in your account. That information is based on IAM identities.
- **PCI DSS Compliance**: IAM supports the processing, storage, and transmission of credit card data by a merchant or service provider, and has been validated as being compliant with Payment Card Industry (PCI) Data Security Standard (DSS).
- **Integrated with many AWS services**: For a list of AWS services that work with IAM, see AWS services that work with IAM.
- **Eventually consistent**: IAM, like many other AWS services, is eventually consistent. IAM achieves high availability by replicating data across multiple servers within AWS data centers around the world. If a request to change some data is successful, the change is committed and safely stored. However, the change must be replicated across IAM, which can take some time. Such changes include creating or updating users, groups, roles, or policies. We recommend that you do not include such IAM changes in the critical, high-availability code paths of your application. Instead, make IAM changes in a separate initialization or setup routine that you run less frequently. Also, be sure to verify that the changes have been propagated before production workflows depend on them.
- **Free to use**: AWS Identity and Access Management (IAM) and AWS Security Token Service (AWS STS) are features of your AWS account offered at no additional charge. You are charged only when you access other AWS services using your IAM users or AWS STS temporary security credentials.

142.1.2 Benefits

- **Set permission guardrails and fine-grained access**: Set and manage guardrails with broad permissions, and move toward least privilege by using fine-grained access controls for your workloads.
- **Manage workload and workforce identities across your AWS accounts**: Manage identities across single AWS accounts or centrally connect identities to multiple AWS accounts.
- **Use temporary security credentials and permission sets to access your AWS resources**: Grant temporary security credentials for workloads that access your AWS resources using IAM and grant your workforce access with AWS IAM Identity Center.
- **Analyze access and validate IAM policies as you move toward least privilege**: Generate least-privilege policies, verify external and unused access to resources, and continually analyze the required permissions.

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 14 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/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/

142.5 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.
- IAM API Reference: Describes all the API operations for AWS Identity and Access Management in detail. Also provides sample requests, responses, and errors for the supported web services protocols.
- IAM in the AWS CLI Reference: Describes the AWS CLI commands that you can use to administer IAM. Provides syntax, options, and usage examples for each command.
- 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.

143 AWS IoT Analytics

143.1 Service Overview

AWS IoT Analytics is a fully-managed IoT analytics service that collects, pre-processes, enriches, stores, and analyzes IoT device data at scale. IoT Analytics can perform simple ad hoc queries as well as complex analysis, and is a simpler way to run IoT analytics for understanding the performance of devices, predicting device failures, and machine learning. It is designed specifically for IoT and automatically captures and stores the message timestamp making time-series analysis easy.

IoT Analytics can also enrich the data with device-specific metadata such as device type and location using the AWS IoT registry. AWS IoT Analytics stores data in an IoT-optimized data store so you can run queries on large datasets.

With AWS IoT Analytics, you can also execute custom authored code containers. AWS IoT Analytics automates the execution of your custom analyses created in Jupyter Notebook or your own tools (such as Matlab, Octave, etc.) for scheduled execution.

143.1.1 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 analyze**: 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.
- **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 analyze 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. These templates contain AWS-authored machine learning models and visualizations to help with use cases related to device failure profiling and forecasting events such as low usage.
- **Bring Your Custom Container**: AWS IoT Analytics will import your custom authored code containers, built in AWS IoT Analytics or a third party, such as Matlab, or Octave, giving you more time to focus on what sets you apart from your competition.
- **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.
- **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.

Use or disclosure of data contained on this sheet is subject to the restriction on the title page of this document.

- **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.

143.1.2 Benefits

- **Predictive maintenance**: AWS IoT Analytics provides templates to build predictive maintenance models and apply them to your devices. For example, you can use AWS IoT Analytics to predict when heating and cooling systems are likely to fail on connected cargo vehicles so the vehicles can be rerouted to prevent shipment damage.
- **Proactive replenishing of supplies**: AWS IoT Analytics lets you build IoT applications that can monitor inventories in real time.
- **Process efficiency scoring**: With AWS IoT Analytics, you can build IoT applications that constantly monitor the efficiency of different processes and take action to improve the process.
- **Smart agriculture**: AWS IoT Analytics can enrich IoT device data with contextual metadata using AWS IoT registry data or public data sources so that your analysis factors in time, location, temperature, altitude, and other environmental conditions. With that analysis, you can write models that output recommended actions for your devices to take in the field. For example, to determine when to water, irrigation systems might enrich humidity sensor data with data on rainfall, enabling more efficient water usage.

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

Please see the AWS UK G Cloud 14 Pricing Document accompanying this service in the Digital Marketplace.

143.4 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/

143.5 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.
- API Reference: Describes all the API operations for AWS IoT Analytics in detail. Also provides sample requests, responses, and errors for the supported web services protocols.

- AWS IoT Analytics in the AWS CLI Reference: Describes the AWS CLI commands that you can use to work with AWS IoT Analytics.

144 AWS IoT Core

144.1 Service Overview

AWS IoT Core is a managed cloud platform that lets connected devices easily and securely interact with cloud applications and other devices. AWS IoT Core can support billions of devices and trillions of messages, and can process and route those messages to AWS endpoints and to other devices reliably and securely. With AWS IoT Core, your applications can keep track of and communicate with all your devices, all the time, even when they aren't connected.

144.1.1 Features

- **AWS IoT Device SDK**: The AWS IoT Device SDK lets you connect your hardware device or mobile application to AWS IoT Core. The AWS IoT Device SDK helps your devices 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 help developers 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 verify that devices can securely and efficiently communicate with AWS IoT Core. Currently the Device Gateway supports the MQTT, WebSockets, and HTTPS protocols. For devices that connect using MQTT or WebSockets the Device Gateway will maintain long lived, bidirectional connections, helping these devices send and receive messages at any time with low latency. The Device Gateway is fully managed and scales automatically to support over a billion devices without requiring you to manage any infrastructure.
- **Message Broker**: The Message Broker is a high throughput Pub/Sub messaging agent that securely transmits messages to and from all of your IoT devices and applications with low latency based on the MQTT Version 5.0 Message Standard. The flexible nature of the Message Broker's topic structure helps you send messages to, or receive messages from, a number of devices. It supports messaging patterns ranging from one-to-one command and control messaging, to one-to-one million (or more) broadcast notification systems and everything in between.
- **CoAP Protocol**: Constrained Application Protocol (CoAP) is a specialized Internet application layer protocol that is purpose built for constrained devices like battery powered IoT sensors. CoAP offers familiar HTTP-style request/response messaging in a lightweight User Datagram Protocol (UDP) protocol, making it a suitable protocol choice for IoT use-cases where devices communicate infrequently to the cloud and spend most of their time offline in a low power state. AWS IoT Core offers CoAP/UDP support for cellular devices, such as those using Narrowband IoT (NB-IoT) technology, through partner-developed IoT platforms built on AWS.

- **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 an identity to each device that is consistently formatted regardless of the type of device or how it connects.
- **Device Shadow**: Create a persistent, virtual version, or Device Shadow, of each device. This 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 use the rules engine.
- **Rules Engine**: The Rules Engine helps you build IoT applications that gather, process, analyze, and act on data generated by connected devices at a global scale without having to manage any infrastructure. The Rules Engine evaluates inbound messages published on 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.
- **AWS IoT Core for LoRaWAN**: AWS IoT Core for LoRaWAN helps customers connect wireless devices that use low-power, long-range wide area network (LoRaWAN) technology. Using AWS IoT Core, you can now set up a private LoRaWAN network by connecting your own LoRaWAN devices and gateways to AWS without developing or operating a LoRaWAN Network Server (LNS).
- **Device Location**: AWS IoT Core Device Location helps you track and manage your fleet of IoT devices using their location data, such as latitude and longitude coordinates, without traditional Global Positioning Service (GPS) hardware. When you use the Device Location feature, you don't have to rely on high-power consuming GPS hardware, and can choose an appropriate location technology that works within your engineering constraints.
- **AWS IoT Core for Amazon Sidewalk**: 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. It can also 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.

144.1.2 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.

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 14 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/iot/
- **Service quotas:** https://docs.aws.amazon.com/general/latest/gr/iot-core.html
- **Service FAQs:** https://aws.amazon.com/iot-core/faqs/

144.5 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.

145 AWS IoT Device Defender

145.1 Service Overview

AWS IoT Device Defender is a fully managed IoT security service that enables you to secure your IoT configurations on an ongoing basis. With AWS IoT Device Defender, you get tools to identify and respond to security issues. AWS IoT Device Defender audits your fleet to ensure it adheres to security best practices, continuously monitors your device fleets to detect any abnormal device behavior, alerts you about security issues as they arise, and provides built-in mitigation actions for these security issues.

145.1.1 Features

- **Audit**: AWS IoT Device Defender audits your device-related resources against AWS IoT security best practices. 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 behaviors that may be indicative of a compromise by continuously monitoring high-value security metrics from the device and AWS IoT Core. You can specify normal device behavior for a group of devices by setting up behaviors (rules) for these metrics. AWS IoT Device Defender monitors and evaluates each datapoint reported for these metrics against user-defined behavior (rules) and alerts you if an anomaly is detected.
- **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.

- ML Detect: AWS IoT Device Defender monitors and identifies anomalous datapoints for six cloud-side metrics with ML models and triggers an alarm if an anomaly is detected. AWS IoT Device Defender removes the need to define accurate behaviors 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 to refresh the expected device behaviors based on the latest trailing 14 days. ML Detect makes getting started with monitoring easy.
- **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.
- **Metrics Integration**: With AWS IoT Device Defender ListMetricValues API, you can visualize device-side, cloud-side and custom metrics from connected devices through an open API and easily integrate these metrics into any of your custom dashboards to get a complete overview of your deployments. Visualization of the data is also visible and integrated into Fleet Hub for AWS IoT Device Management.

145.1.2 Benefits

- **Identify Vulnerabilities**: Audit the security posture of IoT resources across your device fleet to easily identify gaps and vulnerabilities.
- **Monitor Traffic**: Use ML models or define your own device behaviors to monitor traffic from a malicious IP or a spike in connection attempts.
- **Receive Alerts and Take Action**: Get security alerts when an audit fails or behavior anomalies are detected. Quickly take actions to minimize operational risk.
- **Predict Security Issues**: Easily mitigate security issues through built-in actions, such as updating a device certificate, quarantining a group of devices, or replacing default policies.

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 14 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-device-defender/
- **Service quotas:** https://docs.aws.amazon.com/general/latest/gr/iot_device_defender.html
- **Service FAQs:** https://aws.amazon.com/iot-device-defender/faq/

145.5 Technical Requirements

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

- Developer Guide: Walks through how to set up the service and integrate with other AWS services.

- API Reference: Describes all of the API operations for service. Also provides sample requests, responses, and errors for the supported web services protocols..
- AWS IoT sections of the AWS CLI Reference: Describes the AWS CLI commands that you can use to administer and manipulate AWS IoT Device Defender.

146 AWS IoT Device Management

146.1 Service Overview

AWS IoT Device Management is a service that makes it easy to securely register, organize, monitor, and remotely manage IoT devices at scale throughout their lifecycle. You can use IoT Device Management to upload and view device information and configuration, organize your device inventory, monitor your fleet of devices, troubleshoot individual devices, and remotely manage devices deployed across many locations including updating device software over-the-air (OTA). With IoT Device Management, you can scale your device fleets and reduce the cost and effort of managing large IoT device deployments.

146.1.1 Features

- **Register connected devices in bulk**: Register new devices with AWS IoT Device Management. Using the IoT management console or API, you can 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 in a few steps in the management console.
- **Organize connected devices into groups**: 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.
- **Index and search your fleet**: 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.
- **Monitor system performance**: Monitoring is an important part of maintaining the reliability, availability, and performance of your IoT solutions. With AWS IoT Device Management, 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 that 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.
- **Manage and perform updates remotely**: Push software and firmware to devices in the field to patch security vulnerabilities and improve device functionality. You can initiate bulk updates, configure rollout schedule, control deployment velocity, set failure thresholds, and define continuous jobs to update device software automatically so that they are always running the latest version.

- **Establish bidirectional communication to remote devices with 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 steps. You can also make multiple concurrent client connections over a single secure tunnel, so you can perform more advanced device troubleshooting, such as issuing remote shell commands to a device while simultaneously debugging a web application on the same device.
- **Manage software package versions**: With AWS IoT Device Management Software Package Catalog, easily track and monitor software package versions, gain valuable insights from a central dashboard, and perform targeted updates on IoT devices running specific software version.
- **Visualize device performance with Fleet Hub**: AWS IoT Device Management includes the ability to create no-code, fully managed web applications to visualize and interact with your device fleets connected to AWS IoT. With Fleet Hub for AWS IoT Device Management, you can easily build these standalone web applications and make them available to users in your organization, even if they don't have AWS accounts.

146.1.2 Benefits

- **Fleet maintenance**: Swiftly onboard and organize your devices into flexible hierarchies to streamline fleet maintenance and workflow updates.
- **Time saving**: Save time by filtering your device search based on specific attributes to make informed decisions.
- **Secure monitoring**: Securely and remotely monitor your device fleet health status, analyze trends, troubleshoot, and push updates at scale.
- **Device health**: Visualize your fleet's health status and remotely perform real-time actions, such as firmware updates and device reboots, using Fleet Hub.

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 14 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-device-management/
- **Service quotas:** https://docs.aws.amazon.com/general/latest/gr/iot_device_management.html#iot_device_management_quotas
- **Service FAQs:** https://aws.amazon.com/iot-device-management/faq/

146.5 Technical Requirements

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

- 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.
- API Reference: Describes all the API operations for AWS IoT Device Management in detail. Provides examples of request and response syntax.
- 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.

147 AWS IoT Events

147.1 Service Overview

AWS IoT Events is a new IoT service that helps companies continuously monitor their equipment and fleets of devices for failure or changes in operation and trigger alerts to respond when events occur. AWS IoT Events recognizes events across multiple sensors to identify operational issues, such as equipment slowdowns, and generates alerts such as notifying support teams of an issue. AWS IoT Events offers a managed complex event detection service on the AWS cloud, accessible through the AWS IoT Events console, a browser-based GUI where you can define and manage your event detectors, or direct ingest APIs, code that allows two applications to communicate with each other. Understanding equipment or a process based on telemetry from a single sensor is often not possible; a complex event detection service will combine multiple sources of telemetry to gain full insight into equipment and processes. You define conditional logic and states inside AWS IoT Events to evaluate incoming telemetry data to detect events in equipment or a process. When AWS IoT Events detects an event, it can trigger pre-defined actions in another AWS service, such as sending alerts through Amazon Simple Notification Service (SNS).

147.1.1 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. If you have many devices using the same detector definition, you can easily launch a new instance of the detector for every unique value of the signal. You can automatically extend your event detection and action capabilities whenever a new device comes online allowing you take advantage of stateful complex event detection at scale. Events can be detected with simple' if-then-else' statements applied to incoming sensor data streams.
- **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). You can select a pre-built action to trigger [on the basis of a detected event], such as sending a notification through Amazon SNS, a fully managed pub/sub messaging

service, to a service technician, or you can build a custom action by setting up a function for AWS Lambda, a compute service that lets you run code without provisioning or managing servers, to send the event data to a third-party application.

- **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 behavior 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, select the severity for this alarm definitio, and configure the notifications to send when an alarm is triggered. Once an alarm has been defined, operators can manage the alarm workflow by taking actions to acknowledge, snooze or disable the alarm. You can also configure additional actions to other AWS services including AWS Lambda, Amazon Simple Queue Service (SQS), and Amazon SNS, to be executed when an alarm triggers. This allows you to integrate alarm notifications with your own ticketing or notification system.

147.1.2 Benefits

- **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.
- **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.

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 14 Pricing Document accompanying this service in the Digital Marketplace.

147.4Service Constraints

Information on service features and constraints is available online here:

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

147.5Technical Requirements

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

- Developer Guide: Walks through how to set up the service and integrate with other AWS services.
- API Reference: Describes all of the API operations for service. Also provides sample requests, responses, and errors for the supported web services protocols..
- AWS IoT sections of the AWS CLI Reference: Describes the AWS CLI commands that you can use to administer and manipulate AWS IoT Device Defender.

148 AWS IoT FleetWise

148.1Service Overview

AWS IoT FleetWise is a managed service that allows you to efficiently collect vehicle data and organize it in the cloud to gain insights about your fleet(s) of vehicles. You can use the data transferred by AWS IoT FleetWise to help you analyze vehicle fleet health to quickly identify potential maintenance issues, make in-vehicle infotainment systems smarter, or use analytics and machine learning (ML) to improve models for autonomous driving and advanced driver assistance systems (ADAS).

148.1.1 Features

- **Vision System Data (Preview)**: You can use AWS IoT FleetWise to easily collect and organize data from vehicle vision systems that include cameras, radars, and lidars. It keeps both structured (object list, sensor metadata, vehicle ID) and unstructured (images, video) vision system data, metadata (event ID, campaign, vehicle), and standard sensor (telemetry data) automatically synchronized in the cloud, making it easier for you to assemble a full picture view of events and gain insights. When collecting vision system data using AWS IoT FleetWise, you can take advantage of the features and interfaces you already use to collect telemetry data. Get started on AWS by defining and modeling a vehicle's vision sensors, alongside its attributes and standard sensors. Your Edge Agent for AWS IoT FleetWise deployed in the vehicle collects data from standard CAN-based vehicle sensors (engine temperature, fuel pressure), as well as from vehicle sub-systems that include vision sensors like cameras, radars, and lidars. You can use the same event- or time-based data collection campaign to collect data signals from both standard sensors and vision systems. In the cloud, you see a unified view of your defined vehicle attributes and other metadata, telemetry data, and structured vision system data, with links to view unstructured vision system data in Amazon S3. Timestamps and campaign and event identifiers keep data automatically synchronized across different formats, simplifying its use in downstream applications and analytics.
- **Vehicle Modeling:** Use AWS IoT FleetWise to build virtual representations of vehicles in the cloud and apply a common data format to structure and label vehicle

attributes, sensors, and signals. AWS IoT FleetWise standardizes vehicle data modeling using the Vehicle Signal Specification (VSS), so that a signal like speed is always represented as “vehicle_speed” and measured in kilometers per hour (km/h). Once the vehicle is modeled, you upload a vehicle network file, such as a CAN database (DBC) file, and AWS IoT FleetWise reads the unique and proprietary data signals sent over the vehicle’s network.

- **Edge Agent:** Create an Edge Agent for AWS IoT FleetWise to facilitate communication between the vehicle and the cloud. You can develop your Edge Agent by using the Edge Agent Reference Implementation for AWS IoT FleetWise, a fully functional piece of embedded software designed for vehicle data collection that can run on most embedded Linux-based platforms. Install your Edge Agent on the vehicle gateway, the telematics control unit (TCU), or other data collection hubs in the vehicle. Visit the Partners page to see a list of reference hardware, which is expected to continue expanding as we validate new reference hardware with the Edge Agent. It’s also straight forward to port the Edge Agent Reference Implementation to your target Linux-based hardware of choice. Find porting information in the Developer Guide. Your Edge Agent should be able to continually receive data collection schemes from AWS IoT FleetWise and collects data accordingly for transfer to the cloud. With your Edge Agent, you control every step of the process from creation to installation, and you maintain full data ownership and control of proprietary information.
- **Cloud-Based Configuration Deployment**: With AWS IoT FleetWise, you can deploy cloud-configured data collection schemes to vehicles and easily update the configuration in the cloud at any point in time as business needs change.
- **Global Signal Catalog**: Use a centralized repository of standardized vehicle sensors and signals for your vehicle models. You can save time by reusing signals when creating new vehicle models.
- **Rules-Based Data Collection**: AWS IoT FleetWise applies the rules you define for transferring only the data signals you want to the cloud. First you define vehicle attributes. Then you define rules and events for when to transfer that data, based on parameters such as changes to temperature or speed for a certain vehicle type. Rules-based data collection reduces the volume of unnecessary data transferred to the cloud, thus providing more useful data.

148.1.2 Benefits

- **Intelligent Data Collection**: Improve data relevance with intelligent data collection that sends only the data you need to the cloud for analysis.
- **Standardized Data**: Analyze standardized data from a fleet of vehicles without developing a custom data collection or logging system.
- **Automatic Data Synchronization**: Gain a unified view of data collected from both standard sensors and vision systems and keep it automatically synchronized in the cloud.

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 14 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/iot-fleetwise/
- **Service quotas:** https://docs.aws.amazon.com/general/latest/gr/iotfleetwise.html
- **Service FAQs:** https://aws.amazon.com/iot-fleetwise/faqs/

148.5 Technical Requirements

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

- Developer Guide: Describes key concepts of AWS IoT FleetWise and provides instructions for using the features of AWS IoT FleetWise.
- API Reference: Describes all the API operations for AWS IoT FleetWise in detail. Also provides sample requests, responses, and errors for the supported web services protocols.
- Edge Agent GitHub: The Edge Agent GitHub page includes source code access, platform support, build dependency information, and other resources.

149 AWS IoT Greengrass

149.1 Service Overview

AWS IoT Greengrass is an IoT open source edge runtime and cloud service that helps you build, deploy, and manage device software. Customers use AWS IoT Greengrass for their IoT applications on millions of devices in homes, factories, vehicles, and businesses. You can program your devices to act locally on the data they generate, execute predictions based on machine learning models, filter and aggregate device data, and only transmit necessary information to the cloud.

AWS IoT Greengrass lets you quickly and easily build intelligent device software. AWS IoT Greengrass enables local processing, messaging, data management, ML inference, and offers prebuilt components to accelerate application development. AWS IoT Greengrass also provides a secure way to seamlessly connect your edge devices to any AWS service as well as to third-party services.

Once software development is complete, AWS IoT Greengrass enables you to remotely manage and operate software on your devices in the field without needing a firmware update. AWS IoT Greengrass helps keep your devices updated and makes them smarter over time.

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 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 14 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/

149.5Technical 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

150.1Service Overview

AWS IoT SiteWise is a managed service that enables industrial enterprises to collect, store, organize, and visualize thousands of sensor data streams across multiple industrial facilities. AWS IoT SiteWise includes AWS IoT SiteWise Edge software that runs on a gateway device that sits onsite in a facility, continuously collects the data from a historian or a specialized industrial server, and sends it to the cloud. Data can also be ingested into AWS IoT SiteWise through AWS IoT Core via the MQTT protocol or by using a REST API. With AWS IoT SiteWise, you can skip months of developing undifferentiated data collection and cataloging solutions. Instead, you can focus on using your data to detect and fix equipment issues, spot inefficiencies, and improve production output.

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 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 14 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 TwinMaker

151.1 Service Overview

AWS IoT TwinMaker is a service that makes it faster and easier to create digital twins of real-world systems and apply them to improve operations. Creating digital twins is a complex process requiring developer expertise that spans operational technology (OT) systems, graph databases, IoT technologies, user interfaces, and 3D web development. To build digital twins, you have to manually connect different types of data from diverse sources (for example, time-series sensor data from equipment, video feeds from cameras, and maintenance records from business applications) and model them into a graph that represents the relationships between the data sources in your physical environment. You then have to build a 3D virtual representation of your physical systems (for example, buildings, factories, equipment, and production lines) and overlay the real-world data onto the 3D visualization. After that, developers have to create applications for end users such as plant operators and maintenance engineers to use their digital twins for everyday operations. AWS IoT TwinMaker makes this process easier by: providing a service that connects to different types of data from multiple sources; automatically creating and managing a digital twin graph that combines and understands the relationships among the connected data sources; and providing a simple web-based tool to create 3D visualizations that display data and insights to create the digital twin.

151.1.1 Features

- **Data Connectors**: AWS IoT TwinMaker provides built-in data connectors for AWS IoT SiteWise for collecting, organizing, and storing equipment and time-series sensor data. AWS IoT TwinMaker also provides a framework for you to easily create custom data connectors to use with other AWS or third-party data sources, such as Amazon Timestream, Snowflake, and Siemens Mindsphere. These data connectors allow your applications to only use the AWS IoT TwinMaker unified data access API to read from and write to the different data stores without needing to query each data source using their own individual API.
- **Model Builder**: To model your physical environment, you can create entities in AWS IoT TwinMaker that are virtual representations of your physical systems, such as a furnace or an assembly line. You can also specify custom relationships between these entities to accurately represent the real-world deployment of these systems. You then connect these entities to your various data stores to form a digital twin graph, which is a knowledge graph that structures and organizes information about the digital twin for easier access and understanding. As you build out this model of your physical environment, AWS IoT TwinMaker automatically creates and updates the digital twin graph by organizing the relationship information in a graph database.
- **Scene Composer**: With AWS IoT TwinMaker, you build a 3D digital twin by using your existing and previously built 3D visual models, such as CAD files, Building Information Modeling (BIM) files, or point cloud scans. Using the AWS IoT TwinMaker scene composer and simple 3D tools, you import these visual assets into a scene and position them to match your physical environment—for example a factory and its equipment. You can then add interactive video and sensor data overlays from the connected data sources, insights from connected machine learning (ML) and simulation services, and maintenance records and operational documents to provide you with a regularly updated, spatially aware visualization of your operations.
- **Applications**: Once you've created the digital twin, AWS IoT TwinMaker provides a low-code experience for building a web application so your plant operators and maintenance engineers can access and interact with the digital twin. AWS IoT

TwinMaker comes with a plug-in for Grafana, a popular open-source dashboard and visualization platform from Grafana Labs. The plug-in provides custom visualization panels, including a 3D scene viewer and dashboard templates, as well as a data-source component to connect to your digital twin data, allowing you to quickly create 3D-enabled applications for your specific needs. The plug-in can also be used to build applications with Amazon Managed Grafana, which is a fully managed service for open-source Grafana.

151.1.2 Benefits

- **Use Your Existing Tools**: Use your existing IoT, video, and enterprise application data where it already lives—without needing to reingest or move the data to another location.
- **Combine Existing 3D Models With Real-World Data**: Save time with an automatically generated knowledge graph that binds your data sources to virtual replicas of physical systems to accurately model real-world environments.
- **Create a Holistic View**: Get an immersive 3D view of your systems and operations to optimize efficiency, increase production, and improve performance

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 14 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://aws.amazon.com/iot-twinmaker/resources/
- **Service quotas:** https://docs.aws.amazon.com/general/latest/gr/iot-twinmaker.html
- **Service FAQs:** https://aws.amazon.com/iot-twinmaker/faqs/

151.5 Technical Requirements

Please refer to https://aws.amazon.com/iot-twinmaker/resources/ and the following links for comprehensive technical documentation regarding this service.

- Developer Guide
- GitHub Repository
- API Reference

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 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 14 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/

152.5Technical 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

153.1Service Overview

AWS Lake Formation makes it easier to centrally govern, secure, and globally share data for analytics and ML. With Lake Formation, you can centralize data security and governance using the AWS Glue Data Catalog, letting you manage metadata and data permissions in one place with familiar database-style features. It also delivers fine-grained data access control, so you can help ensure users have access to the right data, down to the row and column level. You can then scale permissions across your users. Lake Formation also makes it easier to share data internally across your organization, across Regions, and externally using AWS Data Exchange, letting you create a data mesh or meet other data sharing needs with no data movement. And, because Lake Formation tracks data interactions by role and user, it provides comprehensive data access auditing to help ensure the right data was accessed by the right users at the right time.

153.1.1 Features

- **Centralize Permissions Management**: Lake Formation centralizes permission management on your data resources in the AWS Glue Data Catalog, including databases and tables. You can define and manage access by role for your users and applications using familiar database-like grants, bringing the simplicity of data warehouses and databases to your data lake.
- **Define and Manage Fine-Grained 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 with fine-grained access control (FGAC). These policies apply to AWS IAM users and roles and to users and groups when federating through an external identity provider. You can use FGAC to access data secured by Lake Formation within Amazon Redshift Spectrum, Amazon Athena, AWS Glue ETL, and Amazon EMR for Apache Spark.
- **Enforce Permissions With AWS Analytics Services Integration**: AWS Lake Formation helps you consistently enforce permissions across AWS analytics services with native integrations for Amazon Athena, Amazon SageMaker, Amazon Redshift, as well as AWS Glue for data integration and Amazon EMR for big data processing. Integration with AWS Identity and IAM authenticates users and roles, enforcing permissions across AWS analytics and ML services.
- **Integrates With Third-Party Partners**: Lake Formation is integrated with third-party partners so you can extend your permissions management to the engines you prefer, such as Starburst and Dremio. Lake Formation also integrates with Privacera and

Collibra so you can pull permissions or push permissions with Lake Formation and exploit the reach of permissions management capabilities in both Privacera and Collibra. See the documentation for more information on Lake Formation partnerships.

- **Centrally Manage Permissions Across Your Users**: Lake Formation makes it easier to scale permissions across users with tag-based access controls. With tag-based access controls, you can set attributes on data and apply permissions to those attributes to scale. Lake Formation tag-based access control (LF-TBAC) dynamically uses data attributes in the tags to scale permissions as data changes.
- **Scale Dynamic Permissions With AI-Driven Tag Management**: Lake Formation tags can be quickly populated with your own business rules and ontologies such as departments, product lines, data ownership, data sensitivity (for example, public or private), and data classification (for example, Social Security Number, phone numbers). You can dynamically manage your tag values by using integrated AWS services, including AWS Glue Sensitive Data Detection. AWS Glue Sensitive Data Detection can identify a variety of personally identifiable information (PII) and other sensitive data like credit card numbers, helping you tag for data audit purposes or sensitive information.
- **Manage Permissions for Your Data From a Centralized Catalog**: Lake Formation lets you build permissions on databases and tables within the AWS Glue Data Catalog. This allows you to use the AWS Glue Data Catalog as a hub for managing and sharing your data. With AWS Glue Data Catalog federation features, you can extend permissions to data cataloged by your own Hive metastore or with Amazon Redshift data sharing. You can set up and enforce permissions on datasets presented through the AWS Glue Data Catalog, making it easier to control access to your data no matter where it lives.
- **Allow Secure Data Sharing Across Your Organization**: AWS Lake Formation allows for data sharing with zero ETL, making it easier to maintain control of your data while still ensuring users have access. Lake Formation simplifies data sharing, letting you create a data mesh or meet other data-sharing needs. Lake Formation cross-account and cross-Region data-linking capabilities allow users to securely share distributed data lakes across multiple AWS accounts, AWS Organizations, and AWS Regions. Lastly, with Lake Formation data sharing, you can directly control who you are sharing data with, such as selecting the exact IAM principals in other accounts to help ensure data ownership is controlled by the owner once it is shared.
- **Simplify Business-To-Business Data Sharing**: AWS Lake Formation allows business-to-business data sharing external to your organization for licensing or other uses. Lake Formation integrates with AWS Data Exchange — an AWS service that lets you find, subscribe to, and use third-party data in the cloud — so you can share data with external businesses without moving or copying the data.
- **Discover Relevant Data for Analysis**: With Lake Formation permissions on the AWS Glue Data Catalog, users enjoy online, text-based search capabilities to provide them a better understanding of data within the AWS Glue Data Catalog. You can search for relevant data by name, content, sensitivity, or any other defined custom labels.
- **Comprehensive Data Access and Audit Logging**: Lake Formation provides comprehensive audit logs with Amazon CloudTrail to monitor access and compliance with centrally defined policies. You can audit data access history across analytics and ML services that read the data using Lake Formation. This lets you see which users or roles have attempted to access what data and when. You can access audit logs in the same way you access other CloudTrail logs using the CloudTrail APIs and console.

153.1.2 Benefits

- **Manage Permissions**: Manage fine-grained data lake access permissions using familiar database-like features.
- **Scaled Management**: Simplify security management and governance for your users at scale.
- **Deeper Insights**: Quickly gain deeper insights from data securely shared with internal and external users.
- **Data Auditing**: Monitor your data access and help achieve compliance with comprehensive data auditing.

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 14 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. Also, it includes detailed instructions for using the various features.
- API Reference: Describes the API operations for AWS Lake Formation service, including sample requests and responses.
- CLI Guide: Describes the CLI commands for AWS Lake Formation.

154 AWS Lambda

154.1 Service Overview

AWS Lambda lets you run code without provisioning or managing servers. You pay only for the compute time you consume - there is no charge when your code is not running. With Lambda, you can run code for virtually any type of application or backend service - all with zero administration. Just upload your code, and Lambda takes care of everything required to run and scale your code with high availability. You can set up your code to automatically trigger from other AWS services or call it directly from any web or mobile app.

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 appliation 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 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 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 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.

- **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 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 SDK integrates with AWS IAM to ensure secure code access to other AWS services. AWS Lambda runs your code within an Amazon 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 is a service that makes it easier for you to manage your software licenses from software vendors (for example, Microsoft, SAP, Oracle, and IBM) centrally across AWS and your on-premises environments. This provides control and visibility into the

usage of your licenses, enabling you to limit licensing overages and reduce the risk of non-compliance and misreporting.

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 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 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 or disclosure of data contained on this sheet is subject to the restriction on the title page of this document.

- **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 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 Mainframe Modernization

156.1 Service Overview

AWS Mainframe Modernization is an elastic mainframe service and set of development tools for migrating and modernizing mainframe and legacy workloads. Using Mainframe Modernization, system integrators can help discover your mainframe and legacy workloads, assess and analyze migration readiness, and plan migration and modernization projects. Once planning is complete, you can use the Mainframe Modernization built-in development tools to replatform or refactor your mainframe and legacy workloads, test workload performance and functionality, and migrate your data to AWS.

The Mainframe Modernization managed runtime environment simplifies managing the production environment for mainframe and legacy workloads. Mainframe Modernization supports two transformation patterns: replatforming and automated refactoring. Select your best-fit migration pattern based on business objectives.

156.1.1 Features

- **Assess**: AWS Mainframe Modernization service offers application intelligence, knowledge, and analysis for migration teams and developers to better understand large application portfolios. The assessment capability can help you assess, scope, and plan a migration and modernization project. It highlights application dependencies, complexities and facilitates the definition of the modernization strategy, including business and technical evaluations. It highlights application dependencies and complexities, and it facilitates the definition of the modernization strategy, including business and technical evaluations.
- **Refactor**: AWS Mainframe Modernization service offers both automated and manual refactoring capabilities to accelerate the modernization of mainframe and legacy assets. Powered by the AWS Mainframe Modernization Blu Age refactor solution, you can use refactoring to convert legacy application programming languages, create macroservices or microservices, and modernize user interfaces (UIs) and application software stacks. It can convert languages like COBOL, PL/1, NATURAL, RPG/400, and COBOL/400 into agile Java services and web frameworks. It relies on fully automated refactoring, preserving the investment in business functions while

expediting the reliable transition to newer languages, data stores, and frameworks. Learn more about automated refactoring powered by AWS Blu Age automated refactoring.

- **Replatform with Micro Focus**: With AWS Mainframe Modernization service, you can replatform COBOL and PL/I applications onto the mainframe-compatible managed runtime in the cloud. Powered by the AWS Mainframe Modernization Micro Focus replatform solution, this project approach is about porting the application where much of the application source code is recompiled without changes. The managed runtime provides the necessary capabilities to execute both online transactions and batch jobs. It preserves the investment in the language and code, while modernizing the underlying infrastructure for on-demand resources, elastic scaling, and DevOps best practices. Learn more about AWS Mainframe Modernization Replatform with Micro Focus.
- **Replatform with NTT DATA**: Using AWS Mainframe Modernization Replatform with NTT DATA, you can modernize and migrate COBOL applications to the AWS Cloud. With this toolchain you can also modernize mainframe relational databases, data sets, and hierarchical databases. The applications are modernized to the AWS Cloud with minimal changes to the source code. It supports both online transactions and batch jobs. The capability is available on demand with pay-as-you-go pricing. Learn more about AWS Mainframe Modernization Replatform with NTT DATA.
- **Application Testing**: AWS Mainframe Modernization Application Testing accelerates modernization testing timelines, with automation at scale, reducing effort and cost. Built-in scalability of the service allows you to scale from small to very large application testing use cases while leveraging on-demand testing resources in the AWS Cloud. AWS Mainframe Modernization Application Testing is a cloud-native service feature available from the AWS Management Console and provides seamless integration with AWS native resources. Application testing allows you to validate functional equivalence between source applications running on the mainframe and modernized target applications running on the AWS cloud. With AWS Mainframe Modernization Application Testing you can create automated test scenarios, run the scenarios on both the source application and the target applications, and compare the results with automation to ensure that the business logic implemented by the applications is unaltered during the modernization. Application Testing enables cost efficiency, scalability, and agility of testing using automation. Learn more about AWS Mainframe Modernization Application Testing.
- **Data Replication**: AWS Mainframe Modernization Data Replication with Precisely leverages near real-time data replication to migrate data from mainframe and midrange (iSeries) databases to the AWS Cloud. This capability enables a broad set of use cases related to application migration and data-driven innovation using AWS services. Data Replication builds resilient, high-performance data pipelines that replicate data changes via Change Data Capture (CDC), from mainframe and midrange database to target destinations such as Amazon Managed Streaming for Apache Kafka (MSK), Amazon RDS for PostgreSQL, Amazon RDS Aurora for PostgreSQL, Amazon RDS for Oracle, Amazon RDS for Microsoft SQL Server, IBM Db2 LUW on Amazon EC2, and Flat Files. By accelerating access to mainframe and midrange data on AWS, businesses can get modernization projects completed on time and on budget, while extending the value of mission-critical high-investment mainframe and midrange systems including building new business channels, offloading mainframe data processing, and driving business insights and intelligence with AI, data analytics and data lakes. Learn more about AWS Mainframe Modernization Data Replication with Precisely.
- **File Transfer**: AWS Mainframe Modernization File Transfer with BMC is a cloud-native service feature that helps accelerate and simplify modernization and

augmentation projects by providing automation for the transfer and codepage conversion of mainframe data sets and files to the AWS cloud. Using File Transfer you can discover mainframe data sets and files, select the data sets to transfer and the codepage conversion options, and initiate your transfer - all directly from the AWS Management Console. With File Transfer, you can achieve faster mainframe data set transfers to the AWS Cloud with compression, handle parallel transfers and conversion of data sets simultaneously and easily integrate file transfer operations with your existing CI/CD pipelines.

- **Developer IDE**: Mainframe Modernization offers an on-demand integrated development environment (IDE) so developers can write code quicker with smart editing and debugging, instant code compilation, and unit testing. You can also use the IDE to develop and enhance enterprise applications running on the Mainframe Modernization cloud native runtime. The Developer IDE is available on demand for both the refactor and replatform patterns.
- **Cloud Native Runtime**: The AWS Mainframe Modernization service not only facilitates migration and modernization of mainframe applications, but also allows executing these applications on a cloud native fully managed runtime. Similar to any other AWS service, this cloud native runtime is available and operated via AWS Console, AWS CLI, and APIs. It allows customers to leverage best practices of Infrastructure-as-Code (IaC) and DevOps via AWS CloudFormation. It is a fully-managed runtime which automates the management of resources such as compute, middleware, caching, locks, data, queues, transactions, and batch. The automation is built-in and facilitates operations such as installation, deployment, configuration, patching, scaling, versioning, and licensing. The fully-managed runtime is built for business-critical enterprise applications requiring high availability, reliability, and security. It benefits from native integrations with other AWS services such as Amazon CloudWatch for monitoring and logging, AWS IAM for security. It is available on demand for both the refactor and replatform patterns. This helps remove the burden and undifferentiated heavy lifting of managing a complex infrastructure. It allows you to focus on innovating and building new capabilities.

156.1.2 Benefits

- **Reduce Operating Costs**: Migrate and modernize your applications to remove the hardware and staffing costs of traditional mainframes.
- **Manage Modernization and Migration**: Break up and manage your complete migration with infrastructure, software, and tools to refactor and transform legacy applications.
- **Test and Operate Modernized Applications at Scale**: Accelerate modernization and regression testing at scale, with a cloud-native service

156.2 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/.

156.3 Pricing Overview

Please see the AWS UK G Cloud 14 Pricing Document accompanying this service in the Digital Marketplace.

156.4 Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://aws.amazon.com/mainframe-modernization/resources/
- **Service quotas:** https://docs.aws.amazon.com/general/latest/gr/m2.html
- **Service FAQs:** https://aws.amazon.com/mainframe-modernization/faqs/

156.5 Technical Requirements

Please refer to https://docs.aws.amazon.com/m2/latest/userguide/what-is-m2.html and the following links for comprehensive technical documentation regarding this service.

- AWS Mainframe Modernization Documentation

157 AWS Marketplace – BYOL

157.1 Service Overview

AWS Marketplace is a curated digital catalog that customers can use to find, buy, deploy, and manage third-party software, data, and services to build solutions and run their businesses. AWS Marketplace includes thousands of software listings from popular categories such as security, business applications, machine learning, and data products across specific industries, such as healthcare, financial services, and telecommunications. Customers can quickly launch preconfigured software, and choose software solutions in Amazon Machine Images (AMIs), software as a service (SaaS), and other formats.

157.2 AWS Marketplace (BYOL) on the G-Cloud 14 Framework

This Service Definition applies to the 'Bring Your Own Licence' (BYOL) service that is available through the AWS Marketplace under the G-Cloud 14 Framework.

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.

157.2.1 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 14 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).

157.3 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/.

157.4 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 14 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.

157.5 Service Constraints

Other products existing in the AWS Marketplace are NOT covered under G-Cloud 14 terms and are not included as part of this service offering.

157.6 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.

157.7 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.

158 AWS Migration Hub

158.1 Service Overview

AWS Migration Hub delivers a guided end-to-end migration and modernization journey through discovery, assessment, planning, and execution. Access the latest guidance and tools from one location to get automated recommendations, a prescriptive plan and cross-team collaboration and tracking to accelerate your transformation.

158.1.1 Features

- **Journey templates**: Migration Hub journeys make migrations easier to execute and track by showing all of the required steps to successfully migrate to AWS. A journey provides guidance on the recommended tasks for all phases of the migration to be completed. Migration Hub journeys also provide customizable step-by-step run books to execute each subtask, which reduces time for project planning and dependency on cloud experts.
- **Cross-team collaboration**: Migration Hub journeys allows collaboration across partners, practitioners and customers to keep you on-time and on-budget. Start your journey with a pre-defined set of roles (administrator, contributor), and invite other users (including partners and AWS personnel) to collaborate. Based on privileges, users can assign individual tasks to experts, get automatic notifications for delayed or blocked tasks and upload artifacts in one place to ensure context is never lost.
- **Application and server discovery**: Import information about on-premises servers and applications, or perform a deeper discovery using AWS Discovery Agent or AWS Discovery Collector, an agentless approach for VMware environments.
- **Dependency identification**: 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 analyzes your applications to help you determine the optimal strategy and tools to migrate and modernize at scale.
- **Orchestration**: Migration Hub Orchestrator helps accelerate your application migration with predefined workflow templates, for 1) the migrations of SAP NetWeaver–based applications with HANA databases, 2) the rehosting of any

applications to Amazon EC2, 3) the rehosting of SQL Server databases to Amazon EC2, 4) the replatforming of SQL Server databases to Amazon RDS, and 5) the importing of on-premises virtual machine images to AWS , and 6) the replatforming of .NET and Java applications to Amazon Elastic Container Service (ECS) on AWS Fargate.

- **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 is the starting point for incremental application refactoring to microservices. 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.

158.1.2 Benefits

- **Simplified experience**: Discover, assess, analyze, plan, execute and manage from a single location.
- **Expert guidance**: Accelerate migration and modernization projects with prescriptive journey templates.
- **Proven tools**: Use proven specialized services designed to meet your transformation goals.
- **No cost service**: Get started for free using AWS Migration Hub to plan or track a migration to AWS.

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 14 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/migrationhub/
- **Service quotas:** https://docs.aws.amazon.com/migrationhub/latest/ug/limits.html
- **Service FAQs:** https://aws.amazon.com/migration-hub/faqs/

158.5Technical 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.
- AWS Migration Hub Journeys User Guide: AWS Migration Hub Journeys is a service that you can use to plan, perform, and track migrations into AWS.
- Migration Hub API Reference: Describes the API operations available for Migration Hub along with sample requests, responses, and errors for the supported web services protocols.

159 AWS Network Firewall

159.1Service 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 set up 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. 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 or import compatible rules sourced from 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.

159.1.1 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.
- **Diverse ecosystem of partner integrations**: AWS Network Firewall integrates with AWS Partners for integration with central third-party policy orchestration and exporting logs to analytics solutions. AWS Network Firewall supports popular managed threat intelligence feeds for customers who prefer to leverage their existing managed rule providers. AWS Network Firewall can plug into third-party policy orchestration solutions for centrally managing hybrid or multiple firewall vendor architectures. For more visibility, AWS Network Firewall logs and security event information can be sent to third-party analytics solutions, such as Security Information and Event Management (SIEM) software. See a full list of AWS Network Firewall partners
- **Inspect encrypted traffic**: AWS Network Firewall supports Transport Layer Security (TLS) inspection, allowing customers to strengthen their security posture on AWS by improving visibility into encrypted traffic flows. You can use AWS Network Firewall to decrypt TLS sessions and inspect both inbound and outbound Amazon Virtual Private Cloud (VPC) traffic without the need to deploy or manage any additional network security infrastructure. Encryption and decryption happen on the same firewall instance natively, so traffic does not cross any network boundaries.

159.1.2 Benefits

- **Automatically scale**: Automatically scale your network firewall to protect your managed infrastructure
- **Protect workloads**: Protect your unique workloads with a flexible engine that can define thousands of custom rules
- **Centrally manage**: Centrally manage security policies across existing accounts and VPC’s and automatically enforce mandatory policies on new accounts

159.2 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/.

159.3 Pricing Overview

Please see the AWS UK G Cloud 14 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/network-firewall/
- **Service quotas:** https://docs.aws.amazon.com/network-firewall/latest/developerguide/quotas.html
- **Service FAQs:** https://aws.amazon.com/network-firewall/faqs/

159.5 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.
- API Reference: Describes the AWS Network Firewall API operations in detail. Also provides sample requests, responses, and errors for the supported web services protocols.

160 AWS Organizations

160.1 Service Overview

AWS Organizations helps you centrally govern your environment as you scale your workloads on AWS. Whether you are a growing startup or a large enterprise, Organizations helps you to programmatically create new accounts and allocate resources, simplify billing by setting up a single payment method for all of your accounts, create groups of accounts to organize your workflows, and apply policies to these groups for governance. In addition, AWS Organizations is integrated with other AWS services so you can define central configurations, security mechanisms, and resource sharing across accounts in your organization.

160.1.1 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 AWS IAM Identity Center 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 Access Manager (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.

160.1.2 Benefits

- **Scalability**: Quickly scale your environment by programmatically creating new AWS accounts for your resources and teams at no additional charge.
- **Governance**: Simplify user-based permission management to give teams the freedom to build while staying within targeted governance boundaries.
- **Cost optimisation**: Manage and optimize costs across your AWS accounts and resources.
- **Security**: Centrally secure and audit your environment across all of your AWS accounts.

160.2 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/.

160.3 Pricing Overview

Please see the AWS UK G Cloud 14 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/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/

160.5Technical 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.

161 AWS Outposts

161.1Service 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, data residency, and application migration with local system interdependencies.

AWS Outposts Family members include AWS Outposts rack and AWS Outposts servers.

The AWS Outposts rack is an industry standard 42U form factor. It provides the same AWS infrastructure, services, APIs, and tools to virtually any datacenter or co-location space. Outposts rack provides AWS compute, storage, database, and other services locally, while still allowing you to access the full range of AWS services available in the Region for a truly consistent hybrid experience. Scale from a single 42U rack to multiple rack deployments of up to 96 racks to create pools of compute and storage capacity.

The AWS Outposts servers come in a 1U or 2U form factor. They provide the same AWS infrastructure, services, APIs, and tools to on-premises and edge locations that have limited space or smaller capacity requirements, such as retail stores, branch offices, healthcare provider locations, or factory floors. Outposts servers provide local compute and networking services.

161.1.1 Features of AWS Outposts rack

- **Form factors**: The Outposts rack is 80 inches tall, 24 inches wide, and 48 inches deep. Inside are hosts, switches, a network patch panel, a power shelf, and blank panels.
- **Installation**: AWS delivers Outposts racks fully assembled and ready to be rolled into final position. Racks are installed by AWS and simply need to be plugged into power and network.
- **Locally supported services**: Amazon Elastic Compute Cloud (EC2), Amazon Elastic Container Service (ECS), Amazon Elastic Kubernetes Service (EKS), Amazon Elastic Block Store (EBS), Amazon EBS Snapshots, Amazon Simple Storage Service (S3), Amazon Relational Database Service (RDS), Amazon Elasticache, Amazon EMR, Application Load Balancer (ALB), Amazon Route 53 Resolver, CloudEndure, and VMware Cloud. Seamlessly extend Amazon Virtual Private Cloud on premises and run select AWS services locally on Outposts rack, and connect to a broad range of services available in the AWS Region.
- **Networking**: Includes integrated networking gear. Supports Local Gateway, which requires Border Gateway Protocol (BGP) over a routed network.
- **Power**: Supports three power configurations: 5 kVA, 10 kVA, or 15 kVA. The configuration of the power shelf depends on the total power draw of the Outpost capacity. Centralized redundant power conversion unit and a direct current (DC) distribution system in the backplane handled by line mate connectors.

161.1.2 Features of AWS Outposts servers

- **Form factors**: The Outposts rack-mountable servers fit inside 19" width, EIA-310 cabinets. The 1U high server is 24” deep, and uses AWS Graviton2 processors. The 2U high server is 30” deep and uses 3rd generation Intel Xeon Scalable processors.
- **Installation**: AWS delivers Outposts servers directly to you, installed by either onsite personnel install or a 3rd-party vendor. Once connected to your network, AWS will remotely provision compute and storage resources.
- **Locally supported services**: Amazon EC2, Amazon ECS, AWS IoT Greengrass, and Amazon Sagemaker Edge Manager. Seamlessly extend Amazon Virtual Private Cloud on premises and run select AWS services locally on Outposts servers, and connect to a broad range of services available in the AWS Region.
- **Networking**: Does not include integrated networking gear. Supports a simplified network integration experience providing a local Layer 2 presence..
- **Power**: Requires 1-2 kVA of power. Supports standard alternating current (AC) and direct current (DC) power options.

161.1.3 Benefits

- **Run AWS Services on premises**: Extend AWS compute, networking, security, and other services on premises for low latency, local data processing, and data residency needs.
- **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.

Use or disclosure of data contained on this sheet is subject to the restriction on the title page of this document.

161.2 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 (for Outposts rack) and here (for Outposts servers) or as may be provided to the customer during the Outpost order process.
- Customers are governed by the service responsibilities as set out in the AWS service terms here. 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.

161.3 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.

161.4 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.

161.5 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/.

161.6 Pricing Overview

Please see the AWS UK G Cloud 14 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.

161.6.1 Operating Systems

Prices do not include operating system (OS) charges. OS charges will be applied based on usage per instance.

161.6.2 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.

161.6.3 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.

161.7 Service Constraints

161.7.1 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 Rack FAQs and Outposts Servers FAQs to understand what will happen if your external network connectivity is unavailable.

161.7.2 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.

161.8Technical Requirements

Please refer to https://docs.aws.amazon.com/outposts/ for comprehensive technical documentation.

162 AWS Panorama

162.1Service Overview

AWS Panorama is a new machine learning (ML) appliance and software development kit (SDK) allowing organizations to bring computer vision (CV) to their on-premises cameras to make automated predictions with high accuracy and low latency. With AWS Panorama, companies can use compute power at the edge (without requiring video streamed to the cloud) to improve their operations. Panorama automates monitoring and visual inspection tasks like evaluating manufacturing quality, finding bottlenecks in industrial processes, and assessing worker safety within their facilities.

162.1.1 Features

- **Set up in minutes**: Once installed and connected to your network, AWS Panorama devices connect to the AWS Management Console where you can register your AWS Panorama device and add video feeds from onsite cameras, deploy trained ML models, and run applications in minutes.
- **Connect IP cameras**: AWS Panorama devices support RTSP-enabled IP cameras, and allow you to add IP cameras that support the ONVIF standard.
- **Parallel multi-model multi-stream support**: AWS Panorama devices support connecting to multiple camera streams at a given time and support running multiple ML models per stream.
- **Inference at the edge**: AWS Panorama 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 offers a flexible option for adding CV to automate tasks that traditionally require human inspection and monitoring.
- **GPU compute**: AWS Panorama devices have 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 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. AWS Panorama 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.

162.1.2 Benefits

- **Computer vision**: Add computer vision (CV) to your existing fleet of cameras with AWS Panorama devices, which integrate seamlessly with your local area network.

- **Analyze video feeds**: Make predictions locally with high accuracy and low latency from a single management interface, where you can analyze video feeds in milliseconds.
- **Operate with limited internet bandwidth**: Process video feeds at the edge, so you can control where your data is stored and operate with limited internet bandwidth.

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 14 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/panorama/
- **Service quotas:** https://docs.aws.amazon.com/panorama/latest/dev/gettingstarted-quotas.html
- **Service FAQs:** https://aws.amazon.com/panorama/faqs/

162.5 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.

163 AWS PrivateLink

163.1 Service Overview

AWS PrivateLink enables customers to access services hosted on AWS in a highly available and scalable manner, while keeping all the network traffic within the AWS network. Service users can privately access services powered by PrivateLink from their Amazon Virtual Private Cloud (VPC) or their on-premises, without using public IPs, and without requiring traffic to traverse across the Internet. Service owners can register their Network Load Balancers to PrivateLink services in order to provide their services to other AWS customers.

163.1.1 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.
- **Regulatory compliance**: Preventing your sensitive data, such as customer records, from traversing the internet helps you maintain compliance with regulations such as HIPAA, EU-US Privacy Shield, and payment card industry (PCI). This is especially critical to customers in the financial services, healthcare, and government sectors. With AWS PrivateLink, traffic between AWS resources, VPCs, and third-party services stays on the AWS network where there are robust controls in place to maintain security and compliance. Controls include compliance alignment with standard financial regulations such as the SEC Rule 17a-4(f) and the Japanese Financial Intelligence Centres (JAFIC).

163.1.2 Benefits

- **Security**: Secure your traffic by using private IP addresses when exchanging data with your software as a service (SaaS) applications.
- **Cost savings**: Connect with simplified network and firewall management rules and reduced data output and NAT costs.
- **Speed**: Accelerate cloud migrations by combining PrivateLink with AWS Direct Connect or a VPN.
- **Compliance**: Deliver SaaS services while staying compliant with HIPAA, EU-US Privacy Shield, PCI, and other regulations.

163.2 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/.

163.3 Pricing Overview

Please see the AWS UK G Cloud 14 Pricing Document accompanying this service in the Digital Marketplace.

163.4 Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/vpc/latest/privatelink/what-is-privatelink.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

163.5 Technical Requirements

Please refer to https://docs.aws.amazon.com/vpc/latest/privatelink/what-is-privatelink.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.

164 AWS Resilience Hub

164.1 Service Overview

AWS Resilience Hub is a central location for you to manage and improve the resilience posture of your applications on AWS. AWS Resilience Hub enables you to define your resilience goals, assess your resilience posture against those goals, and implement recommendations for improvement based on the AWS Well-Architected Framework.

Within AWS Resilience Hub, you can also create and run Amazon Fault Injection Service experiments, which mimic real-life disruptions to your application to help you better understand dependencies and uncover potential weaknesses. AWS Resilience Hub provides a central place with all the AWS services and tools that you need to continuously strengthen your resilience posture. AWS Resilience Hub works with other services to provide recommendations and help you to manage your application resources.

164.1.1 Features

- **Describe**: Describe your applications as resource collections, such as CloudFormation stacks, Terraform state files, AppRegistry applications, or resource groups, or define applications for Kubernetes workloads that are managed on Amazon EKS. Applications can also be described using both resource collections and Amazon EKS clusters.
- **Define**: Define the resilience policies for your applications. These policies include RTO and RPO targets for applications, infrastructure, Availability Zone, and Region disruptions.
- **Assess**: AWS Resilience Hub's assessment uses best practices from the AWS Well-Architected Framework to analyze the components of an application and uncover potential resilience weaknesses. These can be caused by incomplete infrastructure setup, misconfigurations, or situations where additional configuration improvements are needed.
- **Recommend**: AWS Resilience Hub provides actionable recommendations to improve resilience. The resilience assessment also generates code snippets that help you create recovery procedures as AWS Systems Manager documents for your applications, referred to as Standard Operating Procedures (SOPs). AWS Resilience Hub generates a list of recommended Amazon CloudWatch monitors and alarms to help the operator quickly identify any change to the application's resilience posture once deployed.
- **Validate**: After the application and SOPs have been updated to incorporate recommendations from the resilience assessment, you can use AWS Resilience Hub to test and verify that your application can meet its resilience targets before releasing it into production. AWS Resilience Hub is integrated with AWS Fault Injection Simulator (FIS), a chaos engineering service, to provide fault-injection simulations of real-world failures to validate that the application recovers within the defined resilience targets. This can include network errors or too many open connections to a database. AWS Resilience Hub also provides APIs so you can integrate its resilience assessment and testing into your CI/CD pipelines for ongoing resilience validation. Integrating resilience validation into CI/CD pipelines helps ensure that changes to the application's underlying infrastructure do not compromise resilience.

164.1.2 Benefits

- **Reduce outages**: Continuously validate and track application resilience to reduce outages
- **Evaluate resilience targets**: Evaluate resilience targets (Recovery Time Objective and Recovery Point Objective).
- **Proactive issue management**: Identify and resolve issues before they occur in production.
- **Optimize continuity**: Optimize business continuity while reducing recovery costs.

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 14 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/resilience-hub/
- **Service quotas:** https://docs.aws.amazon.com/general/latest/gr/resiliencehub.html
- **Service FAQs**: https://aws.amazon.com/resilience-hub/resources/

164.5 Technical Requirements

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

- User Guide: Describes how to use AWS Resilience Hub to analyze your infrastructure, get recommendations to improve the resiliency of your AWS apps, review resiliency scores, and more.
- API Reference: Describes all the AWS Resilience Hub API operations in detail. Also provides sample requests, responses, and errors for the supported web services protocols.

165 AWS Resource Access Manager (RAM)

165.1 Service Overview

AWS Resource Access Manager (AWS 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 resources with other AWS accounts. This eliminates the need to provision and manage resources in every account. When you share a resource with another account, that account is granted access to the resource and any policies and permissions in that account apply to the shared resource.

165.1.1 Features

- **Discovering your AWS resources**: Search for and discover your resources across the AWS Regions in your AWS account.
- **Grouping your AWS resources**: Group your resources to manage them as a single unit instead of individually. You can create resource groups that consist of all of the resources that are part of an AWS CloudFormation stack. You can also create a group of all resources that are tagged with specific keys and values.
- **Sharing your AWS resources**: Use AWS Resource Access Manager (AWS RAM) to share resources you create in one AWS account with other accounts. If your account is part of an organization, then you can share your resources with all of the accounts in the organization, or with specific organizational units (OUs). Some resource types let you share with individual AWS Identity and Access Management (IAM) users and roles.
- **Tagging your AWS resources**: Attach string key and value pairs to your resources to categorize them for a variety of purposes, including cost allocation and security by using attribute-based access control.

165.1.2 Benefits

- **Reduces your operational overhead**: Create a resource once, and then use AWS RAM to share that resource with other accounts. This eliminates the need to provision duplicate resources in every account, which reduces operational overhead. Within the account that owns the resource, AWS RAM simplifies granting access to every role and user in that account without having to use identity-based permission policies.
- **Provides security and consistency**: Simplify security management for your shared resources by using a single set of policies and permissions. If you were to instead create duplicate resources in all your separate accounts, you would have the task of implementing identical policies and permissions, and then have to keep them identical across all those accounts. Instead, all users of an AWS RAM resource share are managed by a single set of policies and permissions. AWS RAM offers a consistent experience for sharing different types of AWS resources.
- **Provides visibility and auditability**: View the usage details for your shared resources through the integration of AWS RAM with Amazon CloudWatch and AWS CloudTrail. AWS RAM provides comprehensive visibility into shared resources and accounts.

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 14 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/ARG/
- **Service quotas:** https://docs.aws.amazon.com/ram/latest/userguide/service-quotas.html

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

165.5Technical Requirements

Please refer to https://docs.aws.amazon.com/ARG/ 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.

166 AWS Secrets Manager

166.1Service Overview

AWS Secrets Manager is a secrets management service that helps you protect access to your applications, services, and IT resources. This service enables you to easily rotate, manage, and retrieve database credentials, API keys, and other secrets throughout their lifecycle. Using Secrets Manager, you can secure and manage secrets used to access resources in the AWS Cloud, on third-party services, and on-premises.

166.1.1 Features

- **Secure secrets storage**: AWS Secrets Manager encrypts secrets at rest using encryption keys that you own and store in AWS Key Management Service (AWS KMS). When you retrieve a secret, Secrets Manager decrypts the secret and transmits it securely over TLS to your local environment. Secrets Manager integrates with AWS Identity and Access Management (IAM) to control access to the secret using fine-grained IAM policies and resource-based policies.
- **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. 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 secrets used with other AWS or 3P services by modifying sample Lambda functions.
- **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. Specify the AWS Regions where a secret needs 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 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**: Build your applications with security of secrets top of mind. Secrets Manager provides code samples to call Secrets Manager APIs from common programming languages. There are two types of APIs to retrieve secrets:

  - Retrieve a single secret by name or ARN.
  - Retrieve a group of secrets by providing a list of names or ARNs, or filter criteria such as tags.

  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 availability and reduce latency during secrets retrieval.
- **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 created 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 meet compliance requirements.
  - Use AWS Config Rules to help you verify that your secrets are configured in accordance with your organization's security and compliance requirements.
  - 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.
- **Secrets Manager Integration**: AWS services integrate with Secrets Manager to securely manage your credentials. These integrations help you securely exchange credentials with various AWS services. The credentials stored in Secrets Manager are encrypted either using AWS managed KMS keys or customer managed keys. Secrets Manager rotates secrets periodically to keep the security bar high. Once your secrets are stored with Secrets Manager, you will be able to provide the ARN of a secret instead of a plain text credential to an AWS service.

166.1.2 Benefits

- **Store secrets securely**: Centrally store and manage credentials, API keys, and other secrets.
- **Manage access with fine-grained policies**: Use AWS Identity and Access Management (IAM) permissions policies to manage access to your secrets.
- **Automate secrets rotation**: Rotate secrets on demand or on a schedule, without redeploying or disrupting active applications.
- **Audit and monitor secrets usage**: Integrate secrets with AWS logging, monitoring, and notification services.

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 14 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/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/

166.5 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.

167 AWS Security Hub

167.1 Service Overview

AWS Security Hub is a cloud security posture management (CSPM) service that performs automated, continuous security best practice checks against your AWS resources to help you identify misconfigurations, and aggregates your security alerts (i.e. findings) in a standardized format so that you can more easily enrich, investigate, and remediate them.

167.1.1 Features

- **Automated, continuous security best practice checks**: The AWS Foundational Security Best Practices standard comes built into Security Hub. This is a highly-curated set of security best practices vetted by AWS security experts that give you event-based continuous monitoring, or run on a periodic schedule. Each control has a specific severity rating 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.
- **Security standards aligned to regulatory and industry compliance frameworks**: In addition to the AWS Foundational Security Best Practices standard, Security Hub offers additional standards aligned to industry and regulatory frameworks, such as the Payment Card Industry Data Security Standard (PCI DSS), the Center for Internet Security (CIS) AWS Foundations Benchmark, and the National Institute of Standards and Technology (NIST). These standards are also powered by continuous, automated security checks, and you only pay once for the security check—regardless of how many standards it is mapped to.
- **Evaluate your security posture with security score**: 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, or Organization.

- **Fine-tune and centrally configure your CSPM controls to fit the needs of your organization**: Customize Security Hub controls in accordance with the specific security guidelines of your organization, without forgoing the benefits of using managed controls. You can modify the parameter values in many Security Hub controls, reducing the manual efforts of building and testing these across your accounts by hand, while still maintaining security scoring for them. Specify parameters such as the number of days until a resource is considered unused, specific characteristic of a password policy, or a list of high-risk ports. You can also centralize these configurations and capabilities for all or some accounts globally, without needing to update them account-by-account and Region-by-Region.
- **Gain powerful insights with security data visualization**: Customize your Security Hub dashboard according to your specific requirements to more easily identify patterns, vulnerabilities, and threats— leading to faster response. Security Hub's dashboard features a set of AWS managed insights that were carefully chosen to reflect the modern cloud security threat landscape as observed by AWS, and guided by lessons learned from AWS's own security operations. You can select and modify the widgets you want to display, apply and save filters to create contextual views by specific criteria, and prioritize the data and view of your organization's security posture that fits your needs.
- **Cross-Region aggregation of findings**: Designate an aggregator Region and link some or all Regions to give you a centralized view of your findings across your accounts and your linked Regions. Findings are continuously synced between the Regions, so that updates made to a finding in one Region is replicated to the other Region. 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.
- **Aggregate 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. 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.
- **Consolidated control findings and a consolidated controls view**: Simplify how you triage, investigate, and remediate findings by consolidating control findings across multiple standards to more easily identify misconfigurations based on severity and number of failed resources and improve your overall security score. You can also view all enabled controls in one place, along with their compliance status and a summary of passed and failed security checks, and configure each control across all standards in a single action.
- **A single, standardized data format for all of your findings**: Security Hub eliminates time-consuming and resource-intensive data normalization processes by introducing the AWS Security Findings Format (ASFF). With the ASFF, Security Hub 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.

Use or disclosure of data contained on this sheet is subject to the restriction on the title page of this document.

- **Multi-account and AWS Organizations support**: Connect multiple AWS accounts and consolidate findings across those accounts with a few clicks in the 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.
- **Filtering, grouping, and saved searches for your findings**: Filter findings based on fields in the ASFF 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.
- **Automated updating of findings**: Automatically update or suppress findings in near-real time with Security Hub Automation Rules. Security admins can create rules with specific criteria that can automatically be evaluated against every incoming finding and update findings fields if matched. Use automation rules to change the severity or workflow status of specific findings, suppress them, or update their user-defined fields.
- **Automated response, remediation, and enrichment actions**: Create custom automated response, remediation, and enrichment workflows using Security Hub's integration with Amazon EventBridge. 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. Security Hub also supports sending findings to EventBridge on demand via custom actions, and the Security Hub Automated Response and Remediation (ASR) solution provides you with prepackaged EventBridge rules for you to deploy via AWS Cloud Formation.
- **Integrations with ticketing, chat, incident management, investigation, GRC, SOAR, and SIEM tools**: Security Hub 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 send or receive findings from Security Hub.
- **Cost optimization**: You can try Security Hub at no charge with a 30-day free trial. The trial includes the complete Security Hub feature set and security best practice checks. Every AWS account in each Region that is enabled with Security Hub receives a free trial. The free trial will provide you an estimate of your monthly bill if you continue using Security Hub across the same accounts and regions. Security Hub also offers a perpetual free tier of 10,000 findings ingested per account per Region per month. Learn more about Security Hub pricing.
- **Volume pricing**: Security Hub is priced along three dimensions: the quantity of security checks, the quantity of ingested findings, and the quantity of rule evaluations processed per month. With AWS Organizations support, Security Hub allows you to connect multiple AWS accounts and consolidate findings across those accounts to enjoy tiered pricing for your entire organization's security checks, finding ingestions, and automation rule evaluations.
- **Single cost for security checks**: Security Hub will automatically and continuously conduct best practice checks against pre-packaged security standards to evaluate the security posture of your AWS accounts and resources. For checks conducted

against identical controls that are common across the different standards available in Security Hub, you will not be charged for duplicative checks; Security Hub will only charge you once.

167.1.2 Benefits

- **Continuous monitoring**: Detect deviations from security best practices with a single click.
- **Security finding aggregation**: Automatically aggregate security findings in a standardized data format from AWS and partner services.
- **Initiate automated responses**: Accelerate mean time to resolution with automated response and remediation actions.
- **Security posture**: Visualize the security posture of your AWS-based applications.

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 14 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/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/

167.5 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.

168 AWS Service Catalog

168.1 Service Overview

AWS Service Catalog allows IT administrators to create, manage, and distribute catalogs of approved products to end users, who can then access the products they need in a personalized portal. Administrators can control which users have access to each product to enforce compliance with organizational business policies. Administrators can also setup adopted roles so that end users only require IAM access to AWS Service Catalog in order to

deploy approved resources. AWS Service Catalog allows your organization to benefit from increased agility and reduced costs because end users can find and launch only the products they need from a catalog that you control.

168.1.1 Features

- **Products**: A product is a set of AWS cloud resources 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. Using AWS Service Catalog, you can create products in multiple ways including AWS CloudFormation templates or Terraform configurations. The templates and configurations 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.
- **Stack or Resource Group**: AWS Service Catalog products are either AWS CloudFormation type products or Terraform open source type or Terraform Cloud products. AWS CloudFormation type products are launched as an AWS CloudFormation stack, which is a set of resources provisioned for that instance of the product. For Terraform open source and Terraform Cloud type products, the underlying resources are stored and tagged in AWS Resource Groups. AWS CloudFormation stacks and AWS Resource Groups 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.
- **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. A portfolio can contain a mix of products of different types including AWS CloudFormation type products or Terraform open source and Terraform Cloud type products.
- **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. 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 infrastructure as code (IaC) 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. Launch constraints are supported for Terraform open source and Terraform Cloud type products. Template constraints are not yet available for Terraform open source and Terraform Cloud type products.
- **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. Service actions are not available yet for Terraform open source and Terraform Cloud type products.
- **AWS Service Catalog AppRegistry features**: 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.

168.1.2 Benefits

- **Easy to deploy**: Quickly find and deploy approved, self-service cloud resources.
- **Agility**: Stay agile while improving governance over resources across multiple accounts.

- **Streamlining**: Streamline workflows by connecting to ServiceNow and Jira Service Management.
- **Accuracy**: Get up-to-date, accurate application definitions and metadata with AWS Service Catalog AppRegistry.

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 14 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/servicecatalog/
- **Service quotas:** https://docs.aws.amazon.com/servicecatalog/latest/adminguide/limits.html
- **Service FAQs:** https://aws.amazon.com/servicecatalog/faqs/

168.5 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.

169 AWS Shield

169.1 Service Overview

AWS Shield is a managed service that provides protection against Distributed Denial of Service (DDoS) attacks for applications running on AWS. AWS Shield Standard is automatically enabled to all AWS customers at no additional cost. AWS Shield Advanced is an optional paid service. AWS Shield Advanced provides additional protections against more sophisticated and larger attacks for your applications running on Amazon Elastic Compute Cloud (EC2), Elastic Load Balancing (ELB), Amazon CloudFront, AWS Global Accelerator, and Route 53.

169.1.1 Features

- **AWS Shield Standard**: 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 website 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.

- **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.
- **AWS Shield Advanced**: 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, 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. Shield Advanced also gives you 24/7 access to the AWS Shield Response Team (SRT) and protection against DDoS-related spikes in your EC2, ELB, CloudFront, Global Accelerator, and Route 53 charges.
- **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.
- **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.
- **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.
- **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.
- **Specialized support**: For customers on Business or Enterprise support plans, AWS Shield Advanced gives you 24/7 access to the SRT, which can be engaged before, during, or after a DDoS attack. The SRT will help triage the incidents, identify root causes, and apply mitigations on your behalf. The SRT has deep expertise in rapidly responding to and mitigating DDoS attacks across AWS customers.
- **Global availability**: AWS Shield Advanced is available globally on all CloudFront, Global Accelerator, and Route 53 edge locations. You can protect your web applications hosted anywhere in the world by deploying CloudFront in front of your application. Your origin servers can be Amazon Simple Storage Service (S3), EC2, ELB, or a custom server outside of AWS. You can also activate protections directly on Elastic IP or ELB instances in all AWS Regions where Shield Advanced is available.
- **Centralized protection management**: AWS Shield Advanced customers can use AWS Firewall Manager to apply Shield Advanced and AWS WAF protections across their entire organization. The cost of Firewall Manager is included in the Shield Advanced subscription fee. Using Firewall Manager, you can automatically configure policies covering multiple accounts and resources.

169.1.2 Benefits

- **Security**: Automatically detect and mitigate sophisticated network-level distributed denial of service (DDoS) events.
- **Risks**: Customize application protection against DDoS risks through integrations with Shield Response Team (SRT) protocol or AWS WAF.
- **Visibility**: Gain visibility, insights, and cost savings for DDoS events that impact your AWS resources.

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

169.5 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.

170 AWS SimSpace Weaver

170.1 Service Overview

AWS SimSpace Weaver is a service that you can use to build and run large-scale spatial simulations in the AWS Cloud. For example, you can create crowd simulations, large real-world environments, and immersive and interactive experiences.

With SimSpace Weaver, you can distribute simulation workloads across multiple Amazon Elastic Compute Cloud (Amazon EC2) instances. SimSpace Weaver deploys the underlying AWS infrastructure for you, and handles the simulation data management and network communication between the Amazon EC2 instances running your simulation.

170.1.1 Features

- **Managed infrastructure**: AWS SimSpace Weaver handles deploying and managing your simulation infrastructure. You can scale your spatial simulations across up to 10 Amazon Elastic Compute Cloud (Amazon EC2) instances while SimSpace Weaver manages the challenging tasks of provisioning, networking, and deprovisioning infrastructure for you. SimSpace Weaver lets you spend less time worrying about infrastructure and more time developing simulation code and content to fill your virtual worlds.
- **Run your own simulators**: SimSpace Weaver is a flexible service that gives you the opportunity to scale up your own simulation engines. Integrate with either the SimSpace Weaver app SDK for C++ or SimSpace Weaver app SDK for Python to create AWS SimSpace Weaver applications (or simply apps). Then, upload your apps to Amazon Simple Storage Service (Amazon S3) and use them as part of your simulation. Apps can take many different forms and perform various read and write functions based on your simulation code. Within a single simulation, these apps run distributed across your chosen quantity of Amazon EC2 instances while SimSpace Weaver manages the communication and simulation state between them. You can build simulation apps that manage object pathfinding and create specialized apps that send simulation data to a client for visualization purposes.

- **Spatial partitioning**: SimSpace Weaver works by dividing the simulation world into discrete areas called spatial partitions. SimSpace Weaver automatically partitions your simulation world according to your predefined configuration. Each spatial area is managed by a SimSpace Weaver app that is responsible for the simulation logic and behavior of all entities within that area. SimSpace Weaver currently supports partitioning into a uniform 2D grid.
- **Cross-instance data replication**: SimSpace Weaver maintains the global state of all simulated entities, meaning that every entity can be aware of each other's spatial presence and state. SimSpace Weaver also manages the memory and replication of entity data between servers, allowing entities to seamlessly cross boundaries between different simulation areas.
- **Subscriptions**: Applications can read simulation data from other areas of the simulation using the subscription system. This means simulation objects can account for events happening in neighboring partitions or interact with other objects along partition boundaries. Use subscriptions to read data out of the simulation for visualization and interactivity.
- **Messaging**: The SimSpace Weaver app SDK contains Messaging APIs that you can use to send and receive messages between applications in a simulation. Subscriptions provide a pull method to read simulation events; messaging provides a push method to transmit data in the simulation. SimSpace Weaver messages contain the sender in the message header so message recipients can reply. You can use messages to forward information that an application receives from an external client to other applications in the simulation.
- **Multiple spatial domains**: Run simulations with different sets of logic and behaviors in the same simulated environment. Each domain has its own allocated compute resources. Entities can interact and transfer across domains so the simulation still behaves as a single, seamless world. Use multi-domain for more advanced simulations that combine different types of entities that have differing compute needs. For example, a simulation of a virtual city may utilize multi-domain to have a domain dedicated to simulating pedestrians and another dedicated to simulating traffic.
- **Snapshots**: Snapshots allow you to save the state of your simulation entities at any point in time. Snapshot files are uploaded and stored in Amazon S3. Use your snapshot files to launch new simulations with the existing Start Simulation API. Snapshots are an ideal tool for simulations with branching scenarios and addressing backup and restore use cases. Use snapshots to save your simulation progress at regular intervals, especially for long running simulations.
- **AWS SimSpace Weaver Local**: Run simulation code by using your own local hardware to accelerate iteration times and test for free. AWS SimSpace Weaver Local uses the same APIs as the cloud service, so you can seamlessly transition to the cloud for full-scale production without needing to modify any code. AWS SimSpace Weaver Local is currently available for Windows.
- **Unreal and Unity integrations**: SimSpace Weaver offers built-in integrations with Unreal Engine 5 and Unity LTS Release 2021.3.7f1 to accelerate your development with these popular tools.
- **Simulation tick rate (clock speeds)**: The default update rate for SimSpace Weaver simulations is 30 Hz. You can adjust the update rate to 10 Hz, 15 Hz, or allow the clock to run as fast as possible to fit your simulation needs while still keeping all your applications time synchronized.
- **Custom Containers**: You can utilize your own container images to run your simulation applications instead of AWS SimSpace Weaver's stock Amazon Linux (AL2) container. Create your own AL2 based container images with any dependencies required by your simulation applications and upload the images to

Amazon Elastic Container Registry (Amazon ECR). Use this feature to help reduce the file size of your simulation applications, increase upload speeds, and improve iteration times.

- **Pay-as-you-go pricing**: SimSpace Weaver has no upfront commitments. You pay only for the compute resources used by your simulation for however long you use them.

170.1.2 Benefits

- **3D simulations**: Understand possible real-world outcomes and visualize immersive environments with massive, city-scale 3D simulations.
- **Fully managed**: Increase the size, complexity, and number of objects in your simulations, distributed across instances fully managed by SimSpace Weaver.
- **Time saving**: Avoid time-consuming challenges surrounding data management, replication, and time synchronization across servers.
- **Integration**: Integrate seamlessly with popular development toolkits and simulation engines, such as Unreal Engine and Unity.

170.2 Backup/Restore and Disaster Recovery

The service provides backup and recovery. Users can create snapshots to backup your simulation at any point.

170.3 Pricing Overview

Please see the AWS UK G Cloud 14 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/simspaceweaver/
- **Service quotas:** https://docs.aws.amazon.com/simspaceweaver/latest/userguide/service-quotas.html
- **Service FAQs:** https://aws.amazon.com/simspaceweaver/

170.5 Technical Requirements

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

- User Guide: Learn how to use SimSpace Weaver to build and run large-scale simulations and games in the AWS Cloud.
- API Reference: Look up all AWS SDK operations and data types for SimSpace Weaver. The API contains operations for all SimSpace Weaver components.
- Guide Catalog: Find documentation for previous releases of SimSpace Weaver.

171 AWS Snowball Edge

171.1 Service Overview

AWS Snowball Edge is a type of Snowball device with on-board storage and compute power for select AWS capabilities. Snowball Edge can do local processing and edge-computing

workloads in addition to transferring data between your local environment and the AWS Cloud.

171.1.1 Features

- **Built for edge computing & data migration**: AWS Snowball Edge is available in two device types. Snowball Edge Compute Optimized with more computing capabilities, suited for higher performance workloads and Snowball Edge Storage Optimized with more storage, which is suited for large-scale data migrations and capacity-oriented workloads.
- **Powerful computing resources**: Snowball Edge Compute Optimized provides powerful computing resources for use cases such as machine learning, full motion video analysis, analytics, and local computing stacks. These capabilities include 104 vCPUs, 416 GiB of memory, and an optional NVIDIA Tesla V100 GPU. For storage, the device provides 28 TB usable NVMe SSD capacity for Amazon S3 compatible storage or EBS-compatible block volumes. Snowball Edge Compute Optimized devices run Amazon EC2 sbe-c and sbe-g instances, which are equivalent to C5, M5a, G3, and P3 instances.
- **Large-scale data migrations**: Snowball Edge Storage Optimized devices are well suited for large-scale data migrations and recurring transfer workflows, as well as local computing with higher capacity needs. Snowball Edge Storage Optimized provides 80TB HDD or 210TB NVMe capacity for Amazon S3-compatible object storage. For computing resources, the 80TB device provides 40 vCPUs, and 80 GB of RAM.
- **High-speed network**: 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.
- **Fast data migration**: Snowball devices feature high-speed network connections, supporting 10 Gbps to 100 Gbps links with dual 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 migration times. For device specific networking specifications, please see Snowball documentation.
- **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 create a storage cluster using multiple Snowball devices when running edge computing jobs for increased storage durability of on-premise applications. When creating a new job in the Console, select the option to make a storage 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 migration jobs or devices with EKS Anywhere enabled.
- **Kubernetes**: You can create an Amazon EKS Anywhere cluster using one or more Snowball devices for your Kubernetes deployments. When creating a new job in the Console, under the Features and options step, you can choose the number of devices you want in your Kubernetes cluster and to include Amazon EKS Anywhere on them.

Use or disclosure of data contained on this sheet is subject to the restriction on the title page of this document.

- **Amazon S3 compatible storage on Snowball Edge local compute and storage devices**: Applications can work with S3 compatible storage through S3 API compatible endpoint accessed through AWS OpsHub, S3 SDK or CLI. For specific information, see the API documentation.
- **Amazon S3 Adapter on Snowball Edge data migration devices**: You can transfer data with S3 Adapter on Snow through an S3-API compatible endpoint accessed through AWS OpsHub, S3 SDK or CLI. For specific information, see the API documentation.
- **Tape data migration**: You can use a Snowball Edge Storage Optimized device with Tape Gateway to quickly and securely migrate petabyte-scale physical tape data to S3 Glacier Flexible Retrieval or S3 Glacier Deep Archive. You retain your tape-based backup workflows and don't need to use your network connection to AWS for data transfer. After your tape data is migrated to AWS, you can view your virtual tapes through the AWS Storage Gateway console and access data on them through a Tape Gateway that runs on premises as a virtual machine or hardware appliance or on an Amazon Elastic Compute Cloud (Amazon EC2) instance on AWS. Tape Gateway is currently available only on the 80TB Snowball Edge device. Learn more about Tape Gateway.
- **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 devices. This makes it easier to develop applications in EC2, and then run them in disconnected and remote locations. Snowball supports both performance optimized and capacity optimized volume types.
- **Secure, resilient Amazon S3 compatible storage**: Amazon S3 compatible storage is ideal for applications with object storage that require local, real-time decisions or are bound by data residency requirements in rugged, mobile edge, and disconnected environments. Use common Amazon S3 APIs and features to easily develop, deploy and manage applications requiring object storage. Benefit from built-in resiliency, security, performance, and the most frequently used Amazon S3 features such as PUT, GET, Multipart Uploads, Object tagging, bucket operations, and bucket lifecycle policies in disconnected edge environments.

171.1.2 Benefits

- **Speed**: Each Snowball Edge device can transport data at speeds faster than the internet. This transport is done by shipping the data in the appliances through a regional carrier. The appliances are rugged, complete with E Ink shipping labels.
- **Service Status**: Snowball Edge devices come with an on-board LCD display that can be used to manage network connections and get service status information.
- **Clustering**: You can cluster Snowball Edge devices for local storage and compute jobs to achieve data durability across 3 to 16 devices and locally grow or shrink storage on demand.
- **Storage**: Large amounts of storage capacity or compute functionality for devices. This depends on the options you choose when you create your job. Snowball Edge devices have four options for device configurations—Storage Optimized, Compute Optimized, Compute Optimized with GPU, and Import virtual tapes into AWS Storage Gateway.

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 14 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/snowball/latest/developer-guide/whatisedge.html
- **Service quotas:** https://aws.amazon.com/snowball/faqs/
- **Service FAQs:** https://docs.aws.amazon.com/snowball/latest/developer-guide/limits.html

171.5 Technical Requirements

Please refer to https://docs.aws.amazon.com/snowball/latest/developer-guide/whatisedge.html and the following links for comprehensive technical documentation regarding this service.

- AWS Snowball Edge Developer Guide: Provides a conceptual overview of how to use an AWS Snowball Edge device. Includes guidance for local storage and compute, clustering, importing and exporting data into Amazon S3, and other features of an Edge device. Device configuration options include Storage optimized, Compute optimized, and Compute optimized with GPU.

172 AWS Step Functions

172.1 Service Overview

AWS Step Functions is a fully managed service that makes it easier to coordinate the components of distributed applications and microservices using visual workflows. Building applications from individual components that each perform a discrete function helps you scale more easily and change applications more quickly.

Step Functions is a reliable way to coordinate components and step through the functions of your application. Step Functions provides a graphical console to arrange and visualize the components of your application as a series of steps. This makes it easier to build and run multi-step applications.

Step Functions automatically triggers and tracks each step and retries when there are errors, so your application executes in order and as expected. Step Functions logs the state of each step, so when things do go wrong, you can diagnose and debug problems more quickly. You can change and add steps without even writing code, so you can more easily evolve your application and innovate faster.

172.1.1 Features

- **Visual and easy -to-use 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 error handling, timeouts and parallel processing flow for distributed applications**: 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. Learn more about states.
- **Direct integration with 220 AWS services and over 10,000 APIs**: Using AWS Step Functions Service Integrations, you can configure your Step Functions workflow to call over 220 AWS services.
- **Integrate with virtually any SaaS application directly from your workflow**: AWS Step Functions HTTPS endpoints integration enables you to directly integrate HTTP-based services into your workflows. You can build, manage, and orchestrate multi-step workflows composed of AWS services and SaaS applications from a single location. With the HTTPS endpoints integration you can easily coordinate the various components of your application, manage authentication, test integration, and handle errors without needing to write custom code or maintain additional infrastructure. HTTPS endpoints integration improves developer agility and enables you to easily extend the functionality of SaaS applications that support your business.
- **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 Reusability enables robust and efficient application development**: 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.
- **Shorten development cycles with separation of workflow logic**: 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.
- **Built-in state management and workflow progress tracking**: 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 clean up and recovery code. Learn more about Step Functions error handling and how you can handle error conditions using a state machine.
- **Real-time and auditable workflow execution history**: 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 operator dashboard for monitoring and troubleshooting**: 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.
- **High availability and fault tolerant workflows:** AWS Step Functions has built-in fault tolerance and maintains service capacity across multiple Availability Zones in each region to protect applications against individual machine or data center failures. This helps ensure high availability for both the service itself and for the application workflow it operates.
- **Serverless solutions that scale**: AWS Step Functions automatically scales the operations and underlying compute to run the steps of your application for you in response to changing workloads. Step Functions scales automatically to help ensure the performance of your application workflow remains consistent as the frequency of requests increases.
- **High volume orchestration for event processing workflows such as IoT and data ingestion**: AWS Step Functions adds Express Workflows. Express Workflows support event rates greater than 100,000 per second, enabling you to build high volume, short duration workflows. Express Workflows can coordinate AWS Lambda function invocations, AWS IoT Rules Engine actions, and Amazon EventBridge events from over a hundred AWS third-party SaaS event sources.
- **Large-scale parallel processing**: Step Functions can iterate over objects such as images, logs, or CSV files stored in Amazon S3, then launch and coordinate thousands of parallel workflows to process the data. You can more quickly and more easily scale thousands of concurrent workflow executions, helping you analyze millions of logs, iterate over terabytes of data or process documents, images, and video files. Learn more about the Step Functions Map state for building large-scale parallel workflows.
- **Security**: AWS Step Functions is integrated with AWS Identity and Access Management (IAM), and recommends a least-privileged IAM policy for all of the resources used in your workflow. Step Functions supports VPC Endpoints (VPCE) using AWS PrivateLink. you You can access AWS Step Functions from VPC-enabled AWS Lambda functions and other AWS services without traversing the public internet. Learn more about Step Functions security and authentication.
- **Compliance**: AWS Step Functions is a HIPAA eligible service, and can be used with applications containing healthcare-related information such as personal health information (PHI). Step Functions is also compliant with SOC (System & Organization Control) measures, and the results of these third-party audits are available on the AWS SOC Compliance site. AWS Step Functions also conforms to other common compliance standards. Understand the compliance scope by reviewing the AWS Cloud Security site.
- **Pay per use**: With AWS Step Functions, you pay for each transition from one state to the next. Billing is metered by state transition, and you do not pay for idle time, regardless of how long each state persists (up to one year). This keeps Step Functions cost-effective as you scale from a few executions to tens of millions.

172.1.2 Benefits

- **Quick Onboarding**: Start building quickly with Workflow Studio, a simple drag-and-drop interface for expressing complex business logic.

- **Simple Automation**: Automate workflows across over 220 AWS services without maintaining code.
- **Process data on demand**: Use code to process data on demand with large-scale parallel workflows
- **Event-driven architecture visualization**: Visualize and develop resilient workflows for event-driven architectures.

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 14 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/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/

172.5 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.

173 AWS Storage Gateway

173.1 Service Overview

AWS Storage Gateway is a hybrid cloud storage service that gives you on-premises access to virtually unlimited cloud storage. Storage Gateway provides a standard set of storage protocols such as iSCSI, SMB, and NFS, which allow you to use AWS storage without rewriting your existing applications. It provides low-latency performance by caching frequently accessed data on premises, while storing data securely and durably in Amazon cloud storage services.

173.1.1 Features

- **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.

- **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 and access controls, and supplies compliances and certifications that address enterprise customers' real and perceived security concerns when using AWS 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**. 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; and software errors, such as connection timeouts and file share or volume unavailability.
- **Amazon S3 File Gateway**: Amazon S3 File Gateway presents a file interface 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.
- **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.

- **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.
- **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.
- **Storage Gateway deployment options**: As a hybrid cloud service, AWS Storage Gateway is fully managed and consists of both cloud and on-premises components, which can be deployed in several methods based on your on-premises infrastructure needs: as a virtual machine, which can run on VMware ESXi, Microsoft Hyper-V, or Linux KVM on premises; as a hardware appliance on-premises; as a VM in VMware Cloud on AWS; or as an AMI in Amazon EC2.

173.1.2 Benefits

- **Low-latency**: Deliver low-latency data access to on-premises applications while using the agility, economics and security capabilities of AWS in the cloud.
- **On-premises applications access**: Provide on-premises applications access to cloud-backed storage without disruption to your business by maintaining user and application workflows.
- **Unlimited cloud storage**: Offer virtually unlimited cloud storage to users and applications without deploying new storage hardware.
- **Security and compliance**: Support your compliance efforts with key capabilities like encryption, audit logging, and write-once, read-many (WORM) storage.

173.2 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.

173.3 Pricing Overview

Please see the AWS UK G Cloud 14 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/storagegateway/

- **Service quotas:** https://docs.aws.amazon.com/filegateway/latest/files3/fgw-quotas.html
- **Service FAQs:** https://aws.amazon.com/storagegateway/faqs/

173.5 Technical Requirements

Please refer to https://docs.aws.amazon.com/storagegateway/ 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.

174 AWS Supply Chain

174.1 Service Overview

AWS Supply Chain unifies data and provides machine learning–powered actionable insights, built-in contextual collaboration, and demand planning.

174.1.1 Features

- **Data lakes:** Supply Chain sets up a data lake using ML models for supply chains to understand, extract, and transform disparate, incompatible data into a unified data model. The data lake can ingest your data from various data sources, including your existing ERP systems, such as SAP S/4HANA, and supply chain management systems. To add data from variable sources such as EDI 856, Supply Chain uses ML and natural language processing (NLP) to associate data from source systems to the unified data model. EDI 850 and 860 messages are transformed directly with predefined but customizable transformation recipes. You can also load data from other systems to an Amazon Simple Storage Service (Amazon S3) bucket, where it will be automatically ingested into the AWS Supply Chain Data Lake.
- **Real-time visual map:** Supply Chain contextualizes your data in a real-time visual map using a set of interactive, visual end-user interfaces built on a micro frontend (MFE) architecture. Supply Chain then highlights current inventory selection and quantity, as well as the health of inventory at each location (for example, inventory that is at risk for stock out). Inventory managers can drill down into specific facilities and view the current inventory on hand, in transit, and potentially at risk in each location.
- **Insights:** Supply Chain automatically generates insights to potential supply chain risks (for example, overstock or stock-outs) using the comprehensive supply chain data in the data lake and surfaces them in the real-time visual map. Supply Chain

also offers work order insights to provide visibility of maintenance-related materials from sourcing to delivery, providing order status, identifying delivery risks, and providing delivery risk-mitigation options.

- **Recommended actions and collaboration:** Supply Chain automatically evaluates, ranks, and shares various rebalancing options to provide inventory managers and planners with recommended actions to take if a risk is detected. Recommendation options are scored by the percentage of risk resolved, the distance between facilities, and the sustainability impact. Supply chain managers can also drill down to review the impact that each option will have on other distribution centers across the network. Supply Chain also continually learns from the decisions that you make to improve recommendations over time.
- **Demand planning:** AWS Supply Chain Demand Planning generates more accurate demand forecasts, adjusts to market conditions, and empowers demand planners to collaborate across teams to help avoid excess inventory costs and waste. To help remove the manual effort and guesswork around demand planning, Supply Chain uses ML to analyze historical sales data and real-time data (for example, open orders), create forecasts, and continually adjust models to improve accuracy. Supply Chain Demand Planning also continually learns from changing demand patterns and user inputs to offer near real-time forecast updates, allowing companies to proactively adjust supply chain operations.
- **Supply planning:** AWS Supply Chain Supply Planning forecasts and plans purchases of raw materials, components, and finished goods. This capability draws on nearly 30 years of Amazon experience in developing and honing AI/ML supply planning models, and considers economic factors, such as holding and liquidation costs. Supply Chain Supply Planning uses the comprehensive, standardized data from the Supply Chain Data Lake, including demand forecasts generated by Supply Chain Demand Planning (or any other demand planning system). Your organization benefits from improved service levels and reduced inventory costs by being able to better respond to variations in demand and disruptions in supply. Manufacturing customers can create supply plans for components and finished products across multiple levels in their bill of materials and improve in-stock and order fill rates by dynamically calculating inventory targets, while accounting for demand variability, actual vendor lead times, and ordering frequency.
- **N-Tier Visibility:** AWS Supply Chain N-Tier Visibility extends visibility and insights beyond your organization to your external trading partners. This visibility lets you align and confirm orders with suppliers, improving the accuracy of planning and execution processes. Invite, onboard, and collaborate with your trading partners in just a few steps to confirm supply plans and obtain order commitments. Commitments and confirmations are received from partners and written to the Supply Chain Data Lake. This data can then be used to identify material or component shortages and update supply plans with new information and provide better-informed insights.
- **Sustainability:** AWS Supply Chain Sustainability uses the same underlying technology as N-Tier Visibility to provide a more secure and efficient way for sustainability professionals to obtain the documents and datasets they need from their supplier network. These capabilities help you provide environmental and social governance (ESG) information based on a single, auditable record of the data.

174.1.2 Benefits

- **Mitigate overstock and stock out risks:** Mitigate overstock and stock-out risks to improve customer experiences while lowering excess inventory costs.

- **Quickly gain visibility:** Quickly gain visibility across your supply chain without replatforming, upfront licensing fees, or long-term commitments.
- **ML-powered actionable insights:** Make more informed supply chain decisions with machine learning (ML)–powered actionable insights.
- **Streamline collaboration on supply plans and collecting sustainability data:** More securely and easily collaborate with partners on supply plans and order commitments. Identify and mitigate material or component shortages and efficiently collect sustainability data.

174.2 Backup/Restore and Disaster Recovery

This service allows for backup and recovery. The service can reload data from source systems and demand plans. Users control this via a manual or scheduled API call. Users schedule and recover backups through a web interface.

174.3 Pricing Overview

Please see the AWS UK G Cloud 14 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/aws-supply-chain/
- **Service quotas:** https://docs.aws.amazon.com/aws-supply-chain/latest/adminguide/quotas.html
- **Service FAQs:** https://aws.amazon.com/aws-supply-chain/ SA please confirm there are no service FAQS

174.5 Technical Requirements

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

- Administrator Guide**:** Describes how to create instances, add users and groups, choose an administrator, and log into the AWS Supply Chain web application.
- User Guide: Describes how to log into AWS Supply Chain web application, connect to your data sources, create an insight watchlist for inventory risk or lead time deviations, and create forecasts.
- API Reference: Describes the API operations for AWS Supply Chain.

175 AWS Systems Manager

175.1 Service Overview

AWS Systems Manager allows you to centralize operational data from multiple AWS services and automate tasks across your resources on AWS and in multicloud and hybrid environments. 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. You can also take action on each resource group depending on your operational needs. Systems Manager provides a central place to view and manage

your resources on AWS and in multicloud and hybrid environments, so you can have complete visibility and control over your operations.

175.1.1 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.
- **Connect with ITSM and 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.
- **OpsCenter**: OpsCenter provides a central location where operations engineers and IT professionals can view, investigate, and resolve operational issues related to your resources on AWS and in multicloud and hybrid environments. OpsCenter aggregates and standardizes operational issues, referred to as OpsItems, while providing contextually relevant data that helps with diagnosis and remediation. Engineers working on an OpsItem get access to information such as: Event, resource and account details, past OpsItems with similar characteristics, related AWS Config changes, AWS CloudTrail logs, Amazon CloudWatch alarms, stack information, quick-links to access logs and metrics, list of runbooks and recommended runbooks, other information passed to OpsCenter through AWS services.
- **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.
- **Application Manager**: AWS Systems Manager Application Manager helps you investigate and remediate issues with your resources on AWS and in multicloud and hybrid environments in the context of your applications.
- **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.
- **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.
- **Change Manager**: 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**: With Systems Manager Automation, you can author custom runbooks with a low-code visual designer, or choose from over 370 predefined runbooks provided by AWS. You can run Python or PowerShell scripts as part of a runbook in combination with other automation actions such as approvals, AWS API calls, or running commands on your EC2 instances.
- **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**: 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.
- **Compliance**: 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.
- **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**: 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 runbooks 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**: 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.

175.1.2 Benefits

- **Visibility**: Improve visibility and control in the cloud, on premises, and at the edge.
- **Operations**: Shorten the time to detect and resolve operational issues.
- **Compliance**: Maintain instance compliance against your patch, configuration, and custom policies.
- **Automation**: Automate configuration and ongoing management of your applications and resources.

175.2 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.

175.3 Pricing Overview

Please see the AWS UK G Cloud 14 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/systems-manager/
- **Service quotas:** https://docs.aws.amazon.com/general/latest/gr/ssm.html
- **Service FAQs:** https://aws.amazon.com/systems-manager/faq/

175.5 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.

176 AWS Transfer Family

176.1 Service Overview

AWS Transfer Family offers fully managed support for the transfer of files over SFTP, AS2, FTPS, and FTP directly into and out of Amazon S3 or Amazon EFS. You can seamlessly migrate, automate, and monitor your file transfer workflows by maintaining existing client-side configurations for authentication, access, and firewalls — so nothing changes for your customers, partners, and internal teams, or their applications.

176.1.1 Features

- **Fully managed and highly available infrastructure**: The AWS Transfer Family offers fully managed support for the transfer of files over SFTP, FTPS, FTP, and AS2 directly into and out of AWS Storage services. You can migrate, automate, and monitor your file transfer workflows by maintaining existing client-side configurations for authentication, access, and firewalls — so nothing changes for your customers, partners, and internal teams, or their applications.
- **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 for your endpoints, 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.
- **Simple user experience**: An intuitive user interface and API makes it simple for you to configure your SFTP, FTPS, FTP, or AS2 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. You can also setup connectors to access files stored on remote servers hosted by your business partners or your on-premise locations.

- **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.

176.1.2 Benefits

- **Migration**: Easily migrate your file transfer workflows without impacting your existing connections and integrations.
- **Storage**: Store information in Amazon S3 or Amazon EFS, manage workflows, and trigger automated, event-driven tasks with a fully-managed, low-code service.
- **Scalability**: Support thousands of concurrent users with access controls and quickly scale your business-to-business (B2B) file transfers for each line-of-business user.
- **Security**: Meet your security requirements with data encryption, VPC and FIPS endpoints, compliance certifications, PGP, and more.

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 14 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/transfer/
- **Service quotas:** https://docs.aws.amazon.com/general/latest/gr/transfer-service.html
- **Service FAQs:** https://aws.amazon.com/aws-transfer-family/faqs/

176.5 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.
- AWS Transfer Family API reference: Describes all the API operations for AWS Transfer Family in detail. Also provides sample requests, responses, and errors for the supported web service protocols.
- AWS Transfer Family AWS CLI reference: Describes the AWS CLI operations for AWS Transfer Family in detail.

177 AWS Transit Gateway

177.1 Service Overview

AWS Transit Gateway helps you design and implement networks at scale by acting as a cloud router. As your network grows, the complexity of managing incremental connections can slow you down. AWS Transit Gateway connects VPCs and on-premises networks through a central hub.

177.1.1 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 other 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 over 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.
- **Single management portal across cloud and on-premises networks**: Manage your private network that spans the cloud and your premises, from a single pane of glass on the AWS Management Console.
- **Events**: Get notified of network changes, routing changes, and connection status updates.
- **Metrics**: Monitor your global network through performance and traffic metrics, such as bytes in/out, packets in/out, and packets dropped.
- **SD-WAN Network Manager compatibility**: Compatible SD-WAN partners of AWS, such as Cisco, Aruba, Silver Peak, Aviatrix, and Versa have pre-configured AWS Site-to-Site VPNs so that SD-WAN solutions can automate the connection of your remote sites with AWS. With AWS Transit Gateway Network Manager, you get a unified view of your network across AWS and on-premises networks.

177.1.2 Benefits

- **Architecture**: Streamline your architecture to manage it as it grows in complexity.
- **Visibility**: Get better visibility and control over your virtual private clouds and edge connections.
- **Security**: Improve security with inter-Region peering encryption on the AWS global private network.
- **Migration**: Lift and shift your on-premises finance and video multicast applications to the cloud.

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 14 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/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/

177.5 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.

178 AWS Trusted Advisor

178.1 Service Overview

AWS Trusted Advisor helps you optimize costs, increase performance, improve security and resilience, and operate at scale in the cloud. Trusted Advisor continuously evaluates your AWS environment using best practice checks across the categories of cost optimization, performance, resilience, security, operational excellence, and service limits, and recommend actions to remediate any deviations from best practices.

178.1.1 Features

- **Access all checks**: You can access Trusted Advisor in the AWS Management Console at at https://console.aws.amazon.com/trustedadvisor/home. If you have a Basic or Developer Support plan, you can use the Trusted Advisor console to access all checks in the Service Limits category and six checks in the Security category. If you have a Business, Enterprise On-Ramp, or Enterprise Support plan, you can use the Trusted Advisor console and the AWS Trusted Advisor API to access all Trusted Advisor checks.
- **Monitor status**: Use Amazon CloudWatch Events to monitor the status of Trusted Advisor checks. For more information, see Monitoring AWS Trusted Advisor check results with Amazon EventBridge.
- **Recommendations to fix issues**:You can use the Trusted Advisor Recommendations page of the Trusted Advisor console to review check results for your AWS account and then follow the recommended steps to fix any issues. For example, Trusted Advisor might recommend that you delete unused resources to reduce your monthly bill, such as an Amazon Elastic Compute Cloud (Amazon EC2) instance.

- **Perform operations**: Use the AWS Trusted Advisor API to perform operations on your Trusted Advisor checks.

178.1.2 Benefits

- **Align with AWS best practices**: Identify deviations from AWS best practices and get recommended actions to remediate.
- **Prioritize important recommendations**: Prioritized recommendations from your AWS account team, based on your business priorities, critical applications, and urgency of recommendation, available to Enterprise Support customers.
- **Streamline collaboration across your organization**: Get better alignment in your teams through greater visibility, monitoring, and tracking of prioritized recommendations, available to Enterprise Support customers.
- **Optimize your AWS resources at scale**: Get an aggregated view of recommendations across your organization or integrate programmatically with Trusted Advisor APIs.

178.2 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/.

178.3 Pricing Overview

Please see the AWS UK G Cloud 14 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/awssupport/latest/user/trusted-advisor.html
- **Service FAQs:** https://aws.amazon.com/premiumsupport/faqs/

178.5 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

179 AWS Verified Access

179.1 Service Overview

AWS Verified Access verifies each access request in real time and only connects users to the applications that they are allowed to access. This removes broad access to corporate applications, reducing the associated risks. To verify users against specific security requirements, Verified Access integrates with AWS and third-party security services to source information about identity, device security status, and location. IT administrators can use Verified Access to author a set of policies that defines a user's ability to access each application.

Verified Access also simplifies security operations by helping administrators efficiently set and monitor access policies, freeing time to update policies, respond to security and connectivity incidents, and audit for compliance standards.

179.1.1 Features

- **Security:** Using Verified Access, you can configure fine-grained access for your applications, ensuring that application access is granted only when users meet the specified security requirements (for example, user identity and device security status). Built on Zero Trust guiding principles, Verified Access validates every application request before granting access. Verified Access also supports AWS WAF, helping you filter out common threats such as SQL injection and cross-site scripting (XSS).
- **User authentication:** Verified Access is seamlessly integrated with AWS IAM Identity Center, which allows end users to authenticate with SAML-based third-party identity providers (IdPs). If you already have a custom IdP solution that is OpenID Connect compatible, Verified Access can also authenticate users by directly connecting with your IdP.
- **Device posture assessment:** Verified Access is integrated with third-party device management services to provide additional security context. Therefore, you can additionally assess access attempts using the security and compliance state of the user's device.
- **Personalize apps:** Verified Access passes signed identity context, such as user alias, to your applications. This helps you personalize your applications using this context, removing the need to reauthenticate the user at your application. The signed context also protects your applications in case Verified Access is accidentally disabled, as the application can reject the request if it doesn't receive the context.
- **Group applications:** With Verified Access, you can group applications with similar security needs. Each application within a group shares a global policy, achieving a minimum-security bar for the entire group and removing the need to manage individual policies for each application. For example, you can group all "dev" applications and set a group-wide access policy.
- **Access logs:** Verified Access logs every access attempt so that you can quickly respond to security incidents and audit requests. Verified Access supports delivery of these logs to Amazon Simple Storage Service (Amazon S3), Amazon CloudWatch Logs, and Amazon Kinesis Data Firehose. Verified Access supports the Open Cybersecurity Schema Framework (OCSF) logging format, making it easier for you to analyze logs using one of the supported security information and event management (SIEM) and observability providers.

179.1.2 Benefits

- **Secure distributed users:** Verified Access evaluates each request in real time against predefined security requirements to facilitate secure access to applications.
- **Manage corporate application access:** IT administrators can build a set of fine-grained access policies using security signal input like user identity and device security status.
- **Use centralized access logs:** Verified Access evaluates access requests and logs request data, accelerating analysis of and response to security and connectivity incidents.

Use or disclosure of data contained on this sheet is subject to the restriction on the title page of this document.

179.2 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/.

179.3 Pricing Overview

Please see the AWS UK G Cloud 14 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/verified-access/
- **Service quotas:** https://docs.aws.amazon.com/verifiedpermissions/latest/userguide/quotas.html
- **Service FAQs:** https://aws.amazon.com/verified-access/

179.5 Technical Requirements

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

- User Guide: Describes key concepts for Verified Access and provides instructions for using the service.
- API Reference: Describes the API operations for Verified Access.
- EC2 section of the AWS CLI Reference: Describes the AWS CLI commands for Verified Access.

180 AWS VPN

180.1 Service Overview

AWS VPN is comprised of two services: AWS Site-to-Site VPN and AWS Client VPN. AWS Site-to-Site VPN enables you to securely connect your on-premises network or branch office site to your Amazon Virtual Private Cloud (Amazon VPC). AWS Client VPN enables you to securely connect users to AWS or on-premises networks.

180.1.1 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.
- **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). 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.

- **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.
- **Private IP VPN**: Private IP VPN provides the ability to deploy Site-to-site VPN connections over Direct Connect (DX) using private IP addresses. With this feature, you can encrypt DX traffic between your on-premises network and AWS without the need for public IP addresses, thus enabling enhanced security and network privacy at the same time. Private IP VPN can be deployed using AWS Transit Gateway which allows centralized management of customer's AWS Virtual Private Clouds (VPC) and connections to your on-premises networks in a more secured, private and scalable manner.
- **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.
- **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.
- **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.
- **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 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.

180.1.2 Benefits

- **Highly available**: AWS Site-to-Site VPN delivers high availability by using two tunnels across multiple Availability Zones within the AWS global network. You can

stream primary traffic through the first tunnel and use the second tunnel for redundancy if one tunnel goes down, traffic continues to flow.

- **Secure**: With AWS Site-to-Site VPN, you can connect to an Amazon VPC or AWS Transit Gateway the same way you connect to your on-premises servers. AWS Site-to-Site VPN establishes secure and private sessions using IP Security (IPSec).
- **Accelerate applications**: The Accelerated Site-to-Site VPN option improves the performance of your VPN connection by working with AWS Global Accelerator. AWS Global Accelerator is used to intelligently route traffic to the nearest AWS network endpoint with the best performance.
- **Robust monitoring**: AWS Site-to-Site VPN gives you visibility into local and remote network health, and monitors the reliability and performance of your VPN connections by integrating with Amazon CloudWatch.
- **Advanced authentication**: Many organizations require multi-factor authentication (MFA) and federated authentication from their VPN solution. AWS Client VPN supports these and other authentication methods.
- **Elastic**: Traditional on-premises VPN services are limited by the capacity of the hardware that runs them. AWS Client VPN is a pay-as-you-go cloud VPN service that elastically scales up or down based on user demand.
- **Remote access**: Unlike on-premises VPN services, AWS Client VPN allows users to connect to AWS and on-premises networks using a single VPN connection.
- **Fully managed**: AWS Client VPN automatically takes care of deployment, capacity provisioning, and service updates — while you monitor all connections from a single console.

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 14 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/vpn/
- **Service quotas:** https://docs.aws.amazon.com/vpn/latest/clientvpn-admin/limits.html and https://docs.aws.amazon.com/vpn/latest/s2svpn/vpn-limits.html
- **Service FAQs:** https://aws.amazon.com/vpn/faqs/

180.5 Technical Requirements

Please refer to https://docs.aws.amazon.com/vpn/ 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.
- API Reference: Describes the API operations for 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.

181 AWS WAF

181.1 Service Overview

AWS WAF is a web application firewall that helps protect web applications from attacks by allowing you to configure rules that allow, block, or monitor (count) web requests based on conditions that you define. These conditions include IP addresses, HTTP headers, HTTP body, URI strings, SQL injection and cross-site scripting.

181.1.1 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. AWS WAF allows you to create a centralized set of rules that you can deploy across multiple websites.
- **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. 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.
- **Account takeover fraud prevention**: AWS WAF Fraud Control - Account Takeover Prevention is a managed rule group that monitors your application's login page for unauthorized access to user accounts using compromised credentials. You can use the rule group to help protect against credential stuffing attacks, brute force login attempts, and other anomalous login activities. With optional JavaScript and iOS/Android SDKs, you can receive additional telemetry on user devices that attempt to log in to your application to better protect your application against automated login attempts by bots. Account Takeover Prevention is part of Managed Rules for AWS and can be used together with Bot Control to effectively defend your application against bot attacks.
- **Account creation fraud prevention**: Account Creation Fraud Prevention is a managed rule group that monitors your application's sign-up or registration page for creation of fake or fraudulent accounts. You can use the rule group to help protect against abuse such as promotional or sign-up abuse, loyalty or rewards abuse and phishing. With the recommended JavaScript and iOS/Android SDKs, you can receive additional telemetry on user devices that attempt to sign-up to your application to better protect your application against automated attempts by bots.
- **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.
- **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.

181.1.2 Benefits

- **Save time**: Save time with managed rules so you can spend more time building applications.
- **Monitor**: More easily monitor, block, or rate-limit common and pervasive bots.
- **Enhance visibility**: Improve web traffic visibility with granular control over how metrics are emitted.

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 14 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/waf/index.html
- **Service quotas:** https://docs.aws.amazon.com/waf/latest/developerguide/limits.html
- **Service FAQs:** https://aws.amazon.com/waf/faqs/

181.5 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.

182 AWS Wavelength

182.1 Service Overview

AWS Wavelength combines the high bandwidth and ultralow latency of 5G networks with AWS compute and storage services so that developers can innovate and build a new class of applications. Wavelength Zones are available in co-operation with Verizon in the United

States, KDDI in Japan, SK Telecom in South Korea, Vodafone in UK and Germany, and Bell in Canada.

182.1.1 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.

182.1.2 Benefits

- **Easy to use**: Build next-generation applications without any learning curve using familiar AWS services, APIs, and tools.
- **Scalability**: Develop applications once and scale deployments to multiple Wavelength Zones across global 5G networks.
- **Innovation**: Use proven AWS infrastructure and services to accelerate innovative 5G edge application development.

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

Please see the AWS UK G Cloud 14 Pricing Document accompanying this service in the Digital Marketplace.

182.4 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/

182.5 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.

183 AWS Well-Architected Tool

183.1 Service Overview

The AWS Well-Architected Tool lets you review your workloads against current AWS best practices and obtain advice on how to architect your workloads for the cloud. This tool uses the AWS Well-Architected Framework.

183.1.1 Features

- **A wide variety of best practices**: The AWS Well-Architected Tool features a Lens Catalog, which provides you with access to expert-authored review lenses from AWS – above and beyond what's in the AWS Well-Architected Framework. The Lens Catalog is continuously refreshed with the latest technology and industry-specific lenses, so you can dive deep into the areas that are most important to your business, as well as explore best practices that others within your industry are using.
- **Customized reviews**: The AWS Well-Architected Tool offers numerous features which support customized reviews, including custom lenses, profiles, and review templates. Each of these features let you focus on what's most important to your business and also save you time. 1) Custom lenses allow you to add your own pillars, questions, best practices, and improvement plans based on your business needs. 2) Profiles let you pre-define business goals you'd like to accomplish when performing a review. When a profile is applied to a workload, the AWS Well-Architected Tool automatically generates a prioritized list of questions that are most relevant in helping you meet your defined goals. 3) Review templates reduce the need to manually fill in the same answers for best practices that are common across multiple workloads when performing a review, driving consistency and standardization of best practices across teams and workloads.
- **Improved collaboration**: With the AWS Well-Architected Tool, you can easily share your workloads and custom lenses with up to 300 individual IAM users or accounts, or across an entire organization through integration with AWS Organizations, enhancing collaboration amongst teams.
- **Easier discovery**: Integrations with other AWS Services like AWS Trusted Advisor and AWS Service Catalog AppRegistry help you more easily discover the information needed to answer Well-Architected review questions and shorten your review time.

- **In-workflow best practices**: The AWS Well-Architected Tool provides APIs that allow you to extend AWS Well-Architected functionality, best practices, measurements, and learnings into your existing architecture governance processes, applications, and workflows.
- **Support for regulatory and compliance needs**: The AWS Well-Architected Tool is also available in the GovCloud Region (US), which means customers with specific regulatory and compliance requirements, as well as AWS partners, can conduct self-service Well-Architected reviews. Additionally, the AWS Well-Architected Tool is compliant with the Federal Risk and Authorization Management Program (FedRAMP).

183.1.2 Benefits

- **Architectural guidance**: Get access to best practices and an action plan to fine-tune your applications and workloads.
- **In-Workflow best practices**: Extend AWS Well-Architected Tool for your architecture governance processes, applications, and workflows through robust APIs.
- **Measure improvements**: Save point-in-time milestones, implement improvements, and track changes to measure progress.
- **Enhance collaboration**: Tailor your review with custom lenses that you can share with up to 300 individual IAM users or across your entire organization.

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 14 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/wellarchitected/
- **Service quotas:** https://docs.aws.amazon.com/general/latest/gr/wellarchitected.html
- **Service FAQs:** https://aws.amazon.com/well-architected-tool/faqs/

183.5 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.

184 AWS X-Ray

184.1 Service Overview

AWS X-Ray helps to analyze and debug production, distributed applications, such as those built using a microservices architecture. With X-Ray, you can understand how your application and its underlying services are performing to identify and troubleshoot the root cause of performance issues and errors. X-Ray provides an end-to-end view of requests as they travel through your application, and shows a map of your application's underlying components. You can use X-Ray to analyze both applications in development and in production, from simple three-tier applications to complex microservices applications consisting of thousands of services.

184.1.1 Features

- **Simple setup**: AWS X-Ray can be used with applications running on Amazon Elastic Compute Cloud (EC2), Amazon EC2 Container Service (Amazon ECS), AWS Lambda, AWS Elastic Beanstalk. You just integrate the X-Ray SDK with your application and install the X-Ray agent. With AWS Elastic Beanstalk, you only have to integrate the X-Ray SDK with your application since the X-Ray agent is pre-installed on Elastic Beanstalk.
- **End-to-end tracing**: AWS X-Ray provides an end-to-end, cross-service view of requests made to your application. It gives you an application-centric view of requests flowing through your application by aggregating the data gathered from individual services in your application into a single unit called a trace. Use this trace to follow the path of an individual request as it passes through each service or tier in your application so that you can pinpoint where issues are occurring.
- **AWS service and database integrations**: AWS X-Ray supports applications running on Amazon Elastic Compute Cloud (Amazon EC2), Amazon EC2 Container Service (Amazon ECS), AWS Lambda, and AWS Elastic Beanstalk. The X-Ray SDK captures metadata for requests made to MySQL and PostgreSQL databases (self-hosted, Amazon RDS, Amazon Aurora), and Amazon DynamoDB. It also captures metadata for requests made to Amazon Simple Queue Service and Amazon Simple Notification Service.
- **Support for multiple languages**: AWS X-Ray supports tracing for applications that are written in Node.js, Java, and .NET.
- **Request sampling**: You can set the trace sampling rate that suits your production applications or applications in development. X-Ray continually traces requests made to your application and stores a sampling of the requests for your analysis.
- **Service map**: AWS X-Ray creates a map of services used by your application with trace data that you can use to drill into specific services or issues. This provides a view of connections between services in your application and aggregated data for each service, including average latency and failure rates. You can create dependency trees, perform cross-availability zone or region call detections, and more.
- **Server- and Client-side latency detection**: AWS X-Ray allows you to visually detect node and edge latency distribution directly from the service map. You can quickly isolate outliers, graph pattern and trends, drill into traces and filter by built-in keys and custom annotations to better understand performance issues impacting your application and end users.
- **Data annotation and filtering**: AWS X-Ray lets you add annotations to data emitted from specific components or services in your application. You can use this to append business-specific metadata that help you better diagnose issues. You can also view

and filter data for traces by properties such as annotation value, average latencies, HTTP response status, timestamp, database table used, and more.

- **Console and programmatic access**: You can use AWS X-Ray with the AWS Management Console, AWS CLI, and AWS SDKs. The X-Ray API lets you programmatically access the service so you can easily export trace data or ingest the data into your own tools and custom analytics dashboards.
- **Security**: AWS X-Ray is integrated with AWS Identity and Access Management (IAM) so that you can control which users and resources have permission to access your traces and how.

184.1.2 Benefits

- **Easy tracking**: Trace user requests through your application while meeting your security and compliance objectives.
- **Improved performance**: Identify bottlenecks and determine where high latencies are occurring to improve application performance.
- **Removing obstacles**: Remove data silos and get the information you need to improve user experience and reduce downtime.
- **Save time and money**: Debug serverless applications in real time, and monitor both cloud cost and performance metrics.

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

184.5 Technical Requirements

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

- Developer Guide: Describes how to get started and instrument your applications using AWS X-Ray.
- API Reference: Describes all API operations for AWS X-Ray in detail. Also provides details of request and response syntax and errors for the supported web services protocols.
- AWS X-Ray SDK for Java API Reference: Describes all of the classes included in the AWS X-Ray SDK for Java.

185 Elastic Load Balancing

185.1 Service Overview

Elastic Load Balancing lets you Secure your applications with SSL/TLS termination, integrated certificate management, and client certificate authentication. Deliver applications with high availability and automatic scaling. And lastly, Deliver applications with high availability and automatic scaling. Elastic Load Balancing supports Application Load Balancers, Network Load Balancers, Gateway Load Balancers, and Classic Load Balancers.

185.1.1 Features

- **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**: 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).
- **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.
- **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).

185.1.2 Benefits

- **Availability**: A load balancer distributes workloads across multiple compute resources, such as virtual servers. Using a load balancer increases the availability and fault tolerance of your applications.
- **Easy to manage**: You can add and remove compute resources from your load balancer as your needs change, without disrupting the overall flow of requests to your applications.
- **Health check**: You can configure health checks, which monitor the health of the compute resources, so that the load balancer sends requests only to the healthy ones. You can also offload the work of encryption and decryption to your load balancer so that your compute resources can focus on their main work.

Use or disclosure of data contained on this sheet is subject to the restriction on the title page of this document.

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 14 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/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/

185.5 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.

186 FreeRTOS

186.1 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. To support a growing number of use cases, AWS provides software libraries that offer enhanced functionality including connectivity, security, and over-the-air updates. For example, you can use FreeRTOS to securely connect your small, low-powered devices to AWS cloud services like AWS IoT Core or to more powerful edge devices running AWS IoT Greengrass.

186.1.1 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.
- **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 behavior. 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. FreeRTOS manages a secure connection to the cloud using Transport Layer Security (TLS v1.2). The TLS library implements an abstraction layer for the TLS protocol which provides privacy and data integrity between two communicating applications.
- **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.
- **Supported Microcontrollers**: The AWS Partner Device Catalog lists devices that fully support FreeRTOS and its features and capabilities. You can easily and quickly find supported development kits as well as off-the-shelf devices such as gateways, edge servers, sensors, and cameras for immediate project integration. It offers development kits and embedded systems to build new devices, and off-the-shelf-devices such as gateways, edge servers, sensors, and cameras for immediate IoT project integration.
- **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 LTS includes the FreeRTOS kernel and IoT libraries – FreeRTOS+TCP, coreMQTT, coreHTTP, corePKCS11, coreJSON, AWS IoT OTA, AWS IoT Jobs, AWS IoT Device Defender, AWS IoT Device Shadow.

- **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* from 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.

186.1.2 Benefits

- **Time to market**: Accelerate your time to market with a real-time operating system implemented in over 40 architectures.
- **Integrations**: Use a reliable and responsive kernel that more easily integrates with cloud services.
- **Easy to update**: Manage and maintain your designs with modular software libraries, and perform over-the-air (OTA) security patches and firmware updates.
- **Support and maintenance**: Reduce your business risk with FreeRTOS Long-Term Support (LTS) and Extended Maintenance Plan (EMP).

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

Please see the AWS UK G Cloud 14 Pricing Document accompanying this service in the Digital Marketplace.

186.4 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/

186.5 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.
- FreeRTOS Porting Guide: Provides detailed information about porting FreeRTOS to a microcontroller platform.
- FreeRTOS Qualification Guide: Provides detailed information about qualifying a microcontroller platform for FreeRTOS.

187 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.

187.1 Availability

AWS has a large global infrastructure footprint, and this footprint is constantly increasing. 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 33 Regions and 105 Availability Zones throughout the world—including 38 Local Zones and 29 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),), Canada (West), Europe (Ireland), Europe (Frankfurt), Europe (London), Europe (Milan), Europe (Paris), Europe (Stockholm), Europe (Spain), Europe (Zurich), Middle East (Bahrain), Middle East (Israel), Middle East (United Arab Emirates [UAE]), Africa (Cape Town), Asia Pacific (Singapore), Asia Pacific (Tokyo), Asia Pacific (Seoul), Asia Pacific (Mumbai), Asia Pacific (Hyderabad), Asia Pacific (Osaka), Asia Pacific (Hong Kong), Asia Pacific (Jakarta), Australia (Melbourne), Australia (Sydney), South America (Sao Paulo), China (Beijing), and China (Ningxia).

AWS has also announced plans for 6 more AWS Regions in Malaysia, New Zealand, the Kingdom of Saudi Arabia, Thailand, and the AWS European Sovereign Cloud.

Information about each Region can be found at the AWS Global Infrastructure page.

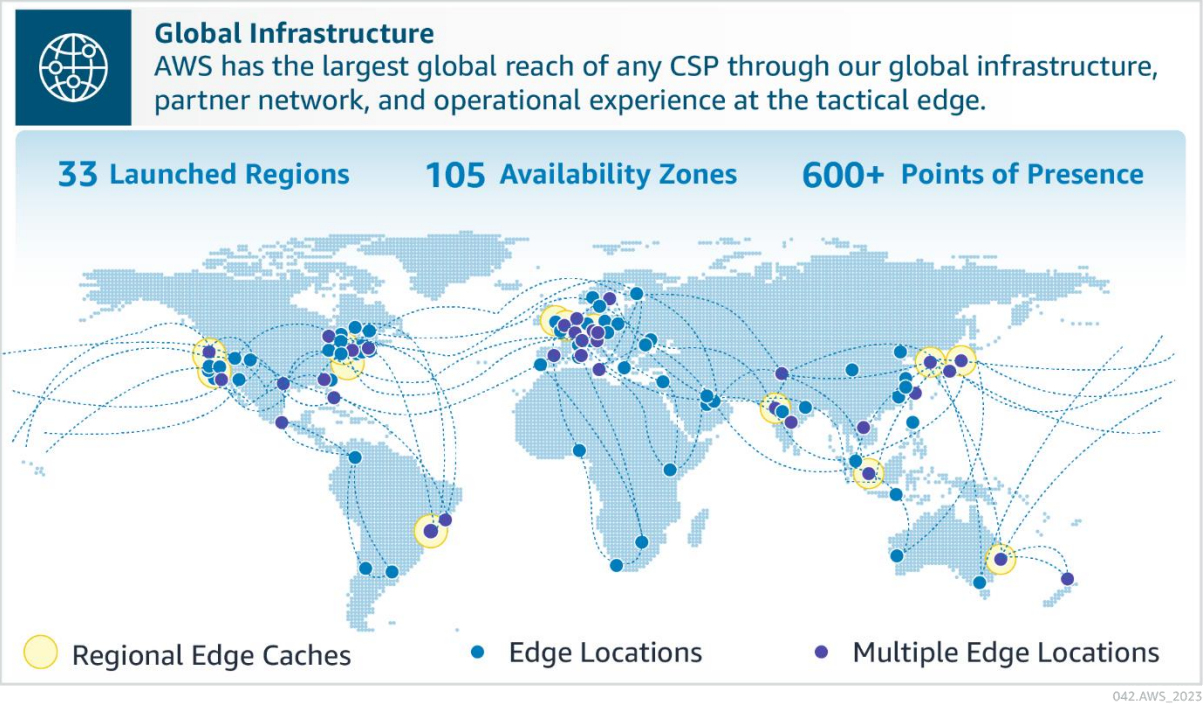

Figure 1. AWS Global Infrastructure. *AWS has a massive global infrastructure footprint and it is constantly increasing. We currently serve 245 countries and territories.*

depicts the current AWS Regions and Availability Zones.

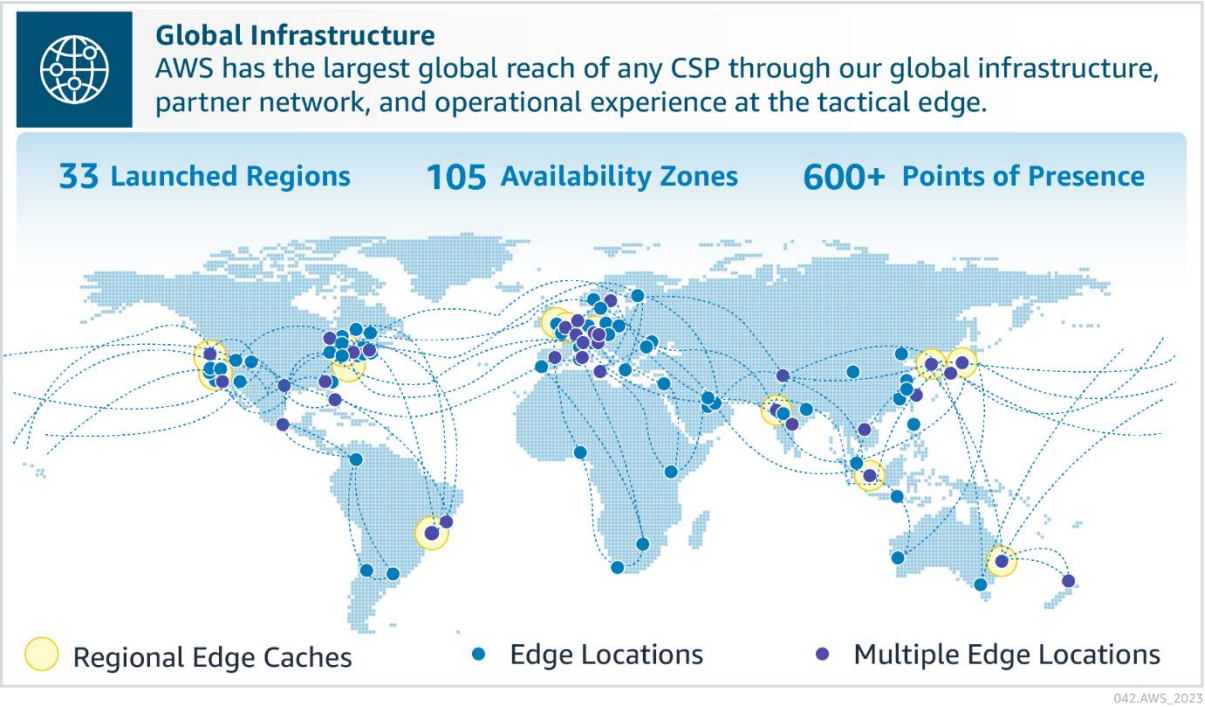

Figure 1. AWS Global Infrastructure. *AWS has a massive global infrastructure footprint and it is constantly increasing. We currently serve 245 countries and territories.*

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.

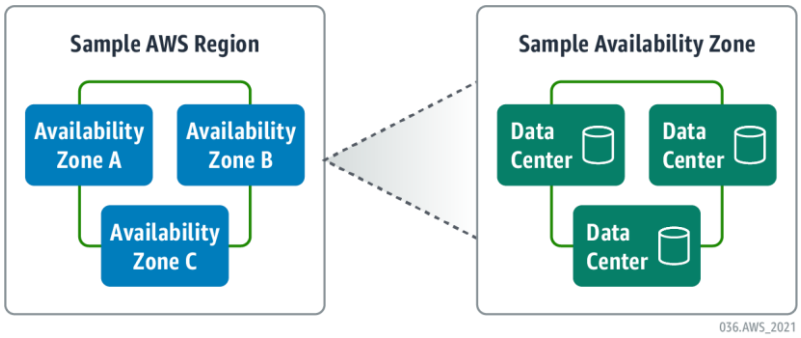

Figure 3. Relationship Between AWS Regions and Availability Zones

187.1.1 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/.

187.1.2 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.

187.1.2.1 Design for Failure

Although rare, failures can occur that affect the availability of resources hosted in a single Availability Zone. In this case, none of your resources would be available. To avoid such a situation, architect across multiple Availability Zones in the same Region to achieve high recovery time objectives, recovery point objectives, and service availability. Availability Zones connect to each other with fast, private fiber-optic networking, so you can architect applications to automatically fail over between zones without interrupting your services or workloads.

For mission-critical applications, it is a good 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 private, high-speed networking or public internet connections. You get an additional layer of business continuity along with low-latency access for services.

187.1.2.2 Automate Failover and Recovery

By monitoring a system for key performance indicators (KPIs), you can trigger automatic actions when a threshold is breached. These KPIs should be a measure of business value and not of the technical aspects of the service operation. Use KPI-based monitoring to automate notifications and track failures, and use automated recovery processes that work around or repair the failure. With sophisticated automation, it is possible to anticipate and remediate failures before they occur.

187.1.2.3 Use AWS Health Dashboard to Monitor and Remediate Operational Events

AWS Health is the authoritative data source for events and changes affecting your AWS Cloud resources. AWS Health provides actionable guidance to help you troubleshoot and remediate operational events, and take actions to prepare for scheduled changes, such as planned lifecycle events. You can view AWS Health information using AWS Health Dashboard.

You can access AWS Health programmatically with Amazon EventBridge and AWS Health API. AWS Health API is available to customers who are on AWS Business, AWS Enterprise On-Ramp, and AWS Enterprise Support plans.

Additionally, Amazon CloudWatch is a service that monitors applications, responds to performance changes, optimizes resource use, and provides insights into operational health. You can create an alarm to monitor any Amazon CloudWatch metric in your account and overlay AWS Health events with CloudWatch metrics to draw more insights when an alarm is activated.

187.1.2.4 Simulate Failures and 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.

187.1.2.5 Use Resources and Reference Architectures

The resources described below help customers to understand AWS 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 offers customers guidance and reference application architectures 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.
- AWS re:Invent: Advanced Continuous Delivery Best Practices video.

187.1.2.6 AWS Well-Architected Framework for Building Reliable Cloud Foundations

Reliability of a workload in the cloud depends on several factors, including design principles such as:

- Automating recovery from failure
- Testing recovery procedures
- Scaling horizontally to increase workload availability
- Automating the addition or removal of resources
- Managing change in automation.

To achieve reliability, you must first evaluate your cloud architectures and implement a consistent approach for resilient and reliable designs. Customers and partners can use the **AWS Well-Architected Framework** to implement design principles that address reliability and resiliency in the cloud.

The AWS Well-Architected Framework is built around six key pillars: operational excellence, security, reliability, performance efficiency, cost optimization, and sustainability. The six pillars together form a framework that guides you through proven architectural practices to design and run workloads in the cloud. The framework codifies the experiences of thousands of customers, helping you assess and improve your cloud-based architectures and mitigate disruptions. The framework includes domain-specific lenses, hands-on labs, and the AWS Well-Architected Tool:

- **Domain-specific lenses** offer guidance to specific industry domains, such as the Government Lens that provides guidance for designing and operating compliant government services on AWS.
- **Hands-on labs** is a collection of workshops and hands-on labs to help you learn, measure, and build using AWS's recommended architectural practices.
- **AWS Well-Architected Tool** helps you review the state of your applications and workloads against proven architectural practices, identify opportunities for improvement, and track progress over time. The AWS Well-Architected Tool is available at zero cost in the AWS Management Console. Customers use the tool to regularly evaluate workloads, identify high-risk issues, and record improvements.

You can use the guidance in the AWS Well-Architected Framework—Reliability Pillar whitepaper to build resilient architectures. The guidance focuses on consistent change

management practices and failure recovery processes that aid resiliency and reliability in the cloud.

187.2 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 learn how migrating on-premises infrastructure to AWS achieves quantifiable business value, refer to The business value of migrating to AWS whitepaper.

187.3 Service Management Details

AWS Cloud services are driven by robust Application Programming Interfaces (APIs). You can integrate AWS Cloud resources with a wide variety of monitoring, management and developer tools. You can use AWS to support tools from familiar vendors such as Microsoft, VMware, BMC Software, Okta, RightScale, Eucalyptus, CA, Xceedium, Symantec, Racemi, and Dell. With this flexibility, you can provision, manage, and monitor all of your IT resources with the tool that best fit your unique needs. Below are various AWS-native management options.

187.3.1 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. You can use the AWS Management Console to perform a number of tasks, from deploying new applications to monitoring the health of your applications. The AWS Management Console helps you manage all aspects of your AWS account, including accessing monthly spending by service, managing security credentials, or setting up new AWS Identity and Access Management (IAM) users. The AWS Management Console supports all AWS Regions so you can provision resources across multiple AWS Regions.

187.3.2 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 API-based services, including Software Development Kits (SDKs), Integrated Development Environment (IDE) toolkits, and Command Line Tools. You can 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. AWS Developer Tools are built to work with AWS, so your team can get set up and start being productive.

We designed AWS Developer Tools to help you build software. They facilitate continuous delivery and infrastructure as code (IaC) for serverless, containers, and Amazon EC2.

187.3.3 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, you 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 a broad partner network 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.

187.4 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:

1. https://aws.amazon.com/api-gateway/sla/
2. https://aws.amazon.com/appstream2/sla/
3. https://aws.amazon.com/athena/sla/
4. https://aws.amazon.com/rds/aurora/sla/
5. https://aws.amazon.com/bedrock/sla/
6. https://aws.amazon.com/braket/sla/
7. https://aws.amazon.com/chime/sla/
8. https://aws.amazon.com/cloudfront/sla/
9. https://aws.amazon.com/cloudsearch/sla/
10. https://aws.amazon.com/cloudwatch/sla/
11. https://aws.amazon.com/codecatalyst/sla/
12. https://aws.amazon.com/codeguru/sla/
13. https://aws.amazon.com/codewhisperer/sla/
14. https://aws.amazon.com/cognito/sla/
15. SLA for Amazon ML Language Services: https://aws.amazon.com/machine-learning/language/sla/. This includes the following services:
    a) Amazon Comprehend
    b) Amazon Comprehend Medical
    c) Amazon Lex
    d) Amazon Polly
    e) Amazon Transcribe
    f) Amazon Transcribe Medical
    g) AWS HealthScribe
    h) Amazon Translate
16. https://aws.amazon.com/compute/sla/
17. https://aws.amazon.com/connect/sla/

18. https://aws.amazon.com/datazone/sla/
19. https://aws.amazon.com/it/detective/sla/
20. https://aws.amazon.com/devops-guru/sla/
21. https://aws.amazon.com/documentdb/sla/
22. https://aws.amazon.com/dynamodb/sla/
23. https://aws.amazon.com/ebs/sla/
24. https://aws.amazon.com/ecs/anywhere/sla/
25. https://aws.amazon.com/efs/sla/
26. https://aws.amazon.com/ecs/sla/
27. https://aws.amazon.com/ecr/sla/
28. https://aws.amazon.com/eks/sla/
29. https://aws.amazon.com/elasticloadbalancing/sla/
30. https://aws.amazon.com/elasticache/sla/
31. https://aws.amazon.com/elasticsearch-service/sla/
32. https://aws.amazon.com/emr/sla/
33. https://aws.amazon.com/eventbridge/sla/
34. https://aws.amazon.com/finspace/sla/
35. https://aws.amazon.com/fsx/sla/ Included Services:
    a) Amazon FSx for Windows File Server
    b) Amazon FSx for Lustre
    c) Amazon FSx for OpenZFS
36. https://aws.amazon.com/guardduty/sla/
37. https://aws.amazon.com/healthlake/sla/
38. https://aws.amazon.com/inspector/sla/
39. https://aws.amazon.com/ivs/sla/
40. https://aws.amazon.com/kendra/sla/
41. https://aws.amazon.com/keyspaces/sla/
42. https://aws.amazon.com/kinesis/sla/
43. https://aws.amazon.com/lightsail/sla-lightsail-instances-and-block-storage/
44. https://aws.amazon.com/lightsail/sla-lightsail-managed-databases/
45. https://aws.amazon.com/location/sla/
46. https://aws.amazon.com/lookout-for-equipment/sla/
47. https://aws.amazon.com/lookout-for-vision/sla/
48. https://aws.amazon.com/machine-learning/language/sla/
49. https://aws.amazon.com/macie/sla/
50. https://aws.amazon.com/managed-blockchain/sla/
51. https://aws.amazon.com/managed-service-apache-flink/sla/
52. https://aws.amazon.com/grafana/sla/
53. https://aws.amazon.com/prometheus/sla/
54. https://aws.amazon.com/msk/sla/
55. https://aws.amazon.com/managed-workflows-for-apache-airflow/sla/
56. https://aws.amazon.com/memorydb/sla/
57. https://aws.amazon.com/messaging/sla/
58. https://aws.amazon.com/monitron/sla/

Use or disclosure of data contained on this sheet is subject to the restriction on the title page of this document.

59. https://aws.amazon.com/amazon-mq/sla/
60. https://aws.amazon.com/neptune/sla/
61. https://aws.amazon.com/opensearch-service/sla/
62. https://aws.amazon.com/pinpoint/sla/
63. https://aws.amazon.com/personalize/sla/
64. https://aws.amazon.com/quicksight/sla/
65. https://aws.amazon.com/rds/proxy/sla/
66. https://aws.amazon.com/redshift/sla/
67. https://aws.amazon.com/rekognition/sla/
68. https://aws.amazon.com/rds/sla/
69. https://aws.amazon.com/route53/sla/
70. https://aws.amazon.com/s3/sla-rtc/
71. https://aws.amazon.com/sagemaker/sla/
72. https://aws.amazon.com/s3/sla/
73. https://aws.amazon.com/swf/sla/
74. https://aws.amazon.com/simpledb/sla/
75. https://aws.amazon.com/textract/sla/
76. https://aws.amazon.com/timestream/sla/
77. https://aws.amazon.com/pinpoint/sla/
78. https://aws.amazon.com/vpc/ipam-sla/
79. https://aws.amazon.com/vpc/sla/
80. https://aws.amazon.com/verified-permissions/sla/
81. https://aws.amazon.com/workdocs/sla/
82. https://aws.amazon.com/workmail/amazon-workmail-service-level-agreement/
83. https://aws.amazon.com/amplify/sla/
84. https://aws.amazon.com/application-migration-service/sla/
85. https://aws.amazon.com/appsync/sla/
86. https://aws.amazon.com/audit-manager/sla/
87. https://aws.amazon.com/backup/sla/
88. https://aws.amazon.com/aws-cost-management/aws-budgets/sla/
89. https://aws.amazon.com/certificate-manager/private-certificate-authority/sla/
90. https://aws.amazon.com/vpn/client-vpn-sla/
91. https://aws.amazon.com/cloud-map/sla/
92. https://aws.amazon.com/cloudhsm/sla/
93. https://aws.amazon.com/cloudtrail/sla/
94. https://aws.amazon.com/codeartifact/sla/
95. https://aws.amazon.com/codebuild/sla/
96. https://aws.amazon.com/codepipeline/sla/
97. https://aws.amazon.com/compute-optimizer/sla/
98. https://aws.amazon.com/config/sla/
99. https://aws.amazon.com/aws-cost-management/aws-cost-explorer/sla/
100. https://aws.amazon.com/data-exchange/sla/
101. https://aws.amazon.com/dms/sla/
102. https://aws.amazon.com/device-farm/sla/

103. https://aws.amazon.com/directconnect/sla/
104. https://aws.amazon.com/directoryservice/sla/
105. https://aws.amazon.com/disaster-recovery/sla/
106. https://aws.amazon.com/mediaconnect/sla/
107. https://aws.amazon.com/mediaconvert/sla/
108. https://aws.amazon.com/medialive/sla/
109. https://aws.amazon.com/mediapackage/sla/
110. https://aws.amazon.com/mediastore/sla/
111. https://aws.amazon.com/mediatailor/sla/
112. https://aws.amazon.com/firewall-manager/sla/
113. https://aws.amazon.com/global-accelerator/sla/
114. https://aws.amazon.com/glue/sla/
115. https://aws.amazon.com/ground-station/sla/
116. https://aws.amazon.com/healthlake/imaging/sla/
117. https://aws.amazon.com/iam/sla/
118. https://aws.amazon.com/transfer/sla/
119. https://aws.amazon.com/iot-analytics/sla/
120. https://aws.amazon.com/iot-core/sla/
121. https://aws.amazon.com/iot-device-defender/sla/
122. https://aws.amazon.com/iot-device-management/sla/
123. https://aws.amazon.com/iot-events/sla/
124. https://aws.amazon.com/greengrass/sla/
125. https://aws.amazon.com/iot-sitewise/sla/
126. https://aws.amazon.com/iot-things-graph/sla/
127. https://aws.amazon.com/iot-twinmaker/sla/
128. https://aws.amazon.com/kms/sla/
129. https://aws.amazon.com/lake-formation/sla/
130. https://aws.amazon.com/lambda/sla/
131. https://aws.amazon.com/mainframe-modernization/sla/
132. https://aws.amazon.com/migration-hub/sla/refactor-spaces/
133. https://aws.amazon.com/network-firewall/sla/
134. https://aws.amazon.com/privatelink/sla/
135. https://aws.amazon.com/resilience-hub/sla/
136. https://aws.amazon.com/secrets-manager/sla/
137. https://aws.amazon.com/security-hub/sla/
138. https://aws.amazon.com/servicecatalog/sla/
139. https://aws.amazon.com/shield/sla/
140. https://aws.amazon.com/simspaceweaver/sla/
141. https://aws.amazon.com/vpn/site-to-site-vpn-sla/
142. https://aws.amazon.com/step-functions/sla/
143. https://aws.amazon.com/systems-manager/sla/
144. https://aws.amazon.com/transit-gateway/sla/
145. https://aws.amazon.com/waf/sla/
146. https://aws.amazon.com/xray/sla/

147. https://aws.amazon.com/elasticache/sla/

See the Supplier Terms document affiliated with this framework catalogue for additional information.

187.5 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.

187.6 Data Restoration/Service Migration

187.6.1 AWS Backup

Protecting your data is an important step towards ensuring that you meet your business and regulatory compliance requirements. Even durable resources are susceptible to threats such as 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 by providing a fully managed, policy-based data protection solution.

AWS Backup is a fully managed service that makes it easy to centralize and automate data protection across on-premises and AWS services. 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 Neptune databases
- Amazon DocumentDB (with MongoDB compatibility) databases
- Amazon Elastic File System (Amazon EFS) file systems
- Amazon FSx for Lustre
- Amazon FSx for Windows File Server
- AWS Storage Gateway volumes
- VMware workloads on-premises and in VMware Cloud™ on AWS

You can use AWS Backup to create backup policies that automate backup schedules and retention management. As illustrated in Figure 4. AWS Backup provides a fully managed, policy-based backup solution, AWS Backup provides a fully managed, policy-based solution, simplifying backup management and enabling you to meet business and regulatory compliance requirements.

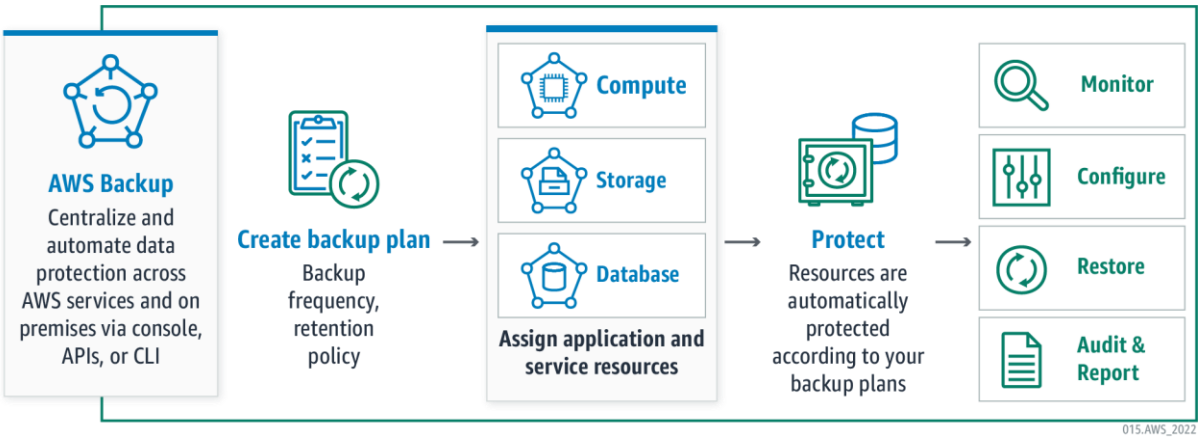

Figure 4. AWS Backup provides a fully managed, policy-based backup solution

187.6.2 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 AWS Elastic 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.

187.6.2.1 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.

187.6.2.2 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 5), and budget. With AWS, customers can build highly resilient applications while taking advantage of flexible, cost-effective infrastructure solutions.

Figure 5. 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 6 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.

[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.

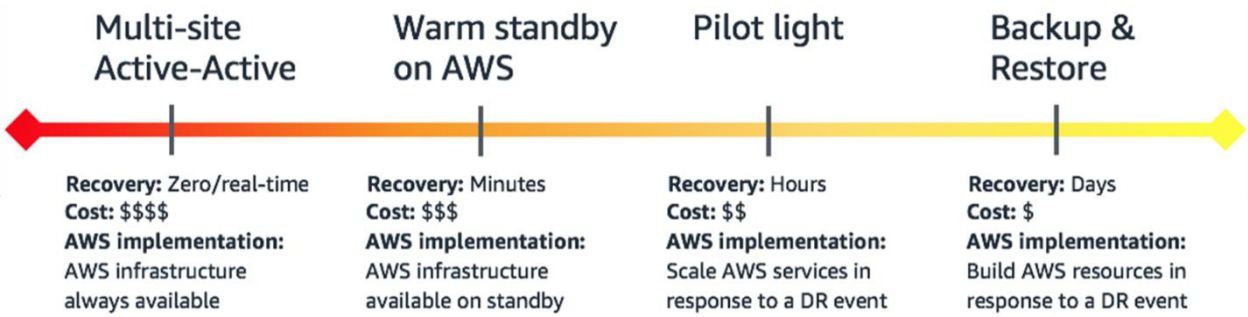

Figure 6. Spectrum of DR Options. *Customers can choose a preferred DR option based on preferences for cost and the time it takes to recover systems.*

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.

187.6.3 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 7 illustrates how AWS Elastic Disaster Recovery operates.

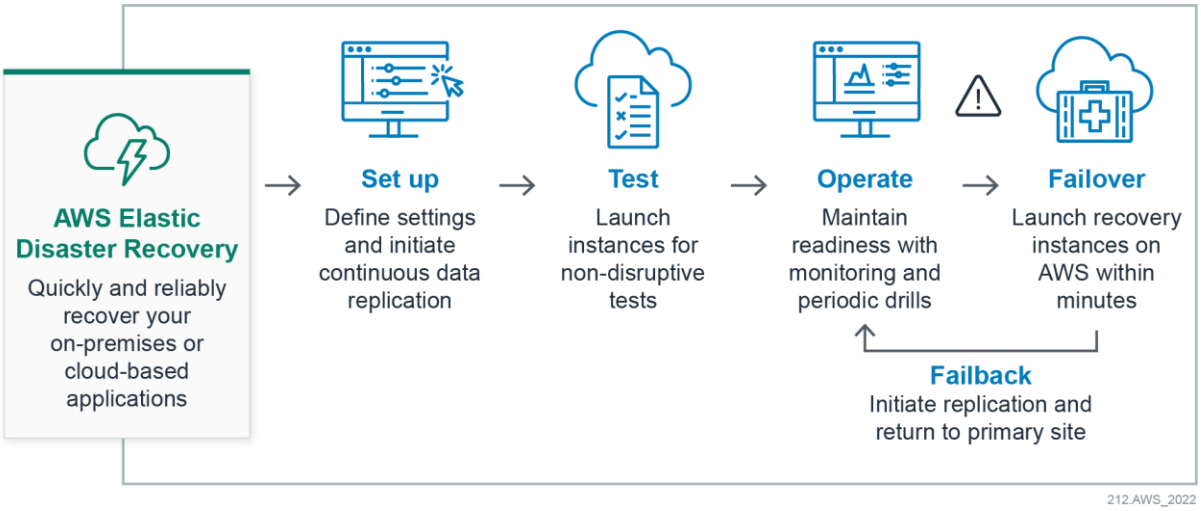

Figure 7. AWS Elastic Disaster Recovery Architecture.*This service simplifies recovery of a wide range of applications on AWS and uses a unified process for drills, recovery, and failback.*

AWS Elastic Disaster Recovery is the recommended service for disaster recovery to AWS. It provides similar capabilities as CloudEndure Disaster Recovery 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.

187.6.3.1 Complementary AWS Cloud Services for Backup/DR

In addition to AWS Elastic 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

187.6.3.2 DR Resources

There are multiple resources to help organizations start using AWS for a DR solution:

- Read the IDC Spotlight: Preparing for Disruption with IT Resilience
- Read the ebook Best practices for implementing disaster recovery in the cloud
- Download the IT Disaster Recovery Plan Checklist to ensure you are on track.
- Learn about affordable, enterprise-grade disaster recovery in our ebook: Top 4 Ways to Reduce Disaster Recovery Costs using the Cloud

187.6.4 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:

- Remote workforce enablement
- Constituent engagement
- Operational continuity
- Real-time analytics
- Process and systems modernization

Error! Reference source not found. 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.

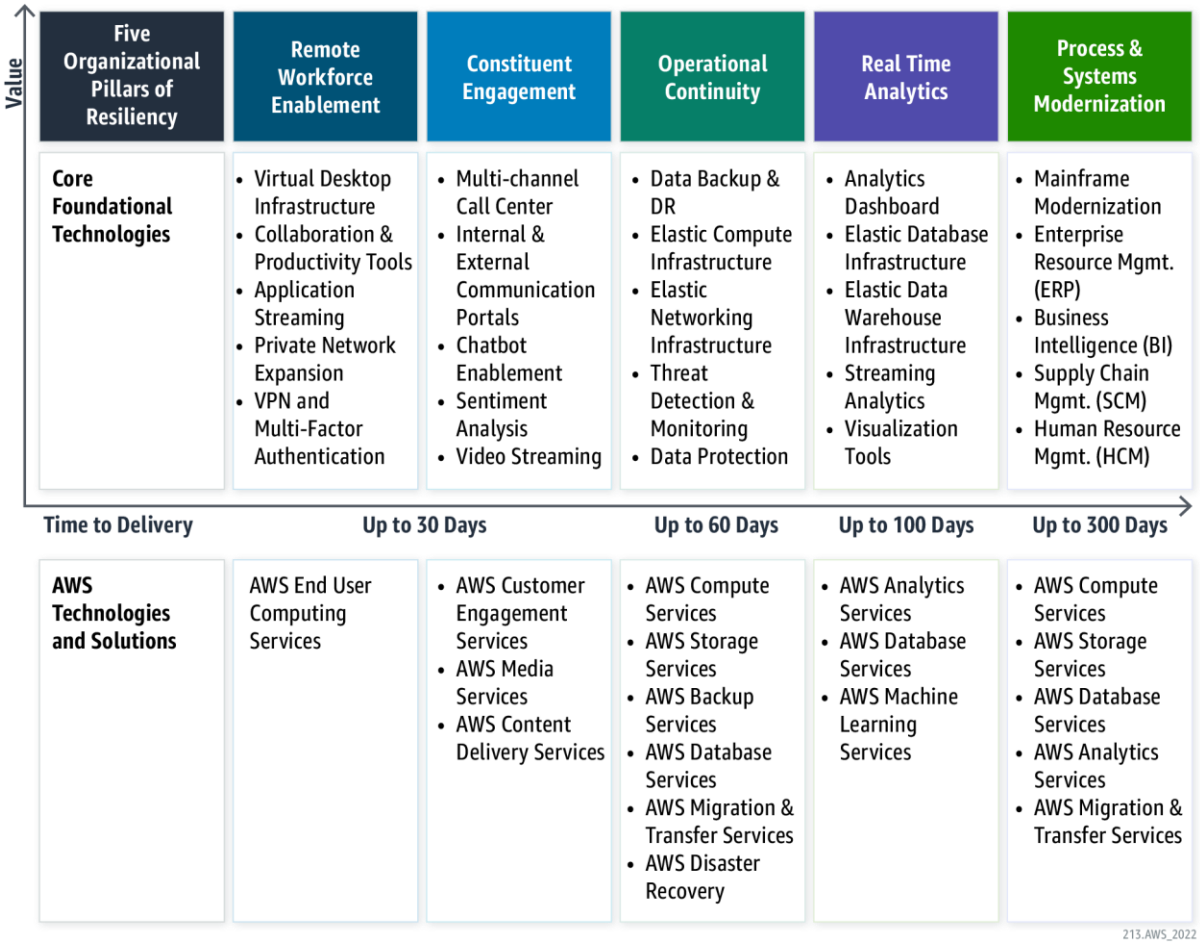

Figure 8. AWS Organizational Resiliency Framework. *This framework can improve an organization's odds of minimizing the impact of an emergency on employees, customers, and partners.*

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.

187.6.4.1 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.

187.6.5 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.

- Cost reduction
- Agility and staff productivity
- Improved security and operational resilience
- Outsourcing and hardware/software end-of-life
- Data center consolidation
- Digital transformation
- Going global quickly, mergers and acquisitions (M&A)
- New technologies (such as AI/ML and IoT).

Independent studies have shown that organizations that migrated to AWS from on premises saw an average of 31% savings in infrastructure,* 62% more efficient IT infrastructure management,* 3x more features delivered each year,* 69% reduction in unplanned downtime,** and 43% fewer security incidents per year.** Organizations can use AWS at every step along the way, to build organizational, operational, and technical capabilities.

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.

187.7 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. Seven Common Migration depicts the seven most common migration patterns that help you create a modern migration strategy.

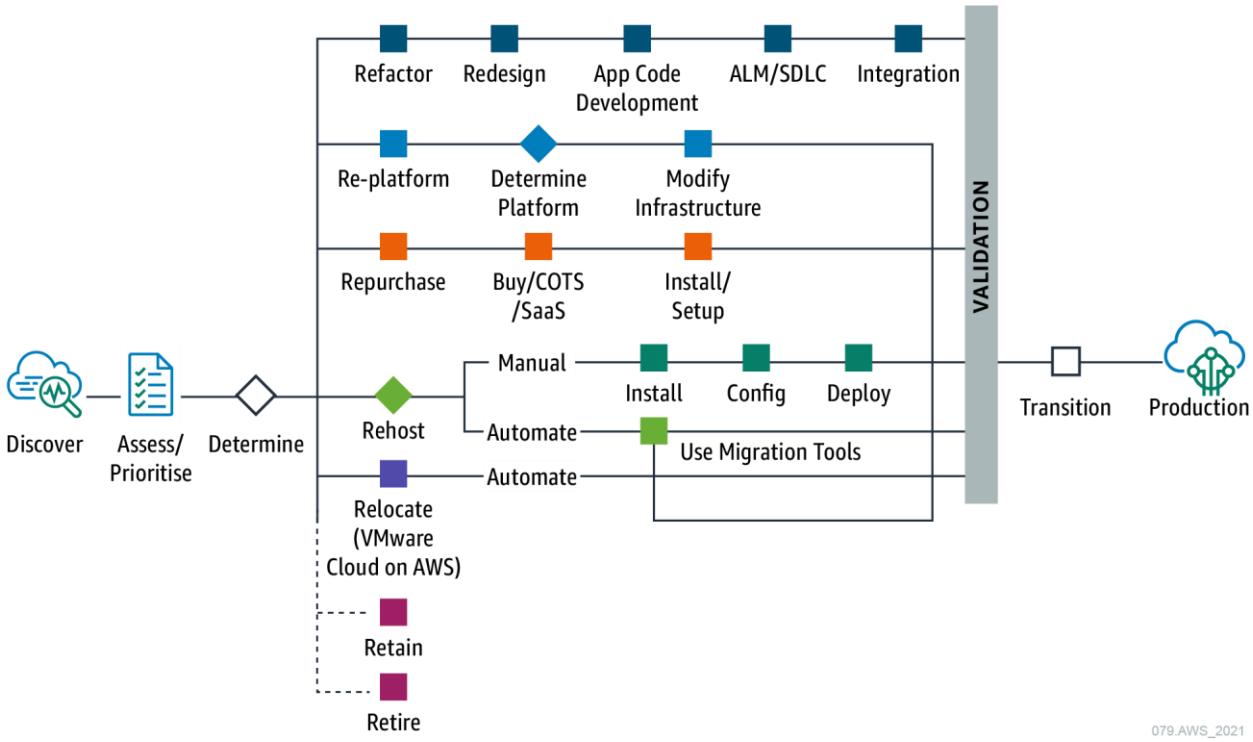

Figure 9. Seven Common Migration Patterns

Rehost (i.e. “lift and shift”) involves moving applications without changing them. A Lift-and-Shift migration is the process of migrating a workload from your on-premises data centers to the AWS Cloud, which requires little to no modification of those workloads. Customers choose a Lift-and-Shift approach to rehost their applications because it gives them the fastest migration execution, the lowest hosting cost, and it minimizes the risks associated with implementing changes. A Lift-and-Shift approach is common for the earlier stages of cloud adoption, and often acts as a transitionary state to a more cloud-native approach (whereby workloads are created and hosted in the cloud).

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.

AWS MGN utilizes continuous, block-level replication and enables short cutover windows measured in minutes. You can use AWS MGN to migrate all your applications and databases that run on supported versions of Windows and Linux operating systems.

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, while Amazon EFS is an elastic file system that uses the Network File System version 4 (NFSv4) protocol.

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. For example, 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. Moreover, part of AWS Database Migration Service is the free AWS Schema Conversion Tool (SCT) that automates the conversion of Oracle PL/SQL and SQL Server T-SQL code to equivalent code in the Amazon Aurora / MySQL dialect of SQL or the equivalent PL/pgSQL code in PostgreSQL.

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.

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.

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 brings the same hardware, software, services, application programming interfaces (APIs), management tools, support, and operating models that are used in the AWS Cloud to your data center, co-location space, or on-premises facility.

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.

187.8 Customer Responsibilities

Customers and AWS share security and compliance responsibilities. Generally, cloud service providers (CSPs) adopt shared responsibility models. These models define certain security tasks that lie with the customer and others that the CSP will take on. AWS's Shared Responsibility Model is designed to relieve your operational burdens while giving you control over your data and enabling you to implement security measures based on your unique needs.

As shown in Figure 10, AWS operates the services and underlying infrastructure that make up the cloud. We secure both the facilities that house the hardware and the foundational cloud technologies, like virtualization tools. You are responsible for managing your data, guest operating systems, and applications, as well as configuring the security of your AWS environment.

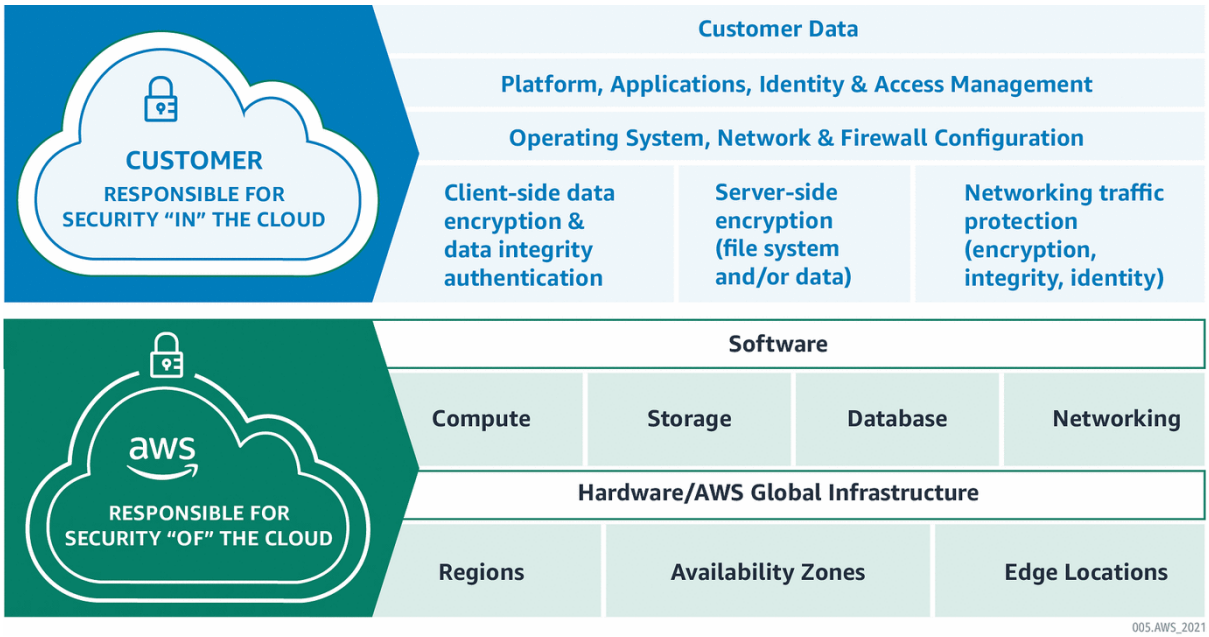

Figure 10. The Shared Responsibility Model *This model establishes that AWS is responsible for the underlying cloud infrastructure, while customers are responsible for their data and applications. The customer's specific responsibilities depend on the AWS services deployed.*

187.8.1 AWS and Customer Responsibilities

As described above, under the shared responsibility model, security and compliance responsibilities are shared between AWS and the customer.

187.8.1.1 AWS Responsibilities (Security of the Cloud)

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

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

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

187.8.1.2 Customer Responsibilities (Security in the Cloud)

AWS customers are responsible for *security in the cloud*. This refers to your applications, content, and systems. You have complete control over which services you use and whom

you empower to access your content and services, including what credentials are required to gain access. When using AWS services, you retain ownership and control of your content.

Because you, rather than AWS, control these important factors, you are responsible for securing any content that you store or process using our services and any content you connect to AWS infrastructure. This means you are responsible for the guest operating system, applications on your compute instances, and content stored and processed in AWS storage, databases, or other services.

At AWS, we want our customers to maintain comprehensive and robust security postures. As all customers are different, achieving this requires you to configure security controls based on your specific needs. For instance, you will need to set up access control policies so that only authorized individuals can access resources. You also need visibility into vulnerabilities that can threaten your applications, content, and systems.

With this in mind, AWS can help you reduce the burden of these responsibilities and properly handle your share of security of the infrastructure. We equip customers with documentation for all our services to guide them through proper configuration for security. We also offer hundreds of tools and features to help you meet your security objectives, including services for network security, configuration management, access control, and data encryption. Additionally, we offer integrated Partner solutions and support in the form of AWS customer enablement services to help you design, implement, and operate your security environment.

187.8.2 Shared Responsibility and IT Controls

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

Inherited Controls

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

Shared Controls

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

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

Customer-Specific Controls

Some controls are the sole responsibility of the customer based on the applications they deploy within the AWS Cloud (customer-specific controls). Examples include service and communications protection or zone security, which may require a customer to route or zone

data within specific security environments. Customers also control how they use their account and what content moves into and out of the account.

187.8.3 Shared Responsibility and Customer Data

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

187.8.3.1 Customer Content

With AWS, customers get the ability to store their content on cloud infrastructure and access it from almost anywhere over the internet. As a customer, you maintain full control of your content that you upload to the AWS services under your AWS account. This means that you maintain responsibility for configuring access to AWS services and resources. We provide an advanced set of access, encryption, and logging features to help you do this effectively. This includes APIs through which you can configure access control permissions for any of the services you develop or deploy in an AWS environment. We do not access or use your content for any purpose without your agreement. We never use your content or derive information from it for marketing or advertising purposes.

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

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

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

11.8.3.1.1 Legal Requests for Customer Content

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

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

- Amazon does not disclose customer information in response to government demands unless it is required to do so to comply with a legally valid and binding order. Unless prohibited from doing so or there is clear indication of illegal conduct in connection with the use of Amazon products or services, Amazon notifies customers before disclosing content information.
- Where Amazon needs to act to protect customers, it does. Amazon has repeatedly challenged government demands for customer information that were considered overbroad, winning decisions that have helped to set the legal standards for protecting customer speech and privacy interests. Amazon also advocates in Congress to modernize outdated privacy laws to require law enforcement to obtain a search warrant from a court to get the content of customer communications. That is the appropriate standard, and it is the standard Amazon and its subsidiaries follow.
- While Amazon recognizes the legitimate needs of law enforcement agencies to investigate criminal and terrorist activity and cooperate with them when they observe legal safeguards for conducting such investigations, it opposes legislation mandating or prohibiting security or encryption technologies that would weaken the security of products, systems, or services its customers use, whether they be individual consumers or business customers. For AWS customers, we offer strong encryption as one of many standard security features, and we provide them the option to manage their own encryption keys. We publish security best practices documents on our website and encourage our clients to use these measures to protect sensitive content.
- Amazon is a member of numerous associations focused on protecting privacy and security, and AWS in particular has achieved a number of internationally recognized certifications and accreditations demonstrating compliance with third-party assurance frameworks. AWS clients have control over their content and where it resides.

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

187.8.3.2 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.

187.8.3.3 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/.

187.8.4 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)