



G-Cloud 12 Amazon Web Services EMEA SARL, UK Branch (AWS) – Cloud Compute Infrastructure Services Service Definition Catalogue

20th July 2020

Points of Contact:

Chris Hayman

Director

AWS Public Sector

aws-gcloud@amazon.com

Neil Hall

Senior Business Contracts Manager

aws-gcloud@amazon.com

Disclaimer: It is the sole responsibility of the Partner to determine how to use the Service Offering information provided in this document. Prices referred to in the document or the incorporated links as referred to throughout the document are for illustration purposes only. Partners should refer to the pricing document on the Digital Marketplace for Supplier's G-Cloud pricing.

Table of Contents

1.0	Introduction.....	6
2.0	AWS Security Assurance	7
3.0	Service Definition – Alexa for Business	9
4.0	Service Definition – Amazon API Gateway.....	11
5.0	Service Definition – Amazon AppStream 2.0.....	13
6.0	Service Definition – Amazon Athena	15
7.0	Service Definition – Amazon Augmented AI.....	17
8.0	Service Definition – Amazon Aurora.....	19
9.0	Service Definition – Amazon Chime	21
10.0	Service Definition – Amazon Cloud Directory.....	23
11.0	Service Definition – Amazon CloudFront.....	25
12.0	Service Definition – Amazon CloudSearch.....	27
13.0	Service Definition – Amazon CloudWatch	29
14.0	Service Definition – Amazon Cognito	32
15.0	Service Definition – Amazon Comprehend.....	34
16.0	Service Definition – Amazon Comprehend Medical.....	35
17.0	Service Definition – Amazon Connect	37
18.0	Service Definition – Amazon DocumentDB (with MongoDB Compatibility)	39
19.0	Service Definition – Amazon DynamoDB	42
20.0	Service Definition – Amazon Elastic Block Store (Amazon EBS)	44
21.0	Service Definition – Amazon Elastic Compute Cloud (Amazon EC2).....	46
22.0	Service Definition – Amazon Elastic Container Registry (Amazon ECR)	63
23.0	Service Definition – Amazon Elastic Container Service (Amazon ECS).....	65
24.0	Service Definition – Amazon Elastic Container Service for Kubernetes (Amazon EKS) 66	
25.0	Service Definition – Amazon Elastic File System (Amazon EFS)	67
26.0	Service Definition – Amazon Elastic Graphics.....	69
27.0	Service Definition – Amazon Elastic Load Balancing	71
28.0	Service Definition – Amazon Elastic Inference	73
29.0	Service Definition – Amazon Elastic Transcoder.....	75
30.0	Service Definition – Amazon ElastiCache	77
31.0	Service Definition – Amazon Elasticsearch Service	80
32.0	Service Definition – Amazon EMR.....	82
33.0	Service Definition - Amazon EventBridge.....	83

34.0	Service Definition – Amazon Forecast.....	86
35.0	Service Definition – Amazon FreeRTOS	89
36.0	Service Definition – Amazon FSx for Lustre	91
37.0	Service Definition – Amazon FSx for Windows File Server	93
38.0	Service Definition – Amazon GameLift	95
39.0	Service Definition – Amazon Glacier	97
40.0	Service Definition – Amazon GuardDuty	98
41.0	Service Definition – Amazon Inspector.....	100
42.0	Service Definition – Amazon Kinesis	101
43.0	Service Definition – Amazon Lex.....	108
44.0	Service Definition – Amazon Lightsail	110
45.0	Service Definition – Amazon Machine Learning	112
46.0	Service Definition – Amazon Macie	113
47.0	Service Definition – Amazon Managed Blockchain	115
48.0	Service Definition – Amazon MQ.....	118
49.0	Service Definition – Amazon Neptune	120
50.0	Service Definition – Amazon Personalize	121
51.0	Service Definition – Amazon Pinpoint.....	123
52.0	Service Definition – Amazon Polly.....	125
53.0	Service Definition – Amazon Quantum Ledger Database (QLDB)	127
54.0	Service Definition – Amazon QuickSight	130
55.0	Service Definition – Amazon Redshift	131
56.0	Service Definition – Amazon Relational Database Service (Amazon RDS)	134
57.0	Service Definition – Amazon Route 53.....	138
58.0	Service Definition – Amazon SageMaker	142
59.0	Service Definition – Amazon Simple Email Service (Amazon SES)	147
60.0	Service Definition – Amazon Simple Notification Service (Amazon SNS)	149
61.0	Service Definition – Amazon Simple Queue Service (Amazon SQS)	150
62.0	Service Definition – Amazon Simple Storage Service (Amazon S3)	151
63.0	Service Definition – Amazon Simple Workflow Service (Amazon SWF).....	154
64.0	Service Definition – Amazon Sumerian	155
65.0	Service Definition – Amazon Textract.....	158
66.0	Service Definition – Amazon Transcribe.....	159
67.0	Service Definition – Amazon Transcribe Medical	161

68.0	Service Definition – Amazon Translate.....	163
69.0	Service Definition – Amazon Virtual Private Cloud (Amazon VPC)	164
70.0	Service Definition – Amazon WorkDocs	166
71.0	Service Definition – Amazon WorkLink.....	167
72.0	Service Definition – Amazon WorkMail.....	170
73.0	Service Definition – Amazon WorkSpaces	171
74.0	Service Definition – Amazon WorkSpaces Application Manager (Amazon WAM) 173	
75.0	Service Definition - AWS Amplify Console	175
76.0	Service Definition - AWS App Mesh	176
77.0	Service Definition – AWS Application Discovery Service	178
78.0	Service Definition – AWS AppSync	179
79.0	Service Definition – AWS Artifact	181
80.0	Service Definition – AWS Auto Scaling	182
81.0	Service Definition – AWS Backup.....	184
82.0	Service Definition – AWS Batch	186
83.0	Service Definition – AWS Certificate Manager	187
84.0	Service Definition – AWS Cloud Map	193
85.0	Service Definition – AWS Cloud9	194
86.0	Service Definition – AWS CloudFormation	196
87.0	Service Definition – AWS CloudHSM	199
88.0	Service Definition – AWS CloudTrail	201
89.0	Service Definition – AWS CodeBuild	203
90.0	Service Definition – AWS CodeCommit.....	204
91.0	Service Definition – AWS CodeDeploy.....	206
92.0	Service Definition – AWS CodePipeline	207
93.0	Service Definition – AWS CodeStar	209
94.0	Service Definition – AWS Config	210
95.0	Service Definition – AWS Data Pipeline	212
96.0	Service Definition – AWS Database Migration Service.....	213
97.0	Service Definition – AWS DataSync	215
98.0	Service Definition - AWS DeepRacer	217
99.0	Service Definition – AWS Device Farm	218
100.0	Service Definition – AWS Direct Connect	220



101.0	Service Definition – AWS Directory Service	222
102.0	Service Definition – AWS Elastic Beanstalk	223
103.0	Service Definition - AWS Elemental MediaConnect	225
104.0	Service Definition – AWS Elemental MediaConvert	227
105.0	Service Definition – AWS Elemental MediaLive	228
106.0	Service Definition – AWS Elemental MediaPackage	229
107.0	Service Definition – AWS Elemental MediaStore	231
108.0	Service Definition – AWS Elemental MediaTailor	232
109.0	Service Definition – AWS Fargate	233
110.0	Service Definition – AWS Firewall Manager	235
111.0	Service Definition – AWS Global Accelerator	237
112.0	Service Definition – AWS Glue	240
113.0	Service Definition – AWS Greengrass	242
114.0	Service Definition – AWS Identity and Access Management (IAM)	243
115.0	Service Definition – AWS IoT	245
116.0	Service Definition – AWS Key Management Service (AWS KMS)	253
117.0	Service Definition – AWS Lake Formation	255
118.0	Service Definition – AWS Lambda	256
119.0	Service Definition – AWS License Manager	257
120.0	Service Definition – AWS Migration Hub	260
121.0	Service Definition – AWS Mobile Hub	262
122.0	Service Definition – AWS OpsWorks	263
123.0	Service Definition – AWS Organizations	265
124.0	Service Definition - AWS Resource Access Manager	267
125.0	Service Definition – AWS RoboMaker	268
126.0	Service Definition – AWS Secrets Manager	271
127.0	Service Definition – AWS Security Hub	272
128.0	Service Definition – AWS Server Migration Service (SMS)	274
129.0	Service Definition – AWS Serverless Application Repository	275
130.0	Service Definition – AWS Service Catalog	277
131.0	Service Definition – AWS Shield	278
132.0	Service Definition – AWS Single Sign-On	280
133.0	Service Definition – AWS Snowball Edge	281
134.0	Service Definition – AWS Snowmobile	283

135.0	Service Definition – AWS Step Functions	285
136.0	Service Definition – AWS Storage Gateway	286
137.0	Service Definition – AWS Systems Manager	288
138.0	Service Definition – AWS Transfer for SFTP (AWS SFTP)	290
139.0	Service Definition – AWS Transit Gateway	293
140.0	Service Definition – AWS Trusted Advisor	295
141.0	Service Definition – AWS Web Application Firewall (AWS WAF)	297
142.0	Service Definition – AWS Well-Architected Tool	299
143.0	Service Definition – AWS X-Ray	300
144.0	Service Definition - Service Quotas	301
145.0	Cross-Service Definitions	303

1.0 Introduction

This document provides service definitions for the Amazon Web Services EMEA SARL, UK Branch (AWS) Service Offerings included in the G-Cloud 12 framework catalogue. We have broken out service definitions in accordance with Invitation to Tender (ITT) requirements.

1.1 How to use the AWS Service Definition Documents

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

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

This AWS Cloud Compute Infrastructure Services Service Definition document describes the key features for each of the different Cloud Compute Services available to Customers on G-Cloud 12 in Lots 1 & 2. Notwithstanding that AWS has combined its service descriptions into a consolidated document for ease of review by Customers, to access the options through a Call-Off Contract the Customer must reference each individual Digital Marketplace Service ID within the Call-Off Contract in order to enable that service as an option that can be procured under their G-Cloud 12 Call-Off Contract. AWS would recommend that Buyers list all of the Digital Marketplace Service ID's for every service described in this document in its Call-Off Contract to enable the option to switch between Services flexibly during the term. For a list of all AWS Digital Marketplace Service ID's, please contact an AWS account representative through aws-gcloud@amazon.com.

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

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

Please note that the options or parameters selected by AWS on this framework are those that most closely align with our existing commercial services. AWS is willing to provide additional information about our services upon request.

2.0 AWS Security Assurance

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

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

British Standard 7858:2012

Buyers selecting AWS Services and expressly requiring AWS conformity to BS7858:2012 acknowledge that AWS scopes BS7858:2012 compliance to those AWS employees with physical access to the 'data layer' zones within datacentres as well as critical permission groups, and those who are directed by the Buyer to access Buyer Data such as Technical Account Managers ("TAMS"). A list of TAMs shall be provided to the Buyer by the Supplier prior to the Start date of the Call-Off Contract and the Buyer shall only contact the listed TAMs in relation to Buyer Data during the Term of the Call-Off Contract. Buyers are obliged in accordance with the Call-Off Contract to encrypt Buyer Data when using AWS Services. Buyer should note that the Supplier does not include Supplier Staff (as defined in the Call-Off Contract) responsible for operating the AWS Services or those with logical access to encrypted Buyer Data for the purposes of its BS7858 compliance.

2.1 Information Assurance

The following subsections provide information relating to information assurance.

2.1.1 ISO 27001 Certification

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

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

2.1.2 NCSC UK Cloud Security Principles

In 2016, National Cyber Security Centre (NCSC) UK published the [Cloud Security Collection](#) documents for public sector organisations that are considering the use of cloud services for handling information classified as OFFICIAL. The collection of guidance documents aims to help public sector organisations make informed decisions about cloud services and choose a cloud service that balances business benefits and security risks. In order to provide you with more information regarding NCSC UK's Cloud Security Principles and to make an informed decision when performing risk assessments, we have published a whitepaper called [Using AWS in the Context of NCSC UK's Cloud Security Principles](#).

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

2.2 GDPR and processing of Personal Data

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

- AWS GDPR Center: <https://aws.amazon.com/de/compliance/gdpr-center/>
- AWS EU Data Protection website: <https://aws.amazon.com/compliance/eudata-protection/>

2.3 AWS re:Start Program

AWS re:Start is an initiative intended to promote the development of much-needed technical skills in the UK. It is a training and job placement programme that was established in late March 2017, coincident with the launch of the AWS London Region. AWS re:Start has been designed to educate young adults, military veterans, members of the military reserve, those leaving the Armed Forces, and service spouses on the latest software development and cloud computing technologies.

AWS re:Start is the result of a close collaboration between the Ministry of Defence, QA Consulting (an APN Training Partner) and The Prince's Trust. In conjunction with members of the AWS Partner Network (APN) and customers, the programme aims to offer work placements to 1,000 people.

AWS re:Start accommodates participants at all levels of experience—even those with no previous technical knowledge. Participants who join AWS re:Start complete technical training classes, led by AWS certified instructors, and gain experience through on-the-job training. They learn about multi-tier architectures, application programming interfaces (APIs), and microservices. Training content for the AWS re:Start program is curated by AWS in collaboration with QA Consulting, which delivers the training courses.

Organisations that have pledged job placements to AWS re:Start include Annalect, ARM, Claranet, Cloudreach, Direct Line Group, EDF Energy, Funding Circle, KCOM, Sage, Tesco Bank, and Zopa. Participants completing the programme are eligible for many different technical positions within these companies, including sought-after entry level positions such as first-line help desk support, IT support analyst, software developer, IT support technician, network engineer, IT recruitment consultant, and IT sales roles.

The programme also provides participants with the fundamental knowledge needed to start working with AWS to build technology start-up businesses. The Young Adults thread of AWS re:Start is delivered via the Prince's Trust Get into Technology initiative. In addition to technical training, the "Get into Technology" programme supports students with mentoring, soft work skills, and help in applying for jobs, including resume writing and interview skills. [Contact us here](#) for information and/or to get involved in the program.

3.0 Service Definition – Alexa for Business

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

3.1 Service Overview

Alexa for Business allows organisations of all sizes to introduce Alexa to their workplace. With Alexa for Business, customers can use the Alexa they know as an intelligent assistant to stay organised and focus on the work that matters. Alexa helps workers be more productive as they move throughout their day, at home and at their

desks as enrolled users with personal devices, and in conference rooms, copy rooms or other shared spaces with shared devices. Alexa for Business includes the tools and controls that administrators need to deploy and manage shared Alexa devices, skills, and users at scale. Top benefits include:

- **Easily Provision and Manage Alexa Devices** - Alexa for Business allows you to provision and manage Alexa devices in your organisation from a centralized console. With Alexa for Business, you can easily provision multiple Alexa devices at the same time, and automatically connect them to your Alexa for Business account. You can specify device locations, enable a set of skills that can be used, and prevent users from tampering with them. This saves time because you don't need to manage these devices individually.
- **Configure Conference Rooms** - Alexa for Business makes it easy for you to configure Alexa to control your conference rooms. Alexa for Business lets you specify the type of conferencing equipment you use and your preferred meeting applications, which allows Alexa to start most meetings, on most devices, in any room. You can use Alexa devices as audio conferencing devices in small conference rooms, or to control equipment in larger rooms. Alexa for Business is an open service, and the Alexa for Business APIs allow you to build skills so that Alexa can work with additional equipment or perform specific tasks in your conference rooms.
- **Manage Users** - Alexa for Business allows you to invite your end users to enrol their personal Alexa account with your Alexa for Business account. This lets them continue to use the Alexa features and skills they've already enabled in their personal Alexa account, as well as the work skills you provide, on any of their devices, at work or at home. Alexa for Business gives you the ability to make work skills available and provide access to your corporate calendar system so that they can use Alexa to manage their calendar.
- **Create Custom Skills** - Alexa for Business lets you build your own private custom skills for your workplace, your employees, or your customers to use. You can make these skills available only to your shared Alexa devices, and your enrolled users. Alexa for Business provides an additional set of APIs that provide information about device location, which lets you add context to your skills. For example, you could build a skill that lets a user report a printer problem to IT, and the skill could use the device location so that IT knows which printer is broken. Building custom skills is easy, and the Alexa Skills Kit provides tools, documentation, and code samples to help you get started.

3.2 Backup/Restore and Disaster Recovery

This requirement is not applicable for Alexa for Business. 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 12 Pricing Document affiliated with this service in the Digital Marketplace.

3.4 Service Constraints

Please see <https://aws.amazon.com/alexaforbusiness/> for more information. Prior to using these Services, Buyer agrees to contact Supplier before using these Services in connection with Buyer Personal Data.

3.5 Technical Requirements

Alexa for Business gives you the tools you need to manage Alexa devices, enrol users, and assign skills. You can build your own voice skills using the Alexa Skills Kit and the Alexa for Business API. You can also make them available as private skills for your organisation.

- **Administration Guide** - Helps you perform several administrative tasks for Alexa for Business, such as getting set up and managing your devices, rooms, skills, and users. [HTML](#) | [PDF](#).
- **API Reference** – Describes the API operations for Alexa for Business. [HTML](#) | [PDF](#)

4.0 Service Definition – Amazon API Gateway

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

4.1 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. With a few clicks in the AWS Management Console, you can create an API that acts as a “front door” for applications to access data, business logic, or functionality from your back-end services, such as workloads running on Amazon EC2, code running on AWS Lambda, or any Web application. Amazon API Gateway handles all the tasks involved in accepting and processing up to hundreds of thousands of concurrent API calls, including traffic management, authorization and access control, monitoring, and API version management. Amazon API Gateway has no minimum fees or start-up costs. You pay only for the API calls you receive and the amount of data transferred out.

Top benefits include:

- **Low-Cost and Efficient** – With Amazon API Gateway, you pay only for calls made to your APIs and data transfer out. There are no minimum fees or up-front commitments.
- **Performance at Any Scale** – With Amazon CloudFront integration, Amazon API

Gateway allows you to take advantage of the worldwide network of edge locations to provide your end users with the lowest possible latency for API requests and responses. Amazon API Gateway also helps you manage traffic through throttling, so that back-end operations can withstand traffic spikes. Additionally, Amazon API Gateway helps you improve the performance of your APIs by caching the output of API calls to avoid calling your backend systems unnecessarily.

- **Easily Monitor API Activity** – After your API is deployed, Amazon API Gateway provides you with a dashboard to visually monitor calls to your services using [Amazon CloudWatch](#), so you see performance metrics and information on API calls, data latency, and error rates.
- **Streamline API Development** – Amazon API Gateway lets you simultaneously run multiple versions of the same API, allowing you to quickly iterate, test, and release new versions.
- **Flexible Security Controls** – Amazon API Gateway provides you with tools to authorise access to your APIs and control service operation access. You can use AWS administration and security tools, such as AWS IAM and Amazon Cognito, to authorise access to your APIs. Amazon API Gateway can verify signed API calls on your behalf using the same technology AWS uses for its own APIs. If you already use [OAuth](#) tokens or other authorization mechanisms, Amazon API Gateway can use AWS Lambda to execute a custom authoriser to help you verify incoming requests.
- **Create RESTful Endpoints for Existing Services** – With Amazon API Gateway, you can create modern, resource-based APIs and then use the dynamic and flexible data transformation capabilities to generate the requests in the language your target services expect. Amazon API Gateway also helps you protect your existing services by setting throttling rules to avoid overwhelming your back-end infrastructure during unpredictable traffic spikes.
- **Run Your APIs Without Servers** – Amazon API Gateway tightly integrates with AWS Lambda to allow you to create completely server-less APIs. With Amazon API Gateway, you create REST APIs that your mobile and web applications can use to call publicly available AWS Cloud services through your code running in AWS Lambda. AWS Lambda runs your code on a high-availability compute infrastructure and performs all the heavy lifting and administration of your compute resources.

4.2 Backup/Restore and Disaster Recovery

This requirement is not applicable for AWS Lambda. 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 12 Pricing Document affiliated with this service in the Digital Marketplace.

4.4 Service Constraints

Please see <https://aws.amazon.com/api-gateway/> for more information.

4.5 Technical Requirements

Amazon API Gateway is a fully managed service that makes it easy for developers to publish, maintain, monitor, and secure APIs at any scale. Create an API to access data, business logic, or functionality from your back-end services, such as applications running on Amazon EC2, code running on AWS Lambda, or any web application.

- **Developer Guide** – Provides a conceptual overview of Amazon API Gateway and includes detailed instructions for using the service. [HTML](#) | [PDF](#) | [Kindle](#)
- **API Reference** – Describes all the Amazon API Gateway operations in detail. Also provides sample requests, responses, and errors for the supported web services protocols. [HTML](#)

5.0 Service Definition – Amazon AppStream 2.0

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

5.1 Service Overview

Amazon AppStream 2.0 is a fully managed application streaming service. You centrally manage your desktop applications on AppStream 2.0 and securely deliver them to any computer. You can easily scale to any number of users across the globe without acquiring, provisioning, and operating hardware or infrastructure.

Top Features include:

- Allows 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 or Windows Server 2019 is supported, and you don't need to make any modifications.
- Runs your applications on virtual machines (VMs) called 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.
- With Amazon AppStream 2.0, your desktop applications can be launched inside an Amazon Virtual Private Cloud (VPC). You can use VPC security groups to provide granular access control to streaming instances, and to manage users'

access to the resources in your VPC, such as your databases, file shares, license servers, and application servers.

- Your Amazon AppStream 2.0 streaming instances 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. Your users are authenticated using a SAML 2.0 provider of your choice, and can access applications that require a connection to your AD domain.
- 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.

Top benefits include:

- Deliver desktop applications to any computer - Users can access the desktop applications they need at any time. AppStream 2.0 streams your applications from AWS to any computer, including Chromebooks, Macs, and PCs.
- Secure applications and data - Applications and data are not stored on users' computers. Your applications are streamed as encrypted pixels and access data secured within your network. AppStream 2.0 runs on AWS, so you benefit from a data center and network architecture built for the most security-sensitive organizations.
- Scale without infrastructure - AppStream 2.0 is fully managed on AWS and globally available with pay-as-you-go pricing. You can easily scale your application streaming to any number of users across the globe without acquiring, provisioning, and operating hardware or infrastructure.
- Centrally manage applications - Each user accesses the same version of your applications. You centrally manage your applications on AppStream 2.0 and can stop managing installations and updates on each user's computer.
- Integrate with your IT - AppStream 2.0 connects to your Active Directory, network, cloud storage, and file shares. Users access applications using their existing credentials and your existing security policies manage access. Extensive APIs integrate AppStream 2.0 with your IT solutions.
- Provide a fluid and responsive user experience - Each user's applications are highly responsive because they run on VMs optimized for their use cases. The NICE DCV protocol automatically adjusts each streaming session to network conditions for a fluid user experience.

5.2 Backup/Restore and Disaster Recovery

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.

Home folders are AppStream 2.0's native persistent storage option. Users can access a home folder on their streaming instance and save content in their folder. Files are stored in an S3 bucket which is automatically created in your AWS account. To learn more, visit [Enable and Administer Home Folders for Your AppStream 2.0 Users](#).

Please refer to the following link for information regarding Amazon S3 resilience:
<https://docs.aws.amazon.com/AmazonS3/latest/dev/disaster-recovery-resiliency.html>

5.3 Pricing Overview

Please see the AWS UK G-Cloud 12 Pricing Document affiliated with this service in the Digital Marketplace.

5.4 Service Constraints

Please see <https://aws.amazon.com/appstream2/> for more information.

5.5 Technical Requirements

The Amazon AppStream web service deploys your application on AWS infrastructure and streams input and output between your application and devices such as personal computers, tablets, and mobile phones. Your application's processing occurs in the cloud, so it can scale to handle vast computational loads. Devices need only display output and return user input, so the client application on the device can be lightweight in terms of file size and processing requirements.

- **Developer Guide** – Helps you get started using Amazon AppStream 2.0. You will learn how to quickly and easily provision and maintain one or more WorkSpaces. [HTML](#) | [PDF](#)
- **API Reference Guide** - Describes the API operations for Amazon AppStream 2.0. [HTML](#) | [PDF](#)

6.0 Service Definition – Amazon Athena

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

6.1 Service Overview

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

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

Top benefits include:

- **Start Querying Instantly: Serverless. No ETL** – Amazon Athena is serverless. You can quickly query your data without having to setup and manage any servers or data warehouses. Just point to your data in Amazon S3, define the schema, and start querying using the built-in query editor. Amazon Athena allows you to tap into all your data in S3 without the need to set up complex processes to extract, transform, and load the data (ETL).
- **Pay Per Query: Only pay for data scanned** – With Amazon Athena, you pay only for the queries that you run. You can save from 30% to 90% on your per query costs and get better performance by compressing, partitioning, and converting your data into columnar formats. Amazon Athena queries data directly in Amazon S3. There are no additional storage charges beyond S3.
- **Open. Powerful. Standard: Built on Presto. Runs standard SQL** – Amazon Athena uses Presto with ANSI SQL support and works with a variety of standard data formats, including CSV, JSON, ORC, Avro, and Parquet. Amazon 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.
- **Fast. Really Fast: Interactive performance even for large datasets** – With Amazon Athena, you don't have to worry about having enough compute resources to get fast, interactive query performance. Amazon Athena automatically executes queries in parallel, so most results come back within seconds.

6.2 Backup/Restore and Disaster Recovery

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

6.3 Pricing Overview

Please see the AWS UK G Cloud 12 Pricing Document affiliated with this service in the Digital Marketplace.

6.4 Service Constraints

Please see <https://aws.amazon.com/athena/> for more information.

6.5 Technical Requirements

Amazon Athena is an interactive query service that makes it easy to analyse data in Amazon S3 using standard SQL. Athena is serverless, so there is no infrastructure to setup or manage, and you pay only for the queries you run. To get started, simply point to your data in S3, define the schema, and start querying using standard SQL.

User Guide – Learn how to use Athena to query data stored in Amazon S3. [HTML](#) | [PDF](#)

7.0 Service Definition – Amazon Augmented AI

Amazon Augmented AI (Amazon A2I) is a service that makes it easy to build the workflows required for human review of Machine Learning predictions. Amazon A2I brings human review to all developers, removing the undifferentiated heavy lifting associated with building human review systems or managing large numbers of human reviewers.

7.1 Service Overview

Top features include:

- Easy integration into existing AWS or other Machine Learning models. Integrate Amazon Textract for document processing and Amazon Rekognition for content moderation with just a few clicks in the Amazon A2I console or a few API parameters.
- Flexibility to work with reviewers inside and outside of your organization
- Easy instructions for reviewers by providing instructional guidance to human reviewers to help ensure consistency. These detailed instructions are available to reviewers within their review interface.
- Use workflows to simplify the human review process. Use conditions to route predictions to reviewers using confidence thresholds, random sampling percentage or a custom workflow.
- Provide reviewers with a web interface consisting of all the instructions and tools they need to complete their tasks. Amazon A2I provides built-in workflows for text extraction and image moderation use cases.

Top benefits include:

- Gives you the flexibility to incorporate human review into Machine Learning applications based on your specific requirements. Low-confidence predictions are sent to humans to review and take action. If needed, you can also require multiple reviewers to review a prediction to achieve consensus.
- Provides options to work with reviewers inside and outside of your organization. Using Amazon A2I, you can easily route reviews to the reviewers you provide or alternatively access a workforce of over 500,000 independent contractors who are already performing machine learning related tasks through Amazon Mechanical Turk.
- Makes it easy to integrate human judgement and Artificial Intelligence into any Machine Learning application, regardless of whether it's run on AWS or on another platform. Amazon A2I is directly integrated into Amazon Textract for document processing and Amazon Rekognition for content moderation, so you can easily add human reviews to these use cases with just a few clicks in the AWS console.

7.2 Backup/Restore and Disaster Recovery

This requirement is not applicable for Amazon Augmented AI. For additional information beyond what is described herein, please refer to <http://aws.amazon.com/documentation/>.

The AWS global infrastructure is built around AWS Regions and Availability Zones. AWS Regions provide multiple physically separated and isolated 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 centre infrastructures.

For more information about AWS Regions and Availability Zones, see AWS Global Infrastructure (<http://aws.amazon.com/about-aws/global-infrastructure/>).

7.3 Pricing Overview

Please see the AWS UK G Cloud 12 Pricing Document affiliated with this service in the Digital Marketplace.

7.4 Service Constraints

Please see <https://docs.aws.amazon.com/sagemaker/latest/dg/a2i-getting-started.html> (<http://aws.amazon.com/documentation/augmented-ai/>) for more information.

7.5 Technical Requirements

Please refer to <https://aws.amazon.com/augmented-ai/> and the following links for comprehensive technical documentation regarding Amazon Augmented AI:

- <https://docs.aws.amazon.com/sagemaker/latest/dg/use-augmented-ai-a2i-human-review-loops.html>
- <https://docs.aws.amazon.com/sagemaker/latest/dg/a2i-permissions-security.html>

8.0 Service Definition – Amazon Aurora

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

This service listing includes the following AWS services:

- Amazon Aurora - PostgreSQL-compatible
- Amazon Aurora - MySQL-compatible

8.1 Service Overview

Amazon Aurora is a MySQL and PostgreSQL-compatible relational database engine that combines the speed and availability of high-end commercial databases with the simplicity and cost effectiveness of open-source databases. Amazon Aurora provides up to five times better performance than MySQL with the security, availability, and reliability of a commercial database at one tenth the cost.

Top features include:

- Fully managed MySQL and PostgreSQL compatible relational database.
- Distributed, fault-tolerant, self-healing storage system that auto-scales.
- Amazon Aurora Global Database is designed for globally distributed applications, allowing a single Amazon Aurora database to span multiple AWS regions.
- Amazon Aurora Serverless is an on-demand, auto-scaling configuration for [Amazon Aurora](#) (MySQL-compatible and PostgreSQL-compatible editions), where the database will automatically start up, shut down, and scale capacity up or down based on your application's needs.
- Amazon Aurora machine learning enables you to add ML-based predictions to applications via the familiar SQL programming language, so you don't need to learn separate tools or have prior machine learning experience.

Top Benefits include:

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

databases, at 1/10th the cost. Amazon Aurora uses a variety of software and hardware techniques to ensure the database engine is able to fully leverage available compute, memory and networking. I/O operations use distributed systems techniques such as quorums to improve performance consistency.

- **Highly Scalable** - You can easily scale your database deployment up and down from smaller to larger instance types as your needs change, or let [Aurora Serverless](#) handle scaling automatically for you. To scale read capacity and performance, you can add up to 15 low latency read replicas across three Availability Zones. Amazon Aurora automatically grows storage as needed, up to 64TB per database instance.
- **Highly Secure** – Amazon Aurora provides multiple levels of security for your database. These include network isolation using [Amazon VPC](#), encryption at rest using keys you create and control through [AWS Key Management Service](#) (KMS) and encryption of data in transit using SSL. On an encrypted Amazon Aurora instance, data in the underlying storage is encrypted, as are the automated backups, snapshots, and replicas in the same cluster.
- **High Availability and Durability** – Amazon Aurora is designed to offer greater than 99.99% availability, replicating 6 copies of your data across 3 Availability Zones and backing up your data continuously to Amazon S3. It transparently recovers from physical storage failures; instance failover typically takes less than 30 seconds. You can also backtrack within seconds to a previous point in time, to recover from user errors. With [Global Database](#), a single Aurora database can span multiple AWS regions to enable fast local reads and quick disaster recovery.
- **Fully Managed** – Amazon Aurora is fully managed by Amazon Relational Database Service (RDS). You no longer need to worry about database management tasks such as hardware provisioning, software patching, setup, configuration, or backups. Aurora automatically and continuously monitors and backs up your database to Amazon S3, enabling granular point-in-time recovery. You can monitor database performance using Amazon CloudWatch, [Enhanced Monitoring](#), or [Performance Insights](#), an easy-to-use tool that helps you quickly detect performance problems.
- **MySQL and PostgreSQL Compatible** – The Amazon Aurora database engine is fully compatible with existing MySQL and PostgreSQL open source databases, and adds compatibility for new releases regularly. This means you can easily migrate MySQL or PostgreSQL databases to Aurora using standard MySQL or PostgreSQL import/export tools or snapshots. It also means the code, applications, drivers, and tools you already use with your existing databases can be used with Amazon Aurora with little or no change.
- **Migration Support** - MySQL and PostgreSQL compatibility make Amazon Aurora a compelling target for database migrations to the cloud. If you're migrating from MySQL or PostgreSQL, see our [migration documentation](#) for a list of tools and options. To migrate from commercial database engines, you can use the [AWS Database Migration Service](#) for a secure migration with minimal downtime.

8.2 Backup/Restore and Disaster Recovery

Amazon Aurora is designed to offer greater than 99.99% availability, increasing MySQL and PostgreSQL performance and availability by tightly integrating the database engine with an SSD-backed virtualized storage layer purpose-built for database workloads. Amazon Aurora's storage is fault-tolerant and self-healing, and disk failures are repaired in the background without loss of database availability. Amazon Aurora is designed to automatically detect database crashes and restart without the need for crash recovery or to rebuild the database cache. If the entire instance fails, Amazon Aurora will automatically fail over to one of up to 15 read replicas.

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 12 Pricing Document affiliated with this service in the Digital Marketplace.

8.4 Service Constraints

Please see

https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/CHAP_Limits.html for more information.

8.5 Technical Requirements

Please refer to <https://docs.aws.amazon.com/rds/index.html> and the following links for comprehensive technical documentation regarding Amazon Aurora.

- https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/CHAP_AuroraOverview.html
- <https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/UsingWithRDS.html>
- <https://aws.amazon.com/rds/aurora/faqs/>

9.0 Service Definition – Amazon Chime

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

9.1 Service Overview

Amazon Chime is a high-quality communications service that transforms online meetings with an easy-to-use app that works seamlessly across all your devices. With Amazon Chime, you can schedule and attend online meetings and video conferences, and chat, call, and collaborate, inside and outside your organisation, all with a single app. Now you can work productively from wherever you are.

Frustration free online meetings that start on time. High-definition video conferencing. A single, easy-to-use communications application that you can trust.

Top benefits include:

Frustration-Free Online Meetings – Make your online meetings effortless with Amazon Chime. Meetings start on time, and they're easy to join. Meetings call you and with a single tap or click, you can join or notify participants that you're running late. A visual roster shows who is in the meeting and who is running late, and allows anyone to mute background noise. Amazon Chime automatically reconnects you if you get dropped. Crystal clear audio and HD video keep you focused on the discussion

Be Mobile – Take your meetings anywhere with a rich, easy-to-use application available for Android, iOS, Mac, and Windows. Meetings and chats are always synchronised, you can join meetings from any device, and you can switch seamlessly between devices anytime – even in the middle of a meeting.

One Application for All Your Communications – Use one application to conduct online meetings, video conferencing, chat, screen sharing, and file sharing, both inside and outside your organisation.

9.2 Backup/Restore and Disaster Recovery

This requirement is not applicable for Amazon Chime. 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 12 Pricing Document affiliated with this service in the Digital Marketplace.

9.4 Service Constraints

Amazon Chime features and default service limits are found at <https://chime.aws/faq/>.

9.5 Technical Requirements

Amazon Chime is a secure, real-time, unified communications service that transforms meetings by making them more efficient and easier to conduct.

Administrator Guide Helps you use Amazon Chime to perform several administrative tasks, such as creating an Amazon Chime account, inviting users, and managing licenses. [HTML](#) | [PDF](#)

User Guide Helps you use Amazon Chime, including joining, attending, and scheduling meetings for your organisation. [HTML](#) | [PDF](#)

10.0 Service Definition – Amazon Cloud Directory

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

10.1 Service Overview

Amazon Cloud Directory enables you to build flexible cloud-native directories for organizing hierarchies of data along multiple dimensions. With Cloud Directory, you can create directories for a variety of use cases, such as organisational charts, course catalogues, and device registries. While traditional directory solutions, such as Active Directory Lightweight Directory Services (AD LDS) and other LDAP-based directories, limit you to a single hierarchy, Cloud Directory offers you the flexibility to create directories with hierarchies that span multiple dimensions. For example, you can create an organisational chart that can be navigated through separate hierarchies for reporting structure, location, and cost centre.

Amazon Cloud Directory automatically scales to hundreds of millions of objects and provides an extensible schema that can be shared with multiple applications. As a fully managed service, Cloud Directory eliminates time-consuming and expensive administrative tasks, such as scaling infrastructure and managing servers. You simply define the schema, create a directory, and then populate your directory by making calls to the Cloud Directory API.

Top benefits include:

Efficiently organise hierarchies of data across multiple dimensions - Many applications require building directories to store hierarchies of data with

multiple dimensions. For example, an organisational chart could have one hierarchy based on reporting structure, a second based on location, and a third based on cost centre. Amazon Cloud Directory enables you to combine different schemas within a single directory to create multiple hierarchies without having to duplicate data.

Scale automatically on managed infrastructure - Amazon Cloud Directory is designed to automatically scale to hundreds of millions of objects. Amazon built Cloud Directory on highly scalable, managed infrastructure that provides high availability and built-in server management. Unlike traditional databases and directories, Cloud Directory does not require upfront capacity planning or continuous, manual optimization.

Search your directory for objects and relationships - Searching large sets of highly connected data has traditionally required complex and computationally expensive queries. With Amazon Cloud Directory's built-in search capabilities, you can search for all the parent objects along a dimension without creating multiple, nested queries. For example, looking up a management chain for an employee can be performed with a single query in Cloud Directory where traditional solutions would require multiple queries.

Easily adapt to changing data requirements - Application data requirements evolve over time. Amazon Cloud Directory is built around a model that explicitly makes room

for change. Cloud Directory has a flexible schema so you can extend your schema with new attributes at any time. Multiple applications can independently extend the schema. This enables applications to share a single directory and avoid duplication of data.

Simplify policy management - With Amazon Cloud Directory, you can attach application-specific policies to objects and define policy inheritance rules along different hierarchies. For example, you can create a custom policy for an organisational chart to allocate budget limits through your cost centre hierarchy, and enforce naming conventions for your device hierarchy. Cloud Directory enables your application to evaluate and enforce these policies on demand to meet the unique requirements of your application. You are not limited by the native policies of a directory.

Integrate with AWS CloudTrail and tagging - Amazon Cloud Directory is integrated with AWS CloudTrail and resource tagging. Using AWS CloudTrail you can log the date, time, and identity of users who accesses your directory data. With resource tagging, you can tag your directories and schemas to better track and manage resources.

10.2 Backup/Restore and Disaster Recovery

This requirement is not applicable for Amazon Cloud Directory. 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 12 Pricing Document affiliated with this service in the Digital Marketplace.

10.4 Service Constraints

Please see <http://aws.amazon.com/documentation/iam/> for more information.

10.5 Technical Requirements

AWS Directory Service provides multiple ways to set up and run Amazon Cloud Directory, Amazon Cognito, and Microsoft AD with other AWS services. Amazon Cloud Directory provides a highly scalable directory store for your application's multihierarchical data. Amazon Cognito helps you create a directory store that authenticates your users either through your own user pools or through federated identity providers. AWS Directory Service for Microsoft Active Directory (Enterprise Edition), also known as Microsoft AD, enables your directory-aware workloads and AWS resources to use a managed Active Directory [in the AWS Cloud](https://aws.amazon.com/documentation/directory-service/). Please refer to <https://aws.amazon.com/documentation/directory-service/> for more information on which directories you can choose from and additional documentation available on this service.

11.0 Service Definition – Amazon CloudFront

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

11.1 Service Overview

Amazon CloudFront is a global Content Delivery Network (CDN) service that accelerates delivery of your websites, APIs, video content, or other web assets. It integrates with other AWS products to give developers and businesses an easy way to accelerate content to end users with no minimum usage commitments.

Top benefits include:

Fast – Using a network of edge locations around the world, Amazon CloudFront caches copies of your static content close to viewers, lowering latency when they download your objects and giving you the high, sustained data transfer rates needed to deliver large popular objects to end users at scale. Requests for your dynamic content are carried back to your origin servers running in AWS (e.g., [Amazon EC2](#), [Elastic Load Balancing](#)) over optimised network paths for a more reliable and consistent experience. These network paths are constantly monitored by Amazon and connections from Amazon CloudFront edge locations to the origin are reused to serve your dynamic content from our CDN with the best possible performance.

Simple – A single API call lets you get started distributing content from your Amazon S3 bucket or Amazon EC2 instance or other origin server through the Amazon CloudFront network. Alternatively, interact with Amazon CloudFront through the AWS Management Console's simple graphical user interface. There is no need to create separate domains for your static and dynamic content. With Amazon CloudFront, you can just use the same domain name to point to all of your website content. Any changes you make to your existing configuration take effect across the entire global network within minutes. Plus, since there's no need to negotiate with a sales person, you can get started quickly and begin delivering your entire website using Amazon CloudFront.

Designed for Use with Other AWS Cloud Services – Amazon CloudFront is designed for use with other AWS Cloud services, including Amazon S3, where you can durably store the definitive versions of your static files, and Amazon EC2, where you can run your application server for dynamically generated content. If you are using Amazon S3 or Amazon EC2 as an origin server, data transferred from the origin server to edge locations (Amazon CloudFront “origin fetches”) will be billed at a lower price than Internet data transfer out of Amazon S3 or Amazon EC2. Amazon CloudFront also integrates with Elastic Load Balancing. For instance, you can deploy your web application on Amazon EC2 servers behind Elastic Load Balancing and use Amazon CloudFront to deliver your entire website.

Cost Effective – Amazon CloudFront passes on the benefits of Amazon's scale to you. You pay only for the content that you deliver through the network, without minimum commitments or up-front fees. This applies for any type of content that you deliver—static, dynamic, streaming media, or a web application with any combination of these.

Elastic – With Amazon CloudFront, you don't need to worry about maintaining expensive web server capacity to meet the demand from potential traffic spikes for your content. The service automatically responds as demand increases or decreases without any intervention from you. Amazon CloudFront also uses multiple layers of caching at each edge location and collapses simultaneous requests for the same object before contacting your origin server. These optimizations further help reduce the need to scale your origin infrastructure as your website becomes more popular.

Reliable – Amazon CloudFront is built using Amazon's highly reliable infrastructure. The distributed nature of edge locations used by Amazon CloudFront automatically routes end users to the closest available location as required by network conditions. Origin requests from the edge locations to AWS origin servers (e.g., Amazon EC2, Amazon S3) are carried over network paths that Amazon constantly monitors and optimises for both availability and performance.

Global – Amazon CloudFront uses a global network of edge locations, located near your end users in the United States, Europe, Asia, South America, and Australia.

11.2 Backup/Restore and Disaster Recovery

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

11.3 Pricing Overview

Please visit <https://aws.amazon.com/pricing/> for further details on pricing AWS services.

11.4 Service Constraints

Please see <https://aws.amazon.com/cloudfront/> for more information.

11.5 Technical Requirements

Amazon CloudFront is a web service that speeds up distribution of your static and dynamic web content, for example, .html, .css, .php, image, and media files, to end users. CloudFront delivers your content through a worldwide network of edge locations. When an end user requests content that you're serving with CloudFront, the user is routed to the edge location that provides the lowest latency, so content is delivered with the best possible performance. If the content is already in that edge location, CloudFront delivers it immediately. If the content is not currently in that edge location, CloudFront retrieves it from an Amazon S3 bucket or an HTTP server (for example, a web server) that you have identified as the source for the definitive version of your content.

Developer Guide – Provides an overview of Amazon CloudFront, detailed feature descriptions, procedures for using the console, and an explanation of the API.

[HTML](#) | [PDF](#) | [Kindle](#)

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.

[HTML](#) | [PDF](#)

12.0 Service Definition – Amazon CloudSearch

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

12.1 Service Overview

Amazon CloudSearch is a managed service in the AWS Cloud that makes it simple and cost-effective to set up, manage, and scale a search solution for your website or application. Amazon CloudSearch supports 34 languages and popular search features such as highlighting, autocomplete, and geospatial search.

With Amazon CloudSearch, you can quickly add rich search capabilities to your website or application. You don't need to become a search expert or worry about hardware provisioning, setup, and maintenance. With a few clicks in the AWS Management Console, you can create a search domain and upload the data that you want to make searchable, and Amazon CloudSearch will automatically provision the required resources and deploy a highly tuned search index.

You can easily change your search parameters, fine tune search relevance, and apply new settings at any time. As your volume of data and traffic fluctuates, Amazon CloudSearch seamlessly scales to meet your needs.

Amazon CloudSearch supports powerful search features such as:

Free text, Boolean, and Faceted search

Autocomplete suggestions

Customizable relevance ranking and query-time rank expressions

Field weighting

Geospatial search

Highlighting

Support for 34 languages

Top benefits include:

Simple – You can configure and manage an Amazon CloudSearch domain through the AWS Management Console, AWS CLI, and AWS SDKs. Simply point to a sample of your data and Amazon CloudSearch will automatically recommend how to configure your domain's indexing options. You can easily add or delete index fields and customise search options such as faceting and highlighting. Configuration changes do not require you to re-upload your data.

Scalable – Amazon CloudSearch offers powerful auto scaling for all search domains. As your data or query volume changes, Amazon CloudSearch can scale your search domain's resources up or down as needed. You can control scaling if you know that you need more capacity for bulk uploads or are expecting a surge in search traffic.

Reliable – Amazon CloudSearch provides automatic monitoring and recovery for your search domains. When Multi-AZ is enabled, Amazon CloudSearch provisions and maintains resources for a search domain in two Availability Zones to ensure high availability. Updates are automatically applied to the search instances in both Availability Zones. Search traffic is distributed across both Availability Zones and the instances in either zone are capable of handling the full load in the event of a failure.

High Performance – Amazon CloudSearch ensures low latency and high throughput performance, even at large scale through automatic sharding and horizontal and vertical auto scaling.

Fully Managed – Amazon CloudSearch is a fully managed custom search service. Hardware and software provisioning, setup and configuration, software patching, data partitioning, node monitoring, scaling, and data durability are handled for you.

Cost-Effective – Amazon CloudSearch is designed to be cost-effective. You pay low hourly rates, and only for the resources you use. Amazon CloudSearch offers low total cost of ownership for your search applications compared to operating a search environment on your own.

Secure – Amazon CloudSearch uses strong cryptographic methods to authenticate users and prevent unauthorised access to your domains. Amazon CloudSearch supports HTTPS and integrates with Identity and Access Management (IAM) to control access to the Amazon CloudSearch configuration service and each domain's document, search, and suggest services.

12.2 Backup/Restore and Disaster Recovery

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

12.3 Pricing Overview

Please see the AWS UK G-Cloud 1 Pricing Document affiliated with this service in the Digital Marketplace.

12.4 Service Constraints

Please see <https://aws.amazon.com/cloudsearch/> for more information.

12.5 Technical Requirements

Amazon CloudSearch is a fully managed service in the cloud that makes it easy to set up, manage, and scale a search solution for your website. Amazon CloudSearch

enables you to search large collections of data such as web pages, document files, forum posts, or product information. With Amazon CloudSearch, you can quickly add search capabilities to your website without having to become a search expert or worry about hardware provisioning, setup, and maintenance. As your volume of data and traffic fluctuates, Amazon CloudSearch automatically scales to meet your needs.

Developer Guide – Describes how to create, manage, and use an Amazon CloudSearch domain to implement a search solution for your website or application. Also provides instructions for using the various features through the AWS Management Console, AWS CLI, and API. [HTML](#) | [PDF](#) | [Kindle](#)

13.0 Service Definition – Amazon CloudWatch

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

This service listing includes the following AWS Services:

- Amazon CloudWatch
- Amazon CloudWatch Events
- Amazon CloudWatch Logs

13.1 Service Overview

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

Top Features include

- Collect and store logs from your resources, applications, and services in near real-time.
- Collect default metrics from more than 70 AWS services, such as Amazon EC2, Amazon DynamoDB, Amazon S3, Amazon ECS, AWS Lambda, and Amazon API Gateway, without any action on your part. Collect custom metrics from your own applications to monitor operational performance, troubleshoot issues, and spot trends.

- Amazon CloudWatch Anomaly Detection applies machine-learning algorithms to continuously analyse data of a metric and identify anomalous behaviour.
- Dashboards enable you to create re-usable graphs and visualize your cloud resources and applications in a unified view.
- Integrated with AWS Identity and Access Management (<https://aws.amazon.com/iam/>) (IAM) so that you can control which users and resources have permission to access your data and how they can access it.

Top benefits include

- **Observability on a single platform across applications and infrastructure** - Modern applications such as those running on microservices architectures generate large volumes of data in the form of metrics, logs, and events. Amazon CloudWatch enables you to collect, access, and correlate this data on a single platform from across all your AWS resources, applications, and services that run on AWS and on-premises servers, helping you break down data silos so you can easily gain system-wide visibility and quickly resolve issues.
- **Easiest way to collect metrics in AWS and on-premises** - Monitoring your AWS resources and applications is easy with CloudWatch. It natively integrates with more than 70 AWS services such as Amazon EC2, Amazon DynamoDB, Amazon S3, Amazon ECS, Amazon EKS, and AWS Lambda, and automatically publishes detailed 1-minute metrics and custom metrics with up to 1-second granularity so you can dive deep into your logs for additional context. You can also use CloudWatch in hybrid cloud architectures by using the CloudWatch Agent or API to monitor your on-premises resources.
- **Improve operational performance and resource optimization** - Amazon CloudWatch enables you to set alarms and automate actions based on either predefined thresholds, or on machine learning algorithms that identify anomalous behaviour in your metrics. For example, it can start Amazon EC2 Auto Scaling automatically, or stop an instance to reduce billing overages. You can also use CloudWatch Events for serverless to trigger workflows with services like AWS Lambda, Amazon SNS, and AWS CloudFormation.
- **Get operational visibility and insight** - To optimize performance and resource utilization, you need a unified operational view, real-time granular data, and historical reference. CloudWatch provides automatic dashboards, data with 1-second granularity, and up to 15 months of metrics storage and retention. You can also perform metric math on your data to derive operational and utilization insights; for example, you can aggregate usage across an entire fleet of EC2 instances.
- **Derive actionable insights from logs** - CloudWatch enables you to explore, analyse, and visualize your logs so you can troubleshoot operational problems with ease. With CloudWatch Logs Insights, you only pay for the queries you run. It scales with your log volume and query complexity giving you answers in seconds. In addition, you can publish log-based metrics, create alarms, and

correlate logs and metrics together in CloudWatch Dashboards for complete operational visibility.

- **Collect and aggregate container metrics and logs** - Container Insights simplifies the collection and aggregation of curated metrics and container ecosystem logs. It collects compute performance metrics such as CPU, memory, network, and disk information from each container as performance events and automatically generates custom metrics used for monitoring and alarming. The performance events are ingested as CloudWatch Logs with metadata about the running environment such as the Amazon EC2 instance ID, Service, Amazon EBS volume mount and ID, etc., to simplify monitoring and troubleshooting. CloudWatch custom metrics are automatically extracted from these ingested logs and can be further analysed using CloudWatch Logs Insights' advanced query language. Container Insights also provides an option to collect application logs (stdout/stderr), custom logs, predefined Amazon EC2 instance logs, Amazon EKS/k8s data plane logs and Amazon EKS control plane logs (<https://docs.aws.amazon.com/eks/latest/userguide/control-plane-logs.html>).
- **High resolution alarms** - Amazon CloudWatch alarms allow you to set a threshold on metrics and trigger an action. You can create high-resolution alarms, set a percentile as the statistic, and either specify an action or ignore as appropriate. For example, you can create alarms on Amazon EC2 metrics, set notifications, and take one or more actions to detect and shut down unused or underutilized instances. Real-time alarming on metrics and events enables you to minimize downtime and potential business impact.
- **Anomaly Detection** - Amazon CloudWatch Anomaly Detection applies machine-learning algorithms to continuously analyse data of a metric and identify anomalous behaviour. 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.
- **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.
- **Compliance and Security** - Amazon CloudWatch is integrated with AWS Identity and Access Management (<https://aws.amazon.com/iam/>) (IAM) so that you can control which users and resources have permission to access your data and how they can access it. Amazon CloudWatch Logs is also PCI and FedRamp compliant. Data is encrypted at rest and during transfer. You can also use AWS KMS encryption to encrypt your log groups for added compliance and security.

13.2 Backup/Restore and Disaster Recovery

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

13.3 Pricing Overview

Please visit <https://aws.amazon.com/pricing/> for further details on pricing AWS services.

13.4 Service Constraints

Please see <http://aws.amazon.com/documentation/cloudwatch/> for more information.

13.5 Technical Requirements

Please refer to <http://aws.amazon.com/documentation/cloudwatch/> and the following links for comprehensive technical documentation regarding Amazon CloudWatch.

- <https://docs.aws.amazon.com/AmazonCloudWatch/latest/monitoring/GettingStarted.html>
- <https://docs.aws.amazon.com/AmazonCloudWatch/latest/monitoring/security.html>
- <https://docs.aws.amazon.com/AmazonCloudWatch/latest/monitoring/cloudwatch-and-eventbridge.html>

14.0 Service Definition – Amazon Cognito

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

14.1 Service Overview

Amazon Cognito lets you easily add user sign-up and sign-in to your mobile and web apps. With Amazon Cognito, you can also authenticate users through social identity providers such as Facebook, Twitter, or Amazon, or by using your own identity solution. In addition, Amazon Cognito enables you to save data locally on users' devices, allowing your applications to work even when the devices are offline. You can then synchronise data across users' devices so that their app experience remains consistent regardless of the device they use.

With Amazon Cognito, you can focus on creating great app experiences instead of worrying about building, securing, and scaling a solution to handle user management, authentication, and syncing across devices.

Top benefits include:

Add User Sign-Up and Sign-In to Your Mobile Web Apps – With Amazon Cognito Identity, you can easily and securely add sign-up and sign-in functionality to your mobile

and web apps. Amazon Cognito Identity is fully managed and can scale to support hundreds of millions of users. You can also implement enhanced security features, such as email and phone number verification and multi-factor authentication. In addition, Amazon Cognito Identity lets you customise workflows through AWS Lambda, such as adding app-specific logic to user registration for fraud detection and user validation.

Federate Identities and Provide Secure Access to AWS Resources – Amazon Cognito Identity allows you to easily federate identities so that your users can sign in through social identity providers such as Facebook, Twitter, Google, and Amazon or using your own identity solution. When you use Amazon Cognito, your app is provided with temporary, limited-privilege credentials that it can use to access AWS resources. You can set granular access permissions on your AWS resources. For example, you can limit access to a folder within an Amazon S3 bucket to a particular app user or enable unauthenticated users to access a restricted set of resources. This means your app can access the resources it needs and that you can follow security best practices by not hardcoding credentials in your app. To learn more, visit our docs.

Store and Sync Across Devices – When you use Amazon Cognito Sync, your app can read and write to a local data store. This means that your app can work in the same way regardless of whether the device is online or offline. You can also save user data, such as user preferences, sign-in, and game state, and then sync this data across a user's devices to make their experience consistent across their devices. Amazon Cognito Sync resolves conflicts and intermittent network connectivity, enabling you to focus on delivering great app experiences instead of creating and managing a user data sync solution.

14.2 Backup/Restore and Disaster Recovery

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

14.3 Pricing Overview

Please visit <https://aws.amazon.com/pricing/> for further details on pricing AWS services.

14.4 Service Constraints

Please see <https://aws.amazon.com/cognito/> for more information.

14.5 Technical Requirements

Amazon Cognito lets you easily add user sign-up and sign-in to your mobile and web apps. With Amazon Cognito, you can also authenticate users through social identity providers such as Facebook, Twitter, or Amazon, or by using your own identity solution.

Developer Guide – Describes how to get started with AWS WAF and AWS Shield Advanced, explains key concepts, and provides step-by-step instructions that show you how to use the features [HTML](#)

15.0 Service Definition – Amazon Comprehend

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

15.1 Service Overview

Amazon Comprehend is a natural language processing (NLP) service that uses machine learning to find insights and relationships in text. Amazon Comprehend identifies the language of the text; extracts key phrases, places, people, brands, or events; understands how positive or negative the text is; and automatically organises a collection of text files by topic.

You can use the Amazon Comprehend APIs to analyse text and use the results in a wide range of applications including voice of customer analysis, intelligent document search, and content personalization for web applications.

The service constantly learns and improves from a variety of information sources, including Amazon.com product descriptions and consumer reviews – one of the largest natural language data sets in the world - to keep pace with the evolution of language..

Top benefits include:

- **Get better answers from your text** - Amazon Comprehend can discover the meaning and relationships in text from customer support incidents, product reviews, social media feeds, news articles, documents, and other sources. For example, you can identify the feature that's most often mentioned when customers are happy or unhappy about your product.
- **Identify topics in a collection of text** - Amazon Comprehend can analyse a collection of documents and other text files (such as social media posts) and automatically organise them by relevant terms or topics. You can then use the topics to deliver personalized content to your customers or to provide richer search and navigation. For example, if you have an extensive collection of news articles, you can automatically group them by subject matter to enable your site to suggest new articles to visitors based on what they've read previously.
- **Continually learning, always improving** - Amazon Comprehend is always learning from new data and continually gains a deeper understanding of entities, key phrases, sentiments, and topics. The learning data covers a wide range of domains including Finance, Health, Media, Telecom, Education, Government, Advertising, and many others. The service also learns from Amazon.com product descriptions and customer reviews, which is one of the largest collections of natural language in the world.
- **Integrated with Amazon S3 and AWS Glue** - Amazon Comprehend is integrated with Amazon S3 and AWS Glue. You can use Amazon S3 to store documents and texts, and directly access them from Amazon Comprehend. In addition, you can use AWS Glue to load documents and texts from various AWS

data stores (Amazon Redshift, Amazon RDS, Amazon DynamoDB, etc.) to Amazon Comprehend, and subsequently load the results from Amazon Comprehend back into these AWS data stores for downstream analytics.

15.2 Backup/Restore and Disaster Recovery

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

15.3 Pricing Overview

Please visit <https://aws.amazon.com/pricing/> for further details on pricing AWS services.

15.4 Service Constraints

Please see <https://aws.amazon.com/comprehend/> for more information.

Prior to using these Services, Buyer agrees to contact Supplier before using these Services in connection with Buyer Personal Data.

15.5 Technical Requirements

Amazon Comprehend uses natural language processing (NLP) to extract insights about the content of documents without the need of any special pre-processing. Amazon Comprehend processes any text files in UTF-8 format. It develops insights by recognizing the entities, key phrases, language, sentiments, and other common elements in a document. Use Amazon Comprehend to create new products based on understanding the structure of documents. With Amazon Comprehend you can search social networking feeds for mentions of products, scan an entire document repository for key phrases, or determine the topics contained in a set of documents.

Developer Guide - Provides a conceptual overview of Amazon Comprehend, includes detailed instructions for using the various features, and provides a complete API reference for developers. [HTML](#) | [PDF](#) | [Kindle](#)

16.0 Service Definition – Amazon Comprehend Medical

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

16.1 Service Overview

Amazon Comprehend Medical is a Natural Language Processing (NLP) service that makes it easy to use machine learning (ML) 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 doctors' notes, clinical trial reports, and patient health records.

Top Features include:

Extract medical information quickly and accurately. Powered by state-of-the-art machine learning models, Amazon Comprehend Medical understands and identifies complex medical information quickly and accurately. For example, Amazon Comprehend Medical can extract “methicillin-resistant Staphylococcus aureus,” often input as “MRSA,” and provide context, such as whether a patient has tested positive or negative, to make the extracted term meaningful.

Protect patient information. Amazon Comprehend Medical provides a number of capabilities to help healthcare providers stay compliant and protect patient data. The service is HIPAA eligible and can identify protected health information (PHI) stored in medical record systems while adhering to the standards for General Data Protection Regulation (GDPR). Amazon Comprehend Medical allows developers to implement data privacy and security solutions by extracting and then identifying relevant patient identifiers as described in HIPAA’s Safe Harbor method of de-identification. Finally, the service does not store or save any customer data.

Lower medical document processing costs. Amazon Comprehend Medical makes it easy to automate and lower the cost of processing and coding unstructured medical text from patient records, billing, and clinical indexing. It offers 2 APIs that developers can integrate into existing workflows and applications with only a few lines of code, costing a penny or less for every 100 characters of analysed text. You pay only for what you use, and there are no minimum fees.

Amazon Comprehend is a fully managed and continuously trained service. you don’t have to manage the scaling of resources, maintenance of code, or maintaining the training data.

Top benefits include:

- Perform medical cohort analysis.
- Support clinical decisions.
- Improve medical coding in revenue cycle management.
- Extract medical information quickly and accurately.
- Lower medical document processing costs.
- Protect patient information.
- Amazon Comprehend Medical is an NLP service that uses machine learning.
- Access to Amazon Comprehend Medical is gained through a simple API call. No machine learning expertise is required, and there are no complicated rules to write or models to train.
- Requests to the Amazon Comprehend Medical APIs and console are made over a secure (SSL) connection.
- You can use AWS Identity and Access Management (IAM) to control which IAM

users have access to specific Amazon Comprehend Medical actions and resources.

16.2 Backup/Restore and Disaster Recovery

This requirement is not applicable to Amazon Comprehend Medical. Access to Amazon Comprehend Medical is gained through an API call, and no data will be stored by the service.

16.3 Pricing Overview

Please visit <https://aws.amazon.com/pricing/> for further details on pricing AWS services.

16.4 Service Constraints

- Amazon Comprehend Medical does **not** store and use text inputs processed by the service.
- Amazon Comprehend Medical does not store data in any other Region.
- Amazon Comprehend Medical does not use customer data to train the models.
- Amazon Comprehend Medical is currently available in the Europe (Ireland), US East (Northern Virginia), US East (Ohio), and US West (Oregon) Regions.
- Amazon Comprehend Medical is not a substitute for professional medical advice, diagnosis, or treatment. You and your end users are responsible for exercising your and their own discretion, experience, and judgment in determining the correctness, completeness, timeliness, and suitability of any information provided by Amazon Comprehend Medical. You and your end users are solely responsible for any decisions, advice, actions, and/or inactions based on the use of Amazon Comprehend Medical.
- Amazon Comprehend Medical may not accurately identify PHI in all circumstances and does not meet the requirements for de-identification of PHI in accordance with HIPAA. You are responsible for reviewing any output provided by Amazon Comprehend Medical to ensure it meets your needs.

16.5 Technical Requirements

Please refer to <https://aws.amazon.com/comprehend/medical/> for comprehensive technical documentation regarding AWS Comprehend Medical.

17.0 Service Definition – Amazon Connect

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

17.1 Service Overview

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

Top benefits include:

- **Easy to set up and manage** – Setting up Amazon Connect is easy. With only a few clicks in the AWS Management Console, agents can begin taking calls within minutes. The drag and drop Contact Flow Editor allows you to create contact flows that define effective and efficient customer interactions without any coding.
- **Scalable and elastic** – There is no infrastructure to deploy or manage, so you can scale your Amazon Connect contact centre up or down, on-boarding up to tens of thousands of agents in response to business cycles and paying only for the time callers are interacting with Amazon Connect plus any associated telephony charges.
- **Pay as you go** – As an on-demand service, you pay for Amazon Connect usage by the minute with no long-term commitments or upfront charges. There is no minimum monthly fee. You are charged based on the number of minutes you use Amazon Connect to engage with your end customers, at the specified per minute rate. Pricing is not based on capacity, agent seats, or maintenance.
- **Reliable** – Amazon Connect runs on AWS-proven [infrastructure](#) operating 42 Availability Zones within 16 geographic regions around the world. This makes Amazon Connect more highly available, fault tolerant, and scalable than would be possible if a contact centre solution was run from a single data centre.
- **Open platform** – Amazon Connect is an open platform so it is easy to integrate with other systems such as Customer Relationship Management (CRM) solutions or the AWS platform. For example, you can use [AWS Lambda](#) to run code for a serverless application or backend service to build contact flow experiences that adapt to your customer needs in real time. And [Amazon S3](#) can be used to store call recordings, where you can control access and lifecycle of the data.

17.2 Backup/Restore and Disaster Recovery

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

17.3 Pricing Overview

Please visit <https://aws.amazon.com/pricing/> for further details on pricing AWS services.

17.4 Service Constraints

Amazon Connect features and default service limits are found at <https://aws.amazon.com/connect/faqs/>.

17.5 Technical Requirements

Amazon Connect is a Contact Centre as a Service (CCaS) solution that offers easy, self-service configuration and enables dynamic, personal, and natural customer engagement at any scale.

- **Administrator Guide** - Describes an Amazon Connect virtual contact centre and helps administrators add users, choose telephony options, manage data storage, and configure encryption settings. [HTML](#) | [PDF](#)
- **User Guide** - Describes key concepts of Amazon Connect and provides instructions for managing, measuring, and scaling a virtual contact centre. [HTML](#) | [PDF](#)

18.0 Service Definition – Amazon DocumentDB (with MongoDB Compatibility)

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

18.1 Service Overview

Amazon DocumentDB (with MongoDB compatibility) is a fast, scalable, highly available, and fully managed document database service that supports MongoDB workloads.

Customers use MongoDB as a document database to store, retrieve, and manage semi-structured data. However, it is hard to build performant, highly available applications that can quickly scale to multiple terabytes and hundreds of thousands of reads and writes per second because of the complexity that comes with setting up and managing MongoDB clusters at scale.

Top features include:

- **MongoDB 3.6 compatible** – Amazon DocumentDB is compatible with MongoDB 3.6 drivers and tools.
- **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 preconfigured

with parameters and settings appropriate for the instance class you have selected.

- **Automatic software patching** – Amazon DocumentDB will keep your database up to date with the latest patches. You can control if and when your cluster is patched via Database Engine Version Management.
- **High throughput and low latency for document queries** – Amazon DocumentDB has a flexible document model, data types, and efficient indexing, and it uses a scale-up, in-memory optimised architecture to allow for fast query evaluation over large sets of documents.
- **Storage that automatically scales** – Amazon DocumentDB will automatically grow the size of your storage volume as your cluster storage needs grow. Your storage volume will grow in increments of 10 GB up to a maximum of 64 TB.

Top benefits include:

- **MongoDB-compatible** – Amazon DocumentDB implements the Apache 2.0 open source MongoDB 3.6 API by emulating the responses that a MongoDB client expects from a MongoDB server, allowing you to use your existing MongoDB drivers and tools with Amazon DocumentDB.
- **Highly available** – Amazon DocumentDB is designed for 99.99% availability and replicates six copies of your data across three AWS Availability Zones.
- **Performance at scale** – Amazon DocumentDB uses a distributed, fault-tolerant, self-healing storage system that automatically scales up to 64 TB per database cluster.
- **Reduced database I/O** – Amazon DocumentDB reduces database I/O by writing only database changes to the storage layer, avoiding slow, inefficient, and expensive data replication across network links.
- **Highly secure** – Amazon DocumentDB provides multiple levels of security for your database, including network isolation using [Amazon VPC](#), encryption at rest using keys you create and control through [AWS KMS](#), and encryption in transit using Transport Layer Security (TLS).
- **Fully managed** – With Amazon DocumentDB, you don't need to worry about database management tasks, such as hardware provisioning, patching, setup, configuration, or backups.
- **Migration support** – Customers can easily migrate their MongoDB databases on premises or on Amazon EC2 to Amazon DocumentDB for free (for six months per instance) with virtually no downtime using the AWS Database Migration Service (AWS DMS).
- **Monitoring and metrics** – Amazon DocumentDB provides [Amazon CloudWatch](#) metrics for your database instances. You can use the AWS Management Console to view over 20 key operational metrics for your cluster, including compute, memory, storage, query throughput, and active connections.

- **Easy scaling of database compute resources** – With a few clicks in the [AWS Management Console](#), you can scale compute and memory resources, powering your cluster 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.
- **Compliance certification** – In addition to being [HIPAA eligible](#), Amazon DocumentDB complies with [PCI DSS](#); [ISO 9001](#), [27001](#), [27017](#), and [27018](#); and [SOC 2](#).

18.2 Backup/Restore and Disaster Recovery

Automated backups are always enabled on Amazon DocumentDB clusters. You can increase your backup window for point-in-time restores up to 35 days. Amazon DocumentDB automatically maintains six copies of your data across three Availability Zones and will automatically attempt to recover your instance in a healthy zone 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 10 GB segments spread across many disks. Each 10 GB chunk of your storage volume is replicated six ways, across three Availability Zones. 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.

Unlike other databases, after a database crash Amazon DocumentDB does not need to replay the redo log from the last database checkpoint (typically five minutes) and confirm that all changes have been applied before making the database available for operations. This reduces database restart times to less than 60 seconds in most cases. Amazon DocumentDB moves the cache out of the database process and makes it available immediately at restart time. This prevents you from having to throttle access until the cache is repopulated to avoid brownouts.

18.3 Pricing Overview

Please visit <https://aws.amazon.com/pricing/> for further details on pricing AWS services.

18.4 Service Constraints

Please see <https://docs.aws.amazon.com/documentdb/latest/developerguide/limits.html> for more information.

18.5 Technical Requirements

Please refer to <https://docs.aws.amazon.com/documentdb/latest/developerguide/what-is.html> for comprehensive technical documentation regarding Amazon DocumentDB.

19.0 Service Definition – Amazon DynamoDB

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

19.1 Service Overview

Amazon DynamoDB is a key-value and document database that delivers single-digit millisecond performance at any scale. It's a fully managed, multi-Region, multi-master database with built-in security, backup and restore, and in-memory caching for internet-scale applications. DynamoDB can handle more than 10 trillion requests per day and support peaks of more than 20 million requests per second.

Top features include:

- DynamoDB Accelerator (DAX) is an in-memory cache that delivers fast read performance for your tables at scale by enabling you to use a fully managed in-memory cache.
- 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 KMS.
- DynamoDB global tables replicate your data automatically across your choice of AWS Regions and automatically scale capacity to accommodate your workloads.
- DynamoDB provides native, server-side support for transactions, simplifying the developer experience of making coordinated, all-or-nothing changes to multiple items both within and across tables.
- Amazon DynamoDB Streams captures a time-ordered sequence of item-level modifications in any DynamoDB table and stores this information in a log for up to 24 hours.

Top benefits include:

- **High performance and scalability** – DynamoDB is a key-value and document database that can support tables of virtually any size with horizontal scaling. This enables DynamoDB to scale to more than 10 trillion requests per day, with peaks greater than 20 million requests per second, over petabytes of storage.
- **High availability and durability** – DynamoDB global tables replicate your data automatically across your choice of AWS Regions and automatically scale capacity to accommodate your workloads. With global tables, your globally distributed applications can access data locally in the selected Regions to get single-digit millisecond read and write performance.
- **Fully managed** – With DynamoDB, there are no servers to provision, patch, or

manage and no software to install, maintain, or operate. DynamoDB automatically scales tables up and down to adjust for capacity and maintain performance. Availability and fault tolerance are built in, eliminating the need to architect your applications for these capabilities. DynamoDB provides both on-demand and provisioned capacity modes so that you can optimise costs by specifying capacity per workload or by paying for only the resources you consume.

- **Secure** – DynamoDB encryption at rest provides an additional layer of data protection by securing your data in the encrypted table, including its primary key, local and global secondary indexes, streams, global tables, backups, and DAX clusters whenever the data is stored in durable media. Organisational policies, industry or government regulations, and compliance requirements often require the use of encryption at rest to increase the data security of your applications.
- **Integration** – DynamoDB integrates with AWS Lambda to provide triggers. Using triggers, you can automatically execute a custom function when item-level changes in a DynamoDB table are detected. With triggers, you can build applications that react to data modifications in DynamoDB tables. The Lambda function can perform any actions you specify, such as sending a notification or initiating a workflow.

19.2 Backup/Restore and Disaster Recovery

Amazon DynamoDB provides on-demand backup capability. It allows you to create full backups of your tables for long-term retention and archival for regulatory compliance needs. You can back up and restore your DynamoDB table data anytime with a single click in the AWS Management Console or with a single API call. Backup and restore actions execute with zero impact on table performance or availability.

On-demand backup and restore scales without degrading the performance or availability of your applications. It uses a new and unique distributed technology that allows you to complete backups in seconds, regardless of table size. You can create backups that are consistent within seconds across thousands of partitions, without worrying about schedules or long-running backup processes. All backups are catalogued, easily discoverable, and retained until explicitly deleted.

Point-in-time recovery (PITR) helps protect your DynamoDB tables from accidental write or delete operations. PITR provides continuous backups of your DynamoDB table data, and you can restore that table to any point in time—up to the second—during 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.

19.3 Pricing Overview

Please visit <https://aws.amazon.com/pricing/> for further details on pricing AWS services.

19.4 Service Constraints

Please see

<https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/Limits.html> for more information.

19.5 Technical Requirements

Please see

<https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/Introduction.html> for comprehensive technical documentation regarding Amazon DynamoDB.

20.0 Service Definition – Amazon Elastic Block Store (Amazon EBS)

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

20.1 Service Overview

Amazon EBS provides persistent, available, and durable block-level storage volumes for use with Amazon EC2 instances in the AWS Cloud. Amazon EBS volumes offer the consistent and low-latency performance needed to run your workloads. With Amazon EBS, you can scale your usage up or down within minutes, all while paying a low price for only what you provision.

Top benefits include:

- **Reliable, Secure Storage** – Each Amazon EBS volume is automatically replicated within its Availability Zone to protect you from component failure.
- **Consistent and Low-Latency Performance** – Amazon EBS General Purpose volumes and Amazon EBS Provisioned IOPS volumes deliver lowlatency through SSD technology and consistent I/O performance scaled to the needs of your application.
- **Backup, Restore, and Innovate** – Back up your data by taking point-in-time snapshots of your Amazon EBS volumes. Boost the agility of your business by using Amazon EBS snapshots to create new Amazon EC2 instances.
- **Geographic Flexibility** – Amazon EBS provides the ability to copy snapshots across AWS Regions, enabling geographical expansion, data centre migration, and disaster recovery.
- **Quickly Scale Up and Easily Scale Down** – Increase or decrease block storage and performance within minutes, enjoying the freedom to adjust as your needs evolve.

20.2 Backup/Restore and Disaster Recovery

An Amazon EBS snapshot is a point-in-time backup copy of an Amazon EBS volume that is stored in Amazon S3. Snapshots are incremental backups, which means that only the blocks on the device that have changed after your most recent snapshot are saved. When you delete a snapshot, only the data exclusive to that snapshot is removed. Active snapshots contain all of the information needed to restore your data (from the time the snapshot was taken) to a new Amazon EBS volume.

When you create a new Amazon EBS volume, you can create it based on an existing snapshot; the new volume begins as an exact replica of the original volume that was used to create the snapshot. New volumes created from existing Amazon S3 snapshots load lazily in the background, so you can begin using them right away. If your instance accesses a piece of data that hasn't yet been loaded, the volume immediately downloads the requested data from Amazon S3 and then continues loading the rest of the volume's data in the background. For more information about creating snapshots, see the [Creating an Amazon EBS Snapshot](#) user guide.

You can share your snapshots with specific individuals or make them public to share them with the entire AWS community. Users with access to your snapshots can create their own Amazon EBS volumes from your snapshot, but your snapshots remain completely intact. For more information about how to share snapshots, see the [Sharing Snapshots](#) user guide.

Amazon EBS snapshots are constrained to the region in which they are created. Once you have created a snapshot of an Amazon EBS volume, you can use it to create new volumes in the same region. For more information, see the [Restoring an Amazon EBS Volume from a Snapshot](#) user guide. You can also copy snapshots across AWS Regions, making it easier to leverage multiple AWS Regions for geographical expansion, data centre migration, and disaster recovery. You can copy any accessible snapshots that are in the "available" status. For more information, see the [Copying an Amazon EBS Snapshot](#) user guide.

20.3 Pricing Overview

Please visit <https://aws.amazon.com/pricing/> for further details on pricing AWS services.

20.4 Service Constraints

Please see <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/AmazonEBS.html> for more information.

20.5 Technical Requirements

Please see <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/AmazonEBS.html> and the following links for comprehensive technical documentation regarding Amazon EBS:

- [Creating an Amazon EBS Snapshot](#)

- [Deleting an Amazon EBS Snapshot](#)
- [Copying an Amazon EBS Snapshot](#)
- [Describing Snapshots](#)
- [Sharing Snapshots](#)

21.0 Service Definition – Amazon Elastic Compute Cloud (Amazon EC2)

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

This service listing includes the following AWS Services:

- Amazon Elastic Compute Cloud (EC2)
- Amazon EC2 Auto Scaling
- Amazon EC2 Image Builder

21.1 Service Overview

Amazon EC2 is a web service that provides resizable compute capacity in the cloud. It is designed to make web-scale cloud computing easier for developers. Amazon EC2's simple interface allows you to obtain and configure capacity with minimal friction. It provides you with complete control of your computing resources and lets you run on Amazon's proven computing environment. Amazon EC2 reduces the time required to obtain and boot new server instances to minutes, allowing you to quickly scale capacity, both up and down, as your computing requirements change. Amazon EC2 changes the economics of computing by allowing you to pay only for capacity that you actually use. Amazon EC2 provides developers the tools to build failure resilient applications and isolate them from common failure scenarios.

Top benefits include:

- **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.
- **Completely Controlled** – You have complete control of your instances. You have root access to each one, and you can interact with them as you would any machine.
- **Flexible Cloud Hosting Services** – You have the choice of multiple instance types, operating systems, and software packages. Amazon EC2 allows you to select a configuration of memory, CPU, instance storage, and the boot partition size that is optimal for your choice of operating system and application.
- **Designed for Use with Other AWS Cloud Services** – Amazon EC2 works in conjunction with Amazon S3, Amazon RDS, and other AWS Cloud services to

provide a complete solution for computing, query processing, and storage across a wide range of applications.

- **Reliable** – Amazon EC2 offers a highly reliable environment where replacement instances can be rapidly and predictably commissioned. The service runs within Amazon’s proven network infrastructure and data centres.
- **Secure** – Amazon EC2 works in conjunction with Amazon VPC to provide security and robust networking functionality for your compute resources.

Lockdown security model prohibits administrative access, eliminating the possibility of human error and tampering. With the AWS Nitro System virtualization resources are offloaded to dedicated hardware and software minimizing the attack surface. AWS supports 89 security standards and compliance certifications including PCI-DSS, HIPAA/HITECH, FedRAMP, GDPR, FIPS 140-2, and NIST 800-171.

- **Inexpensive** – Amazon EC2 passes on to you the financial benefits of Amazon’s scale. You pay a very low rate for the compute capacity you actually consume.

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

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

usage. Savings Plans offer significant savings over On Demand, just like EC2 Reserved Instances, in exchange for a commitment to use a specific amount of compute power (measured in \$/hour) for a one or three year period. You can sign up for Savings Plans for a 1- or 3-year term and easily manage your plans by taking advantage of recommendations, performance reporting and budget alerts in the AWS Cost Explorer.

21.1.1 Instance types

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

21.1.1.1 General Purpose

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

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

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

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

- Burstable CPU, governed by CPU Credits, and consistent baseline performance
- Unlimited mode by default to ensure performance during peak periods and Standard mode option for a predictable monthly cost
- Powered by the AWS Nitro System, a combination of dedicated hardware and lightweight hypervisor

- [AWS Nitro System](#) and high frequency Intel Xeon Scalable processors result in up to a 30% price performance improvement over T2 instances

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

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

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

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

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

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

M6g - [Amazon EC2 M6g](#) instances are powered by Arm-based AWS Graviton2 processors. They deliver up to 40% better price/performance over current generation

M5 instances and offer a balance of compute, memory, and networking resources for a broad set of workloads.

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

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

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

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

- AMD EPYC 7000 series processors with an all core turbo clock speed of 2.5 GHz
- Up to 20 Gbps network bandwidth using Enhanced Networking
- Requires HVM AMIs that include drivers for ENA and NVMe
- Powered by the [AWS Nitro System](#), a combination of dedicated hardware and lightweight hypervisor
- Instance storage offered via EBS or NVMe SSDs that are physically attached to the host server
- With M5ad instances, local NVMe-based SSDs are physically connected to the host server and provide block-level storage that is coupled to the lifetime of the M5a instance

M5n - [M5 instances](#) are ideal for workloads that require a balance of compute, memory, and networking resources including web and application servers, small and mid-sized databases, cluster computing, gaming servers, and caching fleet. The higher bandwidth,

M5n and M5dn, instance variants are ideal for applications that can take advantage of improved network throughput and packet rate performance.

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

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

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

21.1.1.2 **Compute-Optimised**

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

- C5 instances offer a choice of processors based on the size of the instance.
- New C5 and C5d 12xlarge, 24xlarge, and metal instance sizes feature custom 2nd generation Intel Xeon Scalable Processors (Cascade Lake) with a sustained all core Turbo frequency of 3.6GHz and single core turbo frequency of up to 3.9GHz.
- Other C5 instance sizes will launch on the 2nd generation Intel Xeon Scalable Processors (Cascade Lake) or 1st generation Intel Xeon Platinum 8000 series

(Skylake-SP) processor with a sustained all core Turbo frequency of up to 3.4GHz, and single core turbo frequency of up to 3.5 GHz.

- New larger 24xlarge instance size offering 96 vCPUs, 192 GiB of memory, and optional 3.6TB local NVMe-based SSDs
- Requires HVM AMIs that include drivers for ENA and NVMe
- With C5d instances, local NVMe-based SSDs are physically connected to the host server and provide block-level storage that is coupled to the lifetime of the C5 instance
- Elastic Network Adapter (ENA) provides C5 instances with up to 25 Gbps of network bandwidth and up to 14 Gbps of dedicated bandwidth to Amazon EBS.
- Powered by the [AWS Nitro System](#), a combination of dedicated hardware and lightweight hypervisor

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

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

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

- High frequency Intel Xeon E5-2666 v3 (Haswell) processors optimized specifically for EC2

- Default EBS-optimized for increased storage performance at no additional cost
- Higher networking performance with Enhanced Networking supporting Intel 82599 VF
- Requires Amazon VPC, Amazon EBS and 64-bit HVM AMIs

21.1.1.3 *Memory Optimised*

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

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

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

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

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

- 2nd generation Intel Xeon Scalable Processors (Cascade Lake) with a sustained all-core Turbo CPU frequency of 3.1 GHz and maximum single core turbo frequency of 3.5 GHz
- Support for the new Intel Vector Neural Network Instructions (AVX-512 VNNI) which will help speed up typical machine learning operations like convolution, and automatically improve inference performance over a wide range of deep learning workloads
- 25 Gbps of peak bandwidth on smaller instance sizes
- 100 Gbps of network bandwidth on the largest instance size
- Requires HVM AMIs that include drivers for ENA and NVMe
- Powered by the [AWS Nitro System](#), a combination of dedicated hardware and lightweight hypervisor
- Instance storage offered via EBS or NVMe SSDs that are physically attached to the host server
- With R5dn instances, local NVMe-based SSDs are physically connected to the host server and provide block-level storage that is coupled to the lifetime of the R5 instance

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

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

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

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

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

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

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

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

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

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

21.1.1.4 **Accelerated Computing Instances**

P3 - [P3 instances](#) are the latest generation of general purpose GPU instances. Features include:

- Up to 8 NVIDIA Tesla V100 GPUs, each pairing 5,120 CUDA Cores and 640 Tensor Cores
- High frequency Intel Xeon E5-2686 v4 (Broadwell) processors for p3.2xlarge, p3.8xlarge, and p3.16xlarge.
- High frequency 2.5 GHz (base) Intel Xeon P-8175M processors for p3dn.24xlarge.

- Supports NVLink for peer-to-peer GPU communication
- Provides up to 100 Gbps of aggregate network bandwidth.
- EFA support on p3dn.24xlarge instances

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

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

Inf1 - [Amazon EC2 Inf1](#) instances are built from the ground up to support machine learning inference applications.

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

G4 - [G4 instances](#) are designed to help accelerate machine learning inference and graphics-intensive workloads.

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

G3 – [G3 instances](#) are optimized for graphics-intensive applications. Features include:

- High frequency Intel Xeon E5-2686 v4 (Broadwell) processors
- NVIDIA Tesla M60 GPUs, each with 2048 parallel processing cores and 8 GiB of video memory
- Enables NVIDIA GRID Virtual Workstation features, including support for 4 monitors with resolutions up to 4096x2160. Each GPU included in your instance

is licensed for one "Concurrent Connected User"

- Enables NVIDIA GRID Virtual Application capabilities for application virtualization software like Citrix XenApp Essentials and VMware Horizon, supporting up to 25 concurrent users per GPU
- Each GPU features an on-board hardware video encoder designed to support up to 10 H.265 (HEVC) 1080p30 streams and up to 18 H.264 1080p30 streams, enabling low-latency frame capture and encoding, and high-quality interactive streaming experiences
- Enhanced Networking using the Elastic Network Adapter (ENA) with 25 Gbps of aggregate network bandwidth within a Placement Group

F1 – [F1 instances](#) offer customizable hardware acceleration with field programmable gate arrays (FPGAs). Features include:

Instances Features:

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

FPGA Features:

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

21.1.1.5 **Storage-Optimised**

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

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

I3en - [This instance family](#) provides dense Non-Volatile Memory Express (NVMe) SSD instance storage optimized for low latency, high random I/O performance, high sequential disk throughput, and offers the lowest price per GB of SSD instance storage

on Amazon EC2. I3en also offers Bare Metal instances (i3en.metal), powered by the Nitro System, for non-virtualized workloads, workloads that benefit from access to physical resources, or workloads that may have license restrictions.

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

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

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

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

- Powered by 2.3 GHz Intel® Xeon® E5 2686 v4 processors (codenamed Broadwell)
- Up to 16TB of HDD storage
- High disk throughput
- ENA enabled Enhanced Networking up to 25 Gbps

21.1.1.6 *Previous Generation Instances*

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

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

21.1.1.7 *Instance Features*

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

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

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

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

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

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

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

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

- **Bare Metal Instances** – Amazon EC2 bare metal instances provide your applications with direct access to the Intel® Xeon® Scalable processor and memory resources of the underlying server. Bare metal instances come up with up to 448 vCPU (u-xttb1.metal), and up to 24TB RAM (u-24tb1.metal).

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

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

Many Amazon EC2 instances can also include storage from disks that are physically attached to the host computer. This disk storage is referred to as

instance store. Instance store provides temporary block-level storage for Amazon EC2 instances. The data on an instance store volume persists only during the life of the associated Amazon EC2 instance.

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

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

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

21.2 Backup/Restore and Disaster Recovery

Traditional enterprise backup and recovery strategies typically take an agent-based approach whereby the entire contents of a server are backed up over either the Local Area Network (LAN) or the Storage Area Network (SAN). Traditional architectures have required this approach because replacing failed components is complex, time-consuming, and operationally intensive. This has, in turn, created a backup environment that is complex to manage and resource-intensive to operate—requiring technologies such as data de-duplication and virtual tape libraries to cope with ever-increasing workloads.

Amazon EC2 enables the full backup and recovery of a standard server, such as a web server or application server, so you can focus on protecting configuration and stateful data rather than on the server itself. This set of data is much smaller than the aggregate set of server data, which typically includes various application files, operating system files, temporary files, and so on. This change of approach means that regular nightly incremental or weekly full backups can take far less time and consume less storage space.

When a compute instance is started in Amazon EC2, it is based upon an Amazon Machine Image (AMI) and can connect to existing storage volumes (e.g., Amazon EBS). In addition, when launching a new instance, it is possible to pass user data to the instance so that it can be accessed internally as dynamic configuration parameters.

21.3 Pricing Overview

Please see the AWS UK G Cloud 12 Pricing Document affiliated with this service in the Digital Marketplace.

21.4 Service Constraints

Please see <http://aws.amazon.com/documentation/ec2/> for more information on service constraints for Amazon EC2.

21.5 Technical Requirements

Please refer to <http://aws.amazon.com/documentation/ec2/> and the following links for comprehensive technical documentation regarding Amazon EC2.

- **Linux Guide** – Describes key concepts of Amazon EC2 and provides instructions for using the features of Amazon EC2. Available in [HTML](#), [PDF](#), and

[Kindle](#) formats.

- **CLI Reference** – Documents the Amazon EC2 CLI. Available in [HTML](#) and [PDF](#) formats.
- **Amazon EC2 Section of the AWS CLI Reference** – Describes the AWS CLI commands that you can use to administer Amazon EC2. Provides syntax, options, and usage examples for each command. Available in [HTML](#) format.
- **Windows Guide** – Describes key concepts for Amazon EC2 and provides instructions for launching and using your Windows instance. Available in [HTML](#), [PDF](#), and [Kindle](#) format.
- **API Reference** – Documents the Amazon EC2 Query API. Available in [HTML](#) and [PDF](#) format.
- **Amazon EC2 Simple Systems Manager (SSM) API Reference** – Documents the SSM API. Available in [HTML](#) and [PDF](#) format.

22.0 Service Definition – Amazon Elastic Container Registry (Amazon ECR)

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

22.1 Service Overview

Amazon Elastic Container Registry (Amazon ECR) is a fully managed [Docker](#) container registry that makes it easy for developers to store, manage, and deploy Docker container images. Amazon ECR is integrated with [Amazon Elastic Container Service \(Amazon ECS\)](#), 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 reliably deploy containers for your applications.

Top features include:

- Amazon ECR supports Docker Registry HTTP API V2, allowing you to use Docker CLI commands (e.g., push, pull, list, tag) or your preferred Docker tools, maintaining your existing development workflow.
- The service is easily accessible from cloud, on-premises, or local machine Docker environments.
- The service is integrated with Amazon ECS, allowing you to easily store, run, and manage container images for applications running on Amazon ECS.
- Your data is redundantly stored in Amazon S3 across multiple facilities and multiple devices in each facility.
- 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 uses IAM to control and monitor who and what (e.g., EC2 instances) can access your container images.
- Organise your repositories into namespaces based on your team's existing workflows, and set which API actions can be performed on your repository through resource-level policies.

Top benefits include:

- Amazon ECR eliminates the need to operate and scale the infrastructure required to power your container registry, with no software to install and manage or infrastructure to scale.
- The service transfers your container images over HTTPS and automatically encrypts your images at rest. You can configure policies to manage permissions and control access to your images using IAM users and roles without having to manage credentials directly on your EC2 instances.
- Highly scalable, redundant, and durable architecture. Your container images are highly available and accessible, allowing you to reliably deploy new containers for your applications.
- Amazon ECR integrates with Amazon ECS and the Docker CLI, simplifying development and production workflows.
- You pay only for the amount of data you store in your repositories and data transferred to the internet.

22.2 Backup/Restore and Disaster Recovery

This requirement is not applicable to Amazon ECR. For additional information beyond what is described herein, please refer to <http://aws.amazon.com/documentation/>.

22.3 Pricing Overview

Please visit <https://aws.amazon.com/pricing/> for further details on pricing AWS services.

22.4 Service Constraints

Please see

https://docs.aws.amazon.com/AmazonECR/latest/userguide/service_limits.html for more information.

22.5 Technical Requirements

Please refer to <https://docs.aws.amazon.com/AmazonECR/latest/userguide/what-is-ecr.html> and the following links for comprehensive technical documentation regarding Amazon Elastic Container Registry:

https://docs.aws.amazon.com/AmazonECR/latest/userguide/ECR_GetStarted.html

<https://docs.aws.amazon.com/AmazonECR/latest/userguide/Registries.html>

<https://docs.aws.amazon.com/AmazonECR/latest/userguide/Repositories.html>

23.0 Service Definition – Amazon Elastic Container Service (Amazon ECS)

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

23.1 Service Overview

Amazon ECS is a highly scalable, high-performance container management service that supports [Docker](#) containers and allows you to easily run applications on a managed cluster of [Amazon EC2](#) instances. Amazon ECS eliminates the need for you to install, operate, and scale your own cluster management infrastructure. With simple API calls, you can launch and stop Docker-enabled applications, query the complete state of your cluster, and access many familiar features like security groups, [Elastic Load Balancing](#), [Amazon EBS](#) volumes, and [AWS IAM](#) roles. You can use Amazon ECS to schedule the placement of containers across your cluster based on your resource needs and availability requirements. You can also integrate your own scheduler or third-party schedulers to meet business or application specific requirements.

Top benefits include:

- **Easily Manage Clusters for Any Scale** – Amazon ECS eliminates the need to operate cluster management software or design fault-tolerant cluster architectures. There is no software to install, manage, or scale, allowing you to focus on building Dockerised applications.
- **Flexible Container Placement** – You can use Amazon ECS to schedule long running applications and batch jobs. You can also use the APIs to get up-to-date cluster state information to integrate your own custom, as well as third-party, schedulers. Amazon ECS is a shared-state, optimistic concurrency system that supports multiple schedulers on the same cluster for each business or application-specific requirement.
- **Designed for Use with Other AWS Cloud Services** – Amazon ECS is integrated with Elastic Load Balancing, Amazon EBS, Amazon VPC, IAM, and AWS CloudTrail, providing you a complete solution for running a wide range of containerised applications or services.
- **Extensible** – Amazon ECS can easily be integrated or extended through simple APIs. Amazon ECS provides complete visibility and control into your AWS resources, so you can easily integrate and use your own schedulers or connect Amazon ECS into your existing software delivery process (e.g., continuous integration and delivery systems).
- **Performance at Scale** – Amazon ECS is built on technology developed from

many years of experience running highly scalable services. You can launch tens or tens of thousands of Docker containers in seconds using Amazon ECS with no additional complexity.

- **Secure** – Amazon ECS launches your containers on your own Amazon EC2 instances. No compute resources are shared with other customers. Your clusters run in an Amazon VPC, allowing you to use your own Amazon VPC security groups and network Access Control Lists (ACLs). These features provide you a high level of isolation and help you use Amazon ECS to build highly secure and reliable applications.

23.2 Backup/Restore and Disaster Recovery

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

23.3 Pricing Overview

Please visit <https://aws.amazon.com/pricing/> for further details on pricing AWS services.

23.4 Service Constraints

Please see <https://aws.amazon.com/ecs/> for more information.

23.5 Technical Requirements

Amazon ECS is a highly scalable, fast, container management service that makes it easy to run, stop, and manage Docker containers on a cluster of Amazon EC2 instances.

- **Developer Guide** – Describes key concepts of Amazon ECS and provides instructions for using the features of Amazon ECS. [HTML](#) | [PDF](#) | [Kindle](#)
- **API Reference** – Documents the Amazon ECS Query API. [HTML](#) | [PDF](#)
- **ECS section of AWS CLI Reference** – Documents the Amazon ECS commands available in the AWS CLI. [HTML](#)

24.0 Service Definition – Amazon Elastic Container Service for Kubernetes (Amazon EKS)

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

24.1 Service Overview

Amazon Elastic Container Service for Kubernetes (Amazon EKS) is a managed service that makes it easy for customers to run Kubernetes on AWS without needing to install and operate their own Kubernetes clusters. With Amazon EKS, upgrades and high availability are managed by AWS. Amazon EKS runs three Kubernetes masters across

three Availability Zones in order to ensure high availability. Amazon EKS automatically detects and replaces unhealthy masters, and it provides automated version upgrades and patching for the masters.

Top benefits include:

- **No Masters to Manage** - Amazon EKS runs master nodes across multiple AWS Availability Zones, automatically detects and replaces unhealthy masters, and provides on-demand upgrades and patching. You simply provision worker nodes and connect them to the AWS cloud.
- **Secure by Default** - Secure and encrypted communication channels are automatically setup between your worker nodes and the EKS-managed masters, making your infrastructure running on Amazon EKS secure by default.
- **Conformant and Compatible** - We run the latest upstream version of Kubernetes and are Certified Kubernetes Conformant, so applications running on Amazon EKS are fully compatible with applications running on any standard Kubernetes environment.
- **Built with the Community** - We actively work with the community, making contributions to the Kubernetes code base that help Amazon EKS users take advantage of AWS services and features.

24.2 Backup/Restore and Disaster Recovery

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

24.3 Pricing Overview

Please visit <https://aws.amazon.com/pricing/> for further details on pricing AWS services.

24.4 Service Constraints

Amazon EKS features and default service limits are found at <https://aws.amazon.com/eks/faq/>.

24.5 Technical Requirements

Please refer to <http://aws.amazon.com/documentation/> for more details.

25.0 Service Definition – Amazon Elastic File System (Amazon EFS)

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

25.1 Service Overview

Amazon EFS is a file storage service for Amazon EC2 instances. Amazon EFS is easy to use and provides a simple interface that allows you to create and configure file systems quickly and easily. With Amazon EFS, storage capacity is elastic, growing and shrinking automatically as you add and remove files, so your applications have the storage they need, when they need it.

Amazon EFS supports the Network File System version 4 (NFSv4) protocol, so the applications and tools that you use today work seamlessly with Amazon EFS. Multiple Amazon EC2 instances can access an Amazon EFS file system at the same time, providing a common data source for workloads and applications running on more than one instance.

With Amazon EFS, you pay only for the storage used by your file system. You don't need to provision storage in advance and there is no minimum fee or setup cost. Amazon EFS is designed for a wide variety of use cases like content repositories, development environments, and home directories. With on-demand scaling and performance, Amazon EFS is an ideal solution for big data applications.

Top benefits include:

- **Seamless Integration** – Amazon EFS supports the NFSv4 protocol, so the applications and tools that you use today work seamlessly with Amazon EFS. Amazon EFS file systems support standard file system semantics, and [Amazon EC2](#) instances mount Amazon EFS file systems using standard file system mount tools.
- **Scale Up and Down Seamlessly** – Amazon EFS automatically scales your file system storage capacity up or down as you add or remove files without disrupting your applications, ensuring you always have the storage you need while reducing time-consuming administration tasks.
- **Fully Managed Service** – Amazon EFS has a simple interface that allows you to create and configure file systems quickly and easily. The service manages all the file storage infrastructure for you, avoiding the complexity of deploying, patching, and maintaining complex file system deployments.
- **Share File Storage Across Instances** – Multiple Amazon EC2 instances can access an Amazon EFS file system, so applications that scale beyond a single instance can access a file system. Amazon EC2 instances running in multiple Availability Zones within the same region can access the file system, so that many users can access and share a common data source.
- **Consistent, Scalable Performance** – Amazon EFS is SSD-based and is designed to provide the throughput, IOPS, and low latency needed for a broad range of workloads. With Amazon EFS, throughput and IOPS scale as a file system grows, and file operations are delivered with consistent, low latencies.
- **Low Cost** – Amazon EFS provides the capacity you need, when you need it, without having to provision storage in advance. You pay for what you use, with

no minimum commitments or up-front fees.

- **Highly Available and Durable** – Amazon EFS is designed to be durable and highly available. Each Amazon EFS file system object (i.e., directory, file, and link) is redundantly stored across multiple Availability Zones.
- **Secure** – Amazon EFS allows you to tightly control access to your file systems. Only Amazon EC2 instances within the [Amazon VPC](#) you specify can directly access your Amazon EFS file systems. Amazon VPC security groups and network access control lists allow you to manage network access to your Amazon EFS file systems. At the file and directory level, Amazon EFS supports user and group read/write/execute permissions. The service is also integrated with [AWS IAM](#), which can be used to control access to Amazon EFS APIs as well as manage resource-level permissions.

25.2 Backup/Restore and Disaster Recovery

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

25.3 Pricing Overview

Please visit <https://aws.amazon.com/pricing/> for further details on pricing AWS services.

25.4 Service Constraints

Please see <https://aws.amazon.com/efs/> for more information.

25.5 Technical Requirements

Amazon EFS provides file storage for your Amazon EC2 instances. With Amazon EFS, you can create a file system, mount the file system on your Amazon EC2 instances, and then read and write data from your Amazon EC2 instances to and from your file system.

- **User Guide** – Walks through how to set up Amazon EFS and integrate it with other services. Includes the API reference. [HTML](#) | [PDF](#)

26.0 Service Definition – Amazon Elastic Graphics

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

26.1 Service Overview

Amazon Elastic Graphics provides flexible, low-cost, and high-performance graphics acceleration for your Windows instances. Elastic Graphics accelerators come in multiple sizes and are a low-cost alternative to using GPU graphics instance types (such as G2 and G3). You have the flexibility to choose an instance type that meets the compute, memory, and storage needs of your application. Then, choose the accelerator for your instance that meets the graphics requirements of your workload.

Top features include:

- **Accelerated graphics performance** – Amazon Elastic Graphics allows you to easily attach OpenGL 4.3 acceleration capabilities to your Amazon EC2 instances.
- **Different sizes** – Elastic Graphics accelerators come with different sizes, from the smallest eg1.medium with 1 GiB of graphics memory to eg1.2xlarge with 8 GiB of graphics memory.
- **Elastic Graphics driver** – Elastic Graphics accelerators come with a driver, which is an OpenGL acceleration library in your EC2 instances.

Top benefits include:

- **Workstation-class graphics performance** – Elastic Graphics provides accelerators capable of running a variety of graphics workloads, such as 3D modelling and rendering, with similar workstation performance compared to direct-attached GPUs.
- **Optimised performance and cost** – Elastic Graphics accelerators come in multiple frame buffer sizes up to 8 GiB, allowing you to achieve the optimal graphics performance for your workload for the lowest possible cost.
- **Application support** – Elastic Graphics supports OpenGL, a cross-language, cross-platform API for rendering 2D and 3D vector graphics, with a roadmap for certification of additional APIs. Elastic Graphics can be used in conjunction with many desktop streaming protocols, such as NICE DCV, Remote Desktop Protocol (RDP), and Virtual Network Computing (VNC).
- **Pay for use** – You are charged for each second that an Elastic Graphics accelerator is attached to an instance in the running state when the accelerator is in the Ok state. You are not charged for an accelerator attached to an instance that is in the pending, stopping, stopped, shutting down, or terminated state. You are also not charged when an accelerator is in the Unknown or Impaired state.

26.2 Backup/Restore and Disaster Recovery

This requirement is not applicable to Elastic Graphics. For additional information beyond what is described herein, please refer to <http://aws.amazon.com/documentation/>.

26.3 Pricing Overview

Please visit <https://aws.amazon.com/pricing/> for further details on pricing AWS services.

26.4 Service Constraints

Please see <https://docs.aws.amazon.com/AWSEC2/latest/WindowsGuide/elastic-graphics.html#elastic-graphics-limitations> for more information.

26.5 Technical Requirements

Please see <https://docs.aws.amazon.com/AWSEC2/latest/WindowsGuide/elastic-graphics.html> for comprehensive technical documentation regarding Elastic Graphics.

27.0 Service Definition – Amazon Elastic Load Balancing

Elastic Load Balancing automatically distributes incoming application traffic across multiple targets, such as Amazon EC2 instances, containers, IP addresses, and Lambda functions. It can handle the varying load of your application traffic in a single Availability Zone or across multiple Availability Zones. Elastic Load Balancing offers three types of load balancers that all feature the high availability, automatic scaling, and robust security necessary to make your applications fault tolerant.

There are three types of Load Balancer: Application Load Balancer, Network Load Balancer, and Classic Load Balancer.

27.1 Service Overview

Top features include:

- **High availability.** Elastic Load Balancing automatically distributes traffic across multiple targets – Amazon EC2 instances, containers and IP addresses – in a single Availability Zone or multiple Availability Zones.
- **Health checks.** Elastic Load Balancing can detect unhealthy targets, stop sending traffic to them, and then spread the load across the remaining healthy targets.
- **Security features.** Use Amazon Virtual Private Cloud (Amazon VPC) to create and manage security groups associated with load balancers to provide additional networking and security options. You can also create an internal (non-internet-facing) load balancer.
- **TLS termination.** Elastic Load Balancing provides integrated certificate management and SSL/TLS decryption, allowing you the flexibility to centrally manage the SSL settings of the load balancer and offload CPU intensive work from your application.
- **Layer 4 or Layer 7 load balancing.** You can load balance HTTP/HTTPS applications for layer 7-specific features, or use strict layer 4 load balancing for applications that rely on the TCP and UDP protocols.
- **Operational monitoring.** Elastic Load Balancing provides integration with Amazon CloudWatch metrics and request tracing in order to monitor performance of your applications in real time.

Top benefits include:

- **Secure.** Elastic Load Balancing works with Amazon Virtual Private Cloud (VPC) to provide robust security features, including integrated certificate management, user-authentication, and SSL/TLS decryption. Together, they give you the flexibility to centrally manage TLS settings and offload CPU intensive workloads from your applications. Get exactly what you need.
- **Robust monitoring & auditing.** Elastic Load Balancing allows you to monitor your applications and their performance in real time with Amazon CloudWatch metrics, logging, and request tracing. This improves visibility into the behavior of your applications, uncovering issues and identifying performance bottlenecks in your application stack at the granularity of an individual request.
- **Hybrid load balancing.** Elastic Load Balancing offers ability to load balance across AWS and on-premises resources using the same load balancer. This makes it easy for you to migrate, burst, or failover on-premises applications to the cloud.

27.2 Backup/Restore and Disaster Recovery

This requirement is not applicable for Amazon Elastic Load Balancer. For additional information beyond what is described herein, please refer to <http://aws.amazon.com/documentation/>.

The AWS global infrastructure is built around AWS Regions and Availability Zones. AWS Regions provide multiple physically separated and isolated 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 centre infrastructures.

For more information about AWS Regions and Availability Zones, see AWS Global Infrastructure (<http://aws.amazon.com/about-aws/global-infrastructure/>).

27.3 Pricing Overview

Please see the AWS UK G Cloud 12 Pricing Document affiliated with this service in the Digital Marketplace.

27.4 Service Constraints

Please see <https://docs.aws.amazon.com/elasticloadbalancing/> for more information regarding the constraints in place when using Amazon Elastic Load Balancing.

27.5 Technical Requirements

Please refer to <https://docs.aws.amazon.com/elasticloadbalancing/> for comprehensive technical documentation regarding Amazon Elastic Load Balancing

28.0 Service Definition – Amazon Elastic Inference

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

28.1 Service Overview

Top features include:

- **Integrated with Amazon SageMaker, Amazon EC2, and Amazon ECS**

There are multiple ways to run inference workloads on AWS: deploy your model on Amazon SageMaker for a fully managed experience, or run it on Amazon EC2 instances or Amazon ECS tasks and manage it yourself. Amazon Elastic Inference is integrated to work seamlessly with Amazon SageMaker, Amazon EC2, and Amazon ECS, allowing you to add inference acceleration in all scenarios. You can specify the desired amount of inference acceleration when you create your model's HTTPS endpoint in Amazon SageMaker, when you launch your Amazon EC2 instance, and when you define your Amazon ECS task.

- **TensorFlow, Apache MXNet and PyTorch support**

Amazon Elastic Inference is designed to be used with AWS's enhanced versions of TensorFlow Serving, Apache MXNet and PyTorch. These enhancements enable the frameworks to automatically detect the presence of inference accelerators, optimally distribute the model operations between the accelerator's GPU and the instance's CPU, and securely control access to your accelerators using AWS Identity and Access Management (IAM) policies. The enhanced TensorFlow Serving, MXNet and PyTorch libraries are provided automatically in Amazon SageMaker, AWS Deep Learning AMIs, and AWS Deep Learning Containers, so you don't have to make any code change to deploy your models in production. You can also download them separately by following the instructions [here](#).

- **Open Neural Network Exchange (ONNX) format support**

ONNX is an open format that makes it possible to train a model in one deep learning framework and then transfer it to another for inference. This allows you to take advantage of the relative strengths of different frameworks. ONNX is integrated into PyTorch, MXNet, Chainer, Caffe2, and Microsoft Cognitive Toolkit, and there are connectors for many other frameworks including TensorFlow. To use ONNX models with Amazon Elastic Inference, your trained models need to be transferred to the AWS-optimized version of Apache MXNet for production deployment.

- **Choice of single or mixed precision operations**

Amazon Elastic Inference accelerators support both single-precision (32-bit floating point) operations and mixed precision (16-bit floating point) operations. Single precision provides an extremely large numerical range to represent the parameters used by your model. However, most models don't actually need this much precision and calculating numbers that large results in unnecessary loss of performance. To avoid that problem, mixed-precision operations allow you to reduce the numerical range by half to gain up to 8x greater inference performance.

- **Available in multiple amounts of acceleration**

Amazon Elastic Inference is available in multiple throughput sizes ranging from 1 to 32 trillion floating point operations per second (TFLOPS) per accelerator, making it efficient for accelerating a wide range of inference models including computer vision, natural language processing, and speech recognition. Compared to standalone Amazon EC2 P3 instances that start at 125 TFLOPS (the smallest P3 instance available), Amazon Elastic Inference starts at a single TFLOPS per accelerator. This allows you to scale up inference acceleration in more appropriate increments. You can also select from larger accelerator sizes, up to 32 TFLOPS per accelerator, for more complex models.

- **Auto-scaling**

Amazon Elastic Inference can be part of the same Amazon EC2 Auto Scaling group you use to scale your Amazon SageMaker, Amazon EC2, and Amazon ECS instances. When EC2 Auto Scaling adds more EC2 instances to meet the demands of your application, it also scales up the accelerator attached to each instance. Similarly, when Auto Scaling reduces your EC2 instances as demand goes down, it also scales down the attached accelerator for each instance. This makes it easy to scale your inference acceleration alongside your application's compute capacity to meet the demands of your application.

Top benefits include:

- **Reduce inference costs by up to 75%.** Amazon Elastic Inference allows you to choose the instance type that is best suited to the overall compute and memory needs of your application. You can then separately specify the amount of inference acceleration that you need. This reduces inference costs by up to 75% because you no longer need to over-provision GPU compute for inference.
- **Get exactly what you need.** Amazon Elastic Inference can provide as little as a single-precision TFLOPS (trillion floating point operations per second) of inference acceleration or as much as 32 mixed-precision TFLOPS. This is a much more appropriate range of inference compute than the range of up to 1,000 TFLOPS provided by a standalone Amazon EC2 P3 instance. For example, a simple language processing model might require only one TFLOPS to run inference well, while a sophisticated computer vision model might need up to 32 TFLOPS.

- **Respond to changes in demand.** You can easily scale the amount of inference acceleration up and down using Amazon EC2 Auto Scaling groups to meet the demands of your application without over-provisioning capacity. When EC2 Auto Scaling increases your EC2 instances to meet increasing demand, it also automatically scales up the attached accelerator for each instance. Similarly, when it reduces your EC2 instances as demand goes down, it also automatically scales down the attached accelerator for each instance. This helps you pay only for what you need when you need it.

28.2 Backup/Restore and Disaster Recovery

This requirement is not applicable for Amazon Elastic Inference. For additional information beyond what is described herein, please refer to <http://aws.amazon.com/documentation/>.

The AWS global infrastructure is built around AWS Regions and Availability Zones. AWS Regions provide multiple physically separated and isolated 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 centre infrastructures.

For more information about AWS Regions and Availability Zones, see AWS Global Infrastructure (<http://aws.amazon.com/about-aws/global-infrastructure/>).

28.3 Pricing Overview

Please see the AWS UK G Cloud 12 Pricing Document affiliated with this service in the Digital Marketplace.

28.4 Service Constraints

Please see <https://docs.aws.amazon.com/elastic-inference/latest/developerguide/before.html> for more information regarding the constraints in place when using AWS Data Exchange

28.5 Technical Requirements

Please refer to <https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/elastic-inference.html> for comprehensive technical documentation regarding Amazon Elastic Inference.

29.0 Service Definition – Amazon Elastic Transcoder

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

29.1 Service Overview

Amazon Elastic Transcoder is media transcoding in the cloud. It is designed to be a highly scalable, easy-to-use, and cost-effective way for developers and businesses to convert (or “transcode”) media files from their source format into versions that will playback on devices like smartphones, tablets, and PCs.

Top benefits include:

- **Easy to Use** – Amazon Elastic Transcoder is designed to be easy to use. You can get started by using the AWS Management Console, the service API, or the SDKs. System transcoding pre-sets make it easy to get transcoding settings right the first time. We provide predefined pre-sets to create media files that will play on a wide range of devices (like smartphones or tablets), as well as pre-sets to create media files that are optimised for playback on a specific device (like the Amazon Kindle Fire HD or Apple iPod touch). You can also create segmented files and playlists for delivery using the HLS, Smooth, or MPEG-DASH protocol to compatible devices. Developers building applications that need transcoding can use the [AWS SDKs](#) for Java, .NET, Node.js, PHP, Python and Ruby, and the [AWS CLI](#).
- **Elastically Scalable** – Amazon Elastic Transcoder is designed to scale seamlessly with your media transcoding workload. Amazon Elastic Transcoder is architected to handle large volumes of media files and large file sizes. Transcoding pipelines enable you to perform multiple transcodes in parallel.

Amazon Elastic Transcoder leverages other AWS Cloud services like [Amazon S3](#), [Amazon EC2](#), [Amazon DynamoDB](#), [Amazon SWF](#), and [Amazon SNS](#) to give scalability and reliability.
- **Cost Effective** – Amazon Elastic Transcoder uses a content duration-based pricing model: you pay based on the length of the output, in minutes, of the media you are transcoding. For example, if your video’s transcoded output is 30 minutes in duration, you will pay for 30 minutes of transcoding. Similarly, if you create a 20 minute video clip from a 30 minute input file, you will pay for 20 minutes of transcoding. With Amazon Elastic Transcoder, there are no minimum transcoding volumes, monthly commitments, or long-term contracts.
- **Managed** – Amazon Elastic Transcoder enables you to focus on your content rather than on managing transcoding software in a distributed cloud-based environment. The service takes care of scaling and operating the system and manages the process of keeping codecs up to date. Combined with our service API and SDKs, this makes it easy for you to build media solutions that use Amazon Elastic Transcoder.
- **Secure** – Your content is under your control: your assets are in your own Amazon S3 buckets, to which you give us access through AWS IAM roles. This makes it simple to fit seamlessly into your current security and identity framework without giving up control. In building Amazon Elastic Transcoder, we made use of security best practices gained by building other AWS Cloud services. For more

information about security on AWS, please refer to the [AWS Security Center](#). For more information on compliance, including MPAA best practices, please refer to [AWS Compliance](#).

- **Seamless Delivery** – Using Amazon Elastic Transcoder, Amazon S3, and Amazon CloudFront, you can store, transcode, and deliver your content. By setting the Amazon S3 permissions for your CloudFront distribution in Amazon Elastic Transcoder, it is now a simple, one-step process to transcode content with Amazon Elastic Transcoder and deliver the multiple output videos via progressive download or adaptive bitrate streaming (HLS, Smooth, or MPEGDASH) with CloudFront.
- **AWS Integration** – Amazon Elastic Transcoder provides an important media building block for creating end-to-end media solutions on AWS. For example, you can use Amazon Glacier to store master content, Amazon Elastic Transcoder to transcode masters to renditions for distribution stored in Amazon S3, and stream these renditions at scale over the Internet using Amazon CloudFront.

29.2 Backup/Restore and Disaster Recovery

This requirement is not applicable for Amazon Elastic Transcoder. For additional information beyond what is described herein, please refer to <http://aws.amazon.com/documentation/>.

29.3 Pricing Overview

Please visit <https://aws.amazon.com/pricing/> for further details on pricing AWS services.

29.4 Service Constraints

Please see <https://aws.amazon.com/elastictranscoder/> for more information.

29.5 Technical Requirements

Amazon Elastic Transcoder lets you convert media files that you have stored in Amazon S3 into media files in the formats required by consumer playback devices. For example, you can convert large, high-quality digital media files into formats that users can play back on mobile devices, tablets, web browsers, and connected televisions.

- **Developer Guide** – Provides a conceptual overview of Amazon Elastic Transcoder and includes detailed instructions for using its various features.
[HTML](#) | [PDF](#) | [Kindle](#)

30.0 Service Definition – Amazon ElastiCache

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

30.1 Service Overview

Amazon ElastiCache is a web service that makes it easy to deploy, operate, and scale an in-memory cache in the cloud. The service improves the performance of web applications by allowing you to retrieve information from fast, managed, in-memory caches, instead of relying entirely on slower, disk-based databases. ElastiCache supports two open-source, in-memory caching engines:

- **Memcached** – A widely adopted memory object caching system. ElastiCache is protocol compliant with Memcached, so popular tools that you use today with existing Memcached environments will work seamlessly with the service.
- **Redis** – A popular open-source, in-memory key-value store that supports data structures such as sorted sets and lists. ElastiCache supports Master/Slave replication and Multi-AZ, which can be used to achieve cross-AZ redundancy.

Amazon ElastiCache automatically detects and replaces failed nodes, reducing the overhead associated with self-managed infrastructures and provides a resilient system that mitigates the risk of overloaded databases, which slow website and application load times. Through integration with [Amazon CloudWatch](#), Amazon ElastiCache provides enhanced visibility into key performance metrics associated with your Memcached or Redis nodes.

Using Amazon ElastiCache, you can add an in-memory caching layer to your infrastructure in a matter of minutes by using the [AWS Management Console](#).

Top benefits include:

- **Simple to Deploy** – Amazon ElastiCache makes it very easy to deploy a Memcached or Redis-compliant cache environment. Use the [AWS Management Console](#) or simple API calls to access the capabilities of a production-ready cloud cache cluster in minutes without worrying about infrastructure provisioning, or installing and maintaining cache software.
- **Managed** – Amazon ElastiCache automates time-consuming management tasks—such as patch management, failure detection, and recovery—allowing you to pursue higher-value application development.
- **Compatible** – With Amazon ElastiCache, you get native access to the Memcached or Redis in-memory caching environments. This facilitates compatibility with your existing tools and applications.
- **Elastic** – With a simple API call or a few clicks of the AWS Management Console, you can add or delete Cache Nodes to your cloud cache cluster to meet your application load. [Auto Discovery](#) for Memcached enables automatic discovery of Cache Nodes by ElastiCache Clients when the nodes are added to or removed from an Amazon ElastiCache Cluster.
- **Reliable** – Amazon ElastiCache has multiple features that enhance reliability for critical production deployments, including automatic failure detection and recovery. Amazon ElastiCache runs on the same highly reliable infrastructure

used by other AWS Cloud services.

- **Integrated** – Amazon ElastiCache is designed for seamless use with other AWS Cloud services, including Amazon RDS, Amazon DynamoDB, Amazon EC2, Amazon CloudWatch, and Amazon SNS.
- **Secure** – Amazon ElastiCache provides a number of mechanisms to secure your Cache Cluster.
 - Amazon ElastiCache provides interfaces that allow you to configure firewall settings that control network access to your Cache Cluster.
 - Amazon ElastiCache allows you to run your Cache Cluster in [Amazon VPC](#). Amazon VPC enables you to isolate your Cache Cluster by specifying the IP ranges you wish to use for your Cache Nodes, and connect to your existing applications inside Amazon VPC. To learn more about Amazon ElastiCache in Amazon VPC, refer to the [Amazon ElastiCache User Guide](#).
- **Cost Effective** – Amazon ElastiCache saves you the administrative cost of setting up and managing a multi-node Cache Cluster. You can scale up and scale down the number of Cache Nodes in your Cache Cluster to deliver optimum performance as your application usage pattern changes, paying only for the resources you actually consume. The on-demand pricing allows you to pay for memory/compute capacity by the hour with no long-term commitments. This makes the use of Amazon ElastiCache very cost effective and frees you from the costs and complexities of planning, purchasing, and maintaining hardware.
- **Multi-AZ** – Amazon ElastiCache provides replication features for the Redis engine and Multi-AZ functionality. You can take advantage of multiple AWS Availability Zones to gain availability and scale beyond the capacity constraints of a single cache node. In case of primary node loss, Amazon ElastiCache will automatically detect the failure and failover to a read replica to provide higher availability without the need for manual intervention.

30.2 Backup/Restore and Disaster Recovery

Amazon ElastiCache for Redis helps you protect your data by creating snapshots of your clusters. Via a few clicks on the console or simple API calls, you can set up automatic snapshots, as well as initiate a backup whenever you choose. The snapshots can then be used for seeding new ElastiCache for Redis clusters. For additional information beyond what is described herein, please refer to <http://aws.amazon.com/documentation/>.

30.3 Pricing Overview

Please visit <https://aws.amazon.com/pricing/> for further details on pricing AWS services.

30.4 Service Constraints

Please see <https://aws.amazon.com/elasticache/> for more information.

30.5 Technical Requirements

Amazon ElastiCache is a web service that makes it easy to set up, manage, and scale distributed in-memory cache environments in the cloud. It provides a high-performance, resizable, and cost-effective in-memory cache, while removing the complexity associated with deploying and managing a distributed cache environment.

- **User Guide** – Helps you understand the components and features that Amazon ElastiCache provides and how to use them. You'll learn how to **access** ElastiCache through a web-based GUI, Command Line Tools, and the Amazon ElastiCache API. [HTML](#) | [PDF](#) | [Kindle](#)
- **ElastiCache section of the AWS CLI Reference** – Describes the CLI for Amazon **ElastiCache** in detail. Provides basic syntax, options, and usage examples for each command. [HTML](#)
- **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. [HTML](#) | [PDF](#)

31.0 Service Definition – Amazon Elasticsearch Service

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

31.1 Service Overview

Amazon Elasticsearch Service is a managed service that makes it easy to deploy, operate, and scale Elasticsearch in the AWS Cloud. Elasticsearch is a popular open source search and analytics engine for use cases such as log analytics, real-time application monitoring, and click stream analytics. You can set up and configure your Amazon Elasticsearch cluster in minutes from the AWS Management Console. Amazon Elasticsearch Service provisions all the resources for your cluster and launches it. The service automatically detects and replaces failed Elasticsearch nodes, reducing the overhead associated with self-managed infrastructure and Elasticsearch software. Amazon Elasticsearch Service allows you to easily scale your cluster via a single API call or a few clicks in the AWS Management Console. With Amazon Elasticsearch Service, you get direct access to the Elasticsearch open-source API so that code and applications you're already using with your existing Elasticsearch environments will work seamlessly. Top benefits include:

- **Simple to Deploy** – Amazon Elasticsearch Service makes it easy to deploy Elasticsearch in the AWS Cloud. Use the AWS Management Console or simple API calls to access a production-ready Amazon Elasticsearch cluster in minutes without worrying about infrastructure provisioning or installing and maintaining Elasticsearch software.
- **Easy to Administer** – Amazon Elasticsearch Service simplifies time-consuming

management tasks—such as ensuring high availability, patch management, failure detection and node replacement, backups, and monitoring—allowing you to pursue higher-value application development.

- **Scalable** – Amazon Elasticsearch Service enables you to monitor your cluster through Amazon CloudWatch metrics and resize your cluster up or down via a single API call or a few clicks on the AWS Management Console.
- **Integrated with Logstash or Kibana** – Logstash is an open-source data pipeline that helps you process logs and other event data and load them into Elasticsearch. Kibana is an open-source analytics and visualization platform that helps you get a better understanding of your data in Elasticsearch. Amazon Elasticsearch Service includes built-in integration with Kibana and supports integration with Logstash.
- **Cost Effective** – Amazon Elasticsearch Service saves you the administrative costs of setting up and managing Elasticsearch. You can scale up and scale down your cluster to deliver optimum performance as data and usage patterns change, paying only for the resources you actually consume. The on-demand pricing allows you to pay for resources by the hour with no long-term commitments and frees you from the costs and complexities of planning, purchasing, and maintaining hardware.
- **Secure** – You can control access to the Elasticsearch APIs using AWS IAM policies. Using IAM policies you can allow your applications to access your Amazon Elasticsearch clusters securely.

31.2 Backup/Restore and Disaster Recovery

This requirement is not applicable for Amazon Elasticsearch Service. For additional information beyond what is described herein, please refer to <http://aws.amazon.com/documentation/>.

31.3 Pricing Overview

Please visit <https://aws.amazon.com/pricing/> for further details on pricing AWS services.

31.4 Service Constraints

Please see <https://aws.amazon.com/elasticsearch-service/> for more information.

31.5 Technical Requirements

Amazon Elasticsearch Service is a managed service that makes it easy to deploy, operate, and scale Elasticsearch, a popular open-source search and analytics engine. Amazon Elasticsearch Service also offers security options, high availability, data durability, and direct access to the Elasticsearch API.

- **Developer Guide** – Describes how to get started with Amazon Elasticsearch Service, explains key concepts, and provides step-by-step instructions that show

you how to use the features. [HTML](#) | [PDF](#)

32.0 Service Definition – Amazon EMR

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

32.1 Service Overview

Amazon EMR provides a managed Hadoop framework that makes it easy, fast, and cost-effective to process vast amounts of data across dynamically scalable Amazon EC2 instances. You can also run other popular distributed frameworks such as [Apache Spark](#), [HBase](#), [Presto](#), and [Flink](#) in Amazon EMR, and interact with data in other AWS data stores such as Amazon S3 and Amazon DynamoDB.

Amazon EMR securely and reliably handles a broad set of big data use cases, including log analysis, web indexing, data transformations (Extract, Transform, Load [ETL]), ML, financial analysis, scientific simulation, and bioinformatics.

Top features include:

- **Easy to Use** – You can launch an Amazon EMR cluster in minutes. You don't need to worry about node provisioning, cluster setup, Hadoop configuration, or cluster tuning. Amazon EMR takes care of these tasks so you can focus on analysis.
- **Low Cost** – Amazon EMR pricing is simple and predictable: You pay an hourly rate for every instance hour you use. Because Amazon EMR has native support for Amazon EC2 Spot and Reserved Instances, you can also save 50-80% on the cost of the underlying instances.
- **Elastic** – With Amazon EMR, you can provision one, hundreds, or thousands of compute instances to process data at any scale. You can easily increase or decrease the number of instances and you only pay for what you use.
- **Reliable** – You can spend less time tuning and monitoring your cluster. Amazon EMR has tuned Hadoop for the cloud; it also monitors your cluster, retrying failed tasks and automatically replacing poorly performing instances.
- **Secure** – Amazon EMR automatically configures Amazon EC2 firewall settings that control network access to instances, and you can launch clusters in an Amazon VPC, a logically isolated network you define. For objects stored in Amazon S3, you can use Amazon S3 [server-side encryption](#) or Amazon S3 [client-side encryption](#) with Amazon EMR File System (EMRFS), using AWS KMS or customer-managed keys.
- **Flexible** – You have complete control over your cluster. You have root access to every instance, you can easily install additional applications, and you can customise every cluster. Amazon EMR also supports multiple Hadoop distributions and applications.

32.2 Backup/Restore and Disaster Recovery

With Amazon EMR, you can leverage multiple data stores, including Amazon S3, Amazon RDS, Amazon Glacier, and Amazon Redshift. Refer to those services for details on their backup/restore and disaster recovery.

32.3 Pricing Overview

Please visit <https://aws.amazon.com/pricing/> for further details on pricing AWS services.

32.4 Service Constraints

Please see <https://aws.amazon.com/documentation/elastic-mapreduce/> for more information.

32.5 Technical Requirements

Please refer to the links below and <https://aws.amazon.com/documentation/elasticmapreduce/> and the following links for comprehensive technical documentation regarding Amazon EMR.

- **Management Guide** – Provides information about the Amazon EMR service and platform. [HTML](#) | [PDF](#)
- **Release Guide** – Provides information about Amazon EMR releases, including software installed on your Amazon EMR cluster like Hadoop and Spark. [HTML](#) | [PDF](#)
- **API Reference** – Describes all the API operations for Amazon EMR in detail. Also provides sample requests, responses, and errors for the supported web services protocols. [HTML](#) | [PDF](#)

33.0 Service Definition - Amazon EventBridge

Amazon EventBridge is a serverless event bus that makes it easy to connect applications together using data from your own applications, integrated Software-as-a-Service (SaaS) applications, and AWS services. EventBridge delivers a stream of real-time data from event sources, such as Zendesk, Datadog, or Pagerduty, and routes that data to targets like AWS Lambda.

33.1 Service Overview

Top Features include:

- Serverless, fully managed, and scalable event bus that allows applications to communicate using events. There is no infrastructure to manage and no capacity to provision.
- Natively integrated with Software-as-a-service (SaaS) applications from many providers including DataDog, OneLogin, PagerDuty, Savyint, Segment,

SignalFX, SugarCRM, Symantec, Whispir, and Zendesk, with additional integrations planned.

- Directly integrated with over 90 event sources and over 15 targets, including AWS Lambda, Amazon SQS, Amazon SNS, AWS Step Functions, Amazon Kinesis Data Streams, Amazon Kinesis Data Firehose, and more, with additional sources and targets planned.
- Makes it easy for you to build event-driven application architectures. Applications or microservices can publish events to the event bus without awareness of subscribers.
- 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.

Top benefits include:

- EventBridge makes it easy to connect applications together because you can ingest, filter and deliver events without writing custom code. EventBridge automatically routes events from their source to one or more supported AWS service targets, such as AWS Lambda and Amazon Kinesis Data Streams.
- EventBridge ingests data from supported SaaS applications and routes it to AWS service targets through native integration in the AWS management console. With EventBridge, you can use data from your SaaS apps to trigger workflows for customer support, business operations and more.
- You can use Amazon EventBridge as a centralized event bus for your enterprise applications and microservices. You can send events from your own applications to an event bus via the service's PutEvents API. Other applications can then receive events through any of the many supported AWS target services.
- EventBridge simplifies the process of building event-driven architectures. With EventBridge, your event targets don't need to be aware of event sources because you can filter and publish directly to EventBridge. There is no set up required. Event-driven architectures are loosely coupled and distributed, which improves developer agility as well as application resiliency.
- With EventBridge, there are no servers to provision, patch, and manage and there is no software to install, maintain, or operate. EventBridge automatically scales based on the number of events ingested, and you pay only for events published by your AWS or SaaS applications. EventBridge has built-in distributed availability and fault-tolerance.
- You can filter events with rules. A rule matches incoming events for a given event bus and routes them to targets for processing. A single rule can route to multiple targets, all of which are processed in parallel. Rules allow different application components to look for and process the events that are of interest to them. A rule can customize an event before it is sent to the target by passing along only

certain parts or by overwriting it with a constant. You can also have multiple rules that match on the same event, so different microservices or applications can choose to match events based on specific filters.

- Amazon EventBridge extends its predecessor, Amazon CloudWatch Events, and provides a near real-time stream of system events that describe changes to your AWS resources. It allows you to respond quickly to operational changes and take corrective action. You simply write rules to indicate which events are of interest to your application and what automated actions to take when a rule matches an event.
- You can set up scheduled events using the popular Unix cron syntax. Scheduled events are generated on a periodic basis and invoke any of the supported target AWS services.
- You can monitor your event bus using Amazon CloudWatch metrics, such as the number of times an event matches a rule, or the number of times a target is invoked. You can use Amazon CloudWatch Logs to store, monitor, and analyse events that are triggered in your environment. AWS CloudTrail enables you to monitor the calls made to the Amazon EventBridge API.
- Amazon EventBridge is integrated with AWS Identity and Access Management (IAM) so that you can control which users and resources have permission to access your data and how they can access it. EventBridge supports VPC endpoints and encryption in transit using TLS 1.2. Amazon EventBridge is GDPR, SOC, ISO, DoD CC SRG, and FedRamp compliant and is also HIPAA-eligible.

33.2 Backup/Restore and Disaster Recovery

Amazon EventBridge is a managed service. Customers do not need to define their own backup, restore, or disaster recovery strategy.

33.3 Pricing Overview

Please see the AWS UK G Cloud 12 Pricing Document affiliated with this service in the Digital Marketplace.

33.4 Service Constraints

Please see

https://docs.aws.amazon.com/general/latest/gr/aws_service_limits.html#limits_eventbridge for more information.

33.5 Technical Requirements

Please refer to <https://docs.aws.amazon.com/eventbridge/> and the following links for comprehensive technical documentation regarding Amazon EventBridge.

- <https://docs.aws.amazon.com/eventbridge/latest/userguide/what-is-amazon-eventbridge.html>
- <https://docs.aws.amazon.com/eventbridge/latest/userguide/eventbridge-getting-set-up.html>
- <https://docs.aws.amazon.com/eventbridge/latest/userguide/security-eventbridge.html>

34.0 Service Definition – Amazon Forecast

Amazon Forecast is a fully managed service that uses machine learning to deliver highly accurate forecasts.

Amazon Forecast uses machine learning to combine time series data with additional variables to build forecasts. Amazon Forecast requires no machine learning experience to get started. You only need to provide historical data, plus any additional data that you believe may impact your forecasts. For example, the demand for a particular color of a shirt may change with the seasons and store location. This complex relationship is hard to determine on its own, but machine learning is ideally suited to recognize it. Once you provide your data, Amazon Forecast will automatically examine it, identify what is meaningful, and produce a forecasting model capable of making predictions that are up to 50% more accurate than looking at time series data alone.

34.1 Service Overview

Top features include:

- **Works with any historical time series data to create accurate forecasts.** Amazon Forecast can use virtually any historical time series data (e.g., price, promotions, economic performance metrics) to create accurate forecasts for your business. For example, in a retail scenario, Amazon Forecast uses machine learning to process your time series data (such as price, promotions, and store traffic) and combines that with associated data (such as product features, floor placement, and store locations) to determine the complex relationships between them. By combining time series data with additional variables, Amazon Forecast can be 50% more accurate than non-machine learning forecasting tools.
- **Automated machine learning.** No machine learning expertise is required to build an accurate time series-forecasting model that can incorporate time series data from multiple variables at once. Amazon Forecast includes AutoML capabilities that take care of the machine learning for you. Once you provide your data into Amazon S3, Amazon Forecast can automatically load and inspect the data, select the right algorithms, train a model, provide accuracy metrics, and generate forecasts.
- **Based on the same technology used at Amazon.com.** Amazon Forecast includes algorithms that are based on over twenty years of forecasting experience and developed expertise used by Amazon.com. Using AutoML,

Amazon Forecast will automatically select the best algorithm based on your data sets.

- **Easily evaluate the accuracy of your forecasting models.** Amazon Forecast provides comprehensive accuracy metrics to help you understand the performance of your forecasting model and compare it to previous forecasting models you've created that may have looked at a different set of variables or used a different period of time for the historical data. Amazon Forecast allows you to create multiple backtest windows and visualize the metrics, helping you evaluate model accuracy over different start dates.
- **Visualize forecasts.** Amazon Forecasts and their associated accuracy metrics are visualized in easy-to-understand graphs and tables in the service console. Once forecasts are generated, you can navigate to the relevant forecast by picking it from a list of available forecasts. For example, a specific product within your full catalog of products. Visualization allows you to quickly understand the details of each forecast and determine if adjustments are necessary.
- **Integrate with your existing tools.** Amazon Forecast can be easily imported into common business and supply chain applications, such as SAP and Oracle Supply Chain. This makes it easy to integrate more accurate forecasting into your existing business processes with little to no change.
- **Generate probabilistic forecasts.** Unlike most other forecasting solutions that generate point forecasts, Amazon Forecast generates probabilistic forecasts at three different quantiles by default: 10%, 50% and 90%. In addition, you can choose any quantile between 1% and 99%, including the 'mean' forecast. This allows you to choose a forecast that suits your business needs depending on whether the cost of capital (over forecasting) or missing customer demand (under forecasting) is of importance.

Top benefits include:

- **50% more accurate forecasts with machine learning.** Amazon Forecast provides forecasts that are up to 50% more accurate by using machine learning to automatically discover how time series data and other variables like product features and store locations affect each other. You are better able to understand how these complex relationships ultimately affect demand than what looking at time series data alone can deliver. The models that Amazon Forecast builds are unique to your data, which means the predictions are custom fit to your business.
- **Reduce forecasting time from months to hours.** With Amazon Forecast, you can achieve forecasting accuracy levels that used to take months of engineering in as little as a few hours. You can import time series data and associated data into Amazon Forecast from your Amazon S3 database. From there, Amazon Forecast automatically loads your data, inspects it, and identifies the key attributes needed for forecasting. Amazon Forecast then trains and optimizes your custom model, and hosts them in a highly available environment where it can be used to generate your business forecasts. By automatically handling the

complex machine learning required to build, train, tune, and deploy a forecasting model, Amazon Forecast enables you to create accurate forecasts quickly.

- **Create virtually any time series forecast.** Multiple types of time series forecasts are required to run your business, from cash flow to product demand to resource planning. Amazon Forecast allows you to build forecasts for virtually every industry and use case, including retail, logistics, finance, advertising performance, and many more. Using machine learning, Amazon Forecast can work with any historical time series data and use a large library of built-in algorithms to determine the best fit for your particular forecast type automatically.
- **Secure your business data and peace of mind.** Every interaction you have with Amazon Forecast is protected by encryption. Any content processed by Amazon Forecast is encrypted with customer keys through Amazon Key Management Service, and encrypted at rest in the AWS Region where you are using the service. Administrators can also control access to Amazon Forecast through an AWS Identity and Access Management (IAM) permissions policy – ensuring that sensitive information is kept secure and confidential.

34.2 Backup/Restore and Disaster Recovery

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

The AWS global infrastructure is built around AWS Regions and Availability Zones. AWS Regions provide multiple physically separated and isolated 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 centre infrastructures.

For more information about AWS Regions and Availability Zones, see AWS Global Infrastructure (<http://aws.amazon.com/about-aws/global-infrastructure/>).

34.3 Pricing Overview

Please see the AWS UK G Cloud 12 Pricing Document affiliated with this service in the Digital Marketplace.

34.4 Service Constraints

Please see <https://docs.aws.amazon.com/forecast/> for more information regarding the constraints in place when using Amazon Forecast.

34.5 Technical Requirements

Please refer to <https://docs.aws.amazon.com/forecast/> for comprehensive technical documentation regarding Amazon Forecast.

35.0 Service Definition – Amazon FreeRTOS

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

35.1 Service Overview

Amazon FreeRTOS (a:FreeRTOS) is an operating system for microcontrollers that makes small, low-power edge devices easy to program, deploy, secure, connect, and manage. Amazon FreeRTOS is based on the FreeRTOS kernel, a popular open source operating system for microcontrollers, and extends it with software libraries that make it easy to securely connect your small, low-power devices to AWS cloud services like AWS IoT Core or to more powerful edge devices running AWS Greengrass.

A microcontroller (MCU) is a single chip containing a simple processor that can be found in many devices, including appliances, sensors, fitness trackers, industrial automation, and automobiles. Many of these small devices could benefit from connecting to the cloud or locally to other devices. For example, smart electricity meters need to connect to the cloud to report on usage and building security systems need to communicate locally so that a door will unlock when you badge in. Microcontrollers have limited compute power and memory capacity and typically perform simple, functional tasks. Microcontrollers frequently run operating systems which do not have built in functionality to connect to local networks or the cloud, making IoT applications a challenge. Amazon FreeRTOS helps solve this problem by providing both the core operating system (to run the edge device) as well as software libraries that make it easy to securely connect to the cloud (or other edge devices) so you can collect data from them for IoT applications and take action.

Top benefits include:

- **Based on the FreeRTOS Kernel** - Amazon FreeRTOS is open source and is based on the FreeRTOS kernel, the most popular real-time operating system for microcontrollers. With Amazon FreeRTOS, you can take advantage of the large ecosystem of existing tools developed for the FreeRTOS kernel. If you are already using the FreeRTOS kernel on your existing devices, you can choose to use Amazon FreeRTOS and take advantage of the functionality provided by the additional libraries, or continue to use just the kernel operating system.
- **Easily Program, Deploy, and Manage Low Power Connected Devices** - Amazon FreeRTOS includes software libraries that make it easy to program commonly needed IoT capabilities into your device, including libraries that help configure devices to a local network using common connectivity options like WiFi or Ethernet. Amazon FreeRTOS also includes an over-the-air (OTA) update feature, now in beta, to remotely update devices with feature enhancements or security patches.
- **Secure Data and Device Connections** - Amazon FreeRTOS comes with libraries to help secure device data and connections, including support for data encryption and key management. Amazon FreeRTOS includes support for

Transport Layer Security (TLS v1.2) to help your devices connect securely to the cloud. Amazon FreeRTOS also has a code signing feature to ensure your device code is not compromised during deployment and OTA updates (now in beta).

- **Connect to the Cloud or to a Local Edge Device** - Amazon FreeRTOS devices can connect directly to cloud services like AWS IoT Core, or to a local edge device, like AWS Greengrass devices. By connecting Amazon FreeRTOS devices to a Greengrass Core device, you can continue to communicate with the device even if it loses connection to the cloud.
- **Broad Hardware and Technology Ecosystem** - Amazon FreeRTOS gives you the flexibility to easily build IoT solutions on a variety of chipsets and supports a variety of architectures, such as ARM and MIPS. For new devices, you can choose from a variety of qualified chipsets from trusted providers, including Espressif, Microchip, NXP Semiconductors, STMicroelectronics, Texas Instruments, and others from the Amazon FreeRTOS Qualification Program.

35.2 Backup/Restore and Disaster Recovery

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

35.3 Pricing Overview

Please visit <https://aws.amazon.com/pricing/> for further details on pricing AWS services.

35.4 Service Constraints

Please see <https://aws.amazon.com/iot-core/> for more information.

35.5 Technical Requirements

Amazon FreeRTOS is an operating system that makes microcontroller-based edge devices easy to program, deploy, secure, and maintain. Amazon FreeRTOS extends the FreeRTOS kernel, a popular open source operating system for microcontrollers, with software libraries that can be used to securely connect devices (locally and to the cloud) and update them remotely.

- **User Guide** - Provides detailed information about the microcontroller operating system that makes small, low-powered edge devices easy to program, deploy, secure, and maintain. [HTML](#) | [PDF](#)

Additional documentation is available at <https://aws.amazon.com/documentation/freertos/>.

36.0 Service Definition – Amazon FSx for Lustre

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

36.1 Service Overview

Amazon FSx for Lustre provides a high-performance file system optimised for fast processing of workloads such as machine learning, high performance computing (HPC), video processing, financial modelling, and electronic design automation (EDA). These workloads commonly require data to be presented via a fast and scalable file system interface and typically have datasets stored on long-term data stores like Amazon S3.

Top features include:

Low latencies and high throughput – FSx for Lustre can scale up to hundreds of gigabytes per second of throughput and millions of IOPS.

- **Seamless access to your data** – FSx for Lustre works natively with Amazon S3, making it easy to access your Amazon S3 data to run data processing workloads. Your Amazon S3 objects are presented as files in your file system, and you can write your results back to Amazon S3.
- **On-premises integration** – You can use FSx for Lustre for on-premises workloads that need to burst to the cloud due to peak demands or capacity limits.
- **Storage for high-performance workloads** – FSx for Lustre is optimised for a variety of data processing workloads that run from hours to months. FSx for Lustre is optimised for the performance and costs of high-performance workloads, with your long-term storage on Amazon S3 or on-premises data stores.
- **Fully managed** – With a few clicks in the Amazon FSx for Lustre console, CLI, or API, you can create and launch a Lustre file system that is accessible to your compute instances and linked to your Amazon S3 bucket.

Top benefits include:

- **High performance and scalability** – FSx for Lustre is built on Lustre, a popular high-performance file system that is optimised for data processing, with sub-millisecond latencies and throughput that scales to hundreds of gigabytes per second.
- **Seamless access to Amazon S3 or on-premises data** – FSx for Lustre works natively with Amazon S3, making it easy to process your Amazon S3 data with a high-performance POSIX interface.
- **Fully managed** – FSx for Lustre is fully managed, making it easy to launch and run high-performance file systems in the cloud. In minutes, you can create and launch an FSx for Lustre file system by using the [AWS Management Console](#), the AWS CLI, or an AWS SDK.

- **Cost effective** – FSx for Lustre helps you optimise your storage cost for high-performance workloads. It provides cheap and performant storage for processing data, with your long-term data stored in Amazon S3 or other low-cost, long-term data stores.
- **Secure and compliant** – FSx for Lustre automatically encrypts your data at rest. If you are subject to regulatory compliance, FSx for Lustre is PCI DSS and ISO compliant and HIPAA eligible.
- **Automated operations** – FSx for Lustre fully manages and maintains the infrastructure and software powering your file system, performs routine Lustre updates, and detects and addresses hardware issues.
- **Resource-level permissions** – FSx for Lustre is integrated with IAM. This integration means that you can control the actions your IAM users and groups can take to manage your file systems (such as creating and deleting file systems).
- **Access** – 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.
- **Performance monitoring** – FSx for Lustre provides native CloudWatch integration. CloudWatch metrics include storage consumed, number of compute instance connections, throughput, and number of file operations per second.
- **Scaling performance** – FSx for Lustre provides 200 MBps of baseline throughput per TB of storage provisioned. As your file systems scales, your throughput can scale up to hundreds of GB/s.

36.2 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. This durability level corresponds to an average annual expected loss of 0.000000001% of objects. For example, if you store 10,000,000 objects with Amazon S3, you can on average expect to incur a loss of a single object once every 10,000 years. In addition, Amazon S3 Standard, S3 Standard-IA, and S3 Glacier are all designed to sustain data in the event of an entire Amazon S3 Availability Zone loss.

Amazon S3 Standard, S3 Standard-IA, and S3 Glacier storage classes redundantly store your objects on multiple devices across a minimum of three Availability Zones in an Amazon S3 Region before returning SUCCESS. The S3 One Zone-IA storage class stores data redundantly across multiple devices within a single Availability Zone. These

services are designed to sustain concurrent device failures by quickly detecting and repairing any lost redundancy, and they also regularly verify the integrity of your data using checksums.

36.3 Pricing Overview

Please visit <https://aws.amazon.com/pricing/> for further details on pricing AWS services.

36.4 Service Constraints

Please see <https://docs.aws.amazon.com/fsx/latest/LustreGuide/limits.html> for more information.

36.5 Technical Requirements

Please refer to <https://docs.aws.amazon.com/fsx/latest/LustreGuide/what-is.html> for comprehensive technical documentation regarding Amazon FSx for Lustre.

37.0 Service Definition – Amazon FSx for Windows File Server

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

37.1 Service Overview

Amazon FSx for Windows File Server provides a fully managed native Microsoft Windows file system. Built on Windows Server, FSx for Windows File Server provides shared file storage with the compatibility and features that your Windows-based application rely on.

Top features include:

- Full support for the SMB protocol and Windows NTFS
- Active Directory integration and Distributed File System (DFS)
- Fully managed
- Integration with other key AWS Cloud services, including CloudTrail, AWS Direct Connect, AWS VPN, AWS Directory Service for Microsoft AD, IAM, and VPC security groups
- Configurable throughput capacity

Top benefits include:

- **Fully managed** – You no longer have to worry about setting up and provisioning file servers and storage volumes, installing file server software updates and patches, or manually performing backups.
- **Secure and compliant** – FSx for Windows File Server provides multiple levels of security and compliance to help ensure that data is protected. It automatically

encrypts your data at rest and in transit. FSx for Windows File Server is designed to meet the highest security standards, complies with ISO and PCI DSS certifications, and is HIPAA eligible.

- **Native Windows compatibility** – Built on Windows Server, you get native Windows file storage with the features and performance needed to move your applications to AWS at no cost.
- **Fast and flexible performance** – Built on SSD storage, FSx for Windows File Server provides fast performance with up to 2 GB/second throughput per file system, hundreds of thousands of IOPS, and consistent sub-millisecond latencies.
- **Broadly accessible** – By supporting SMB Protocol, FSx for Windows File Server can connect your file system to Amazon EC2, VMware Cloud on AWS, Amazon WorkSpaces, and Amazon AppStream 2.0 instances. It supports all Windows versions, starting from Windows Server 2008 and Windows 7, and current versions of Linux. It can also support on-premises access using AWS Direct Connect or AWS VPN.
- **Cost effective** – You pay for the resources you use, with no minimum commitments or upfront fees. You can launch and delete file systems in minutes, making it easy to respond to changing business needs.

37.2 Backup/Restore and Disaster Recovery

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

37.3 Pricing Overview

Please visit <https://aws.amazon.com/pricing/> for further details on pricing AWS services.

37.4 Service Constraints

Please see <https://docs.aws.amazon.com/fsx/latest/WindowsGuide/limits.html> for more information.

37.5 Technical Requirements

Please refer to <https://docs.aws.amazon.com/fsx/latest/WindowsGuide/what-is.html> for comprehensive technical documentation regarding Amazon FSx for Windows File Server.

38.0 Service Definition – Amazon GameLift

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

38.1 Service Overview

Amazon GameLift is a managed service for deploying, operating, and scaling dedicated game servers for session-based multiplayer games. You can deploy your first game server in the cloud in just minutes, saving up to thousands of engineering hours in upfront software development and lowering the technical risks that often cause developers to cut multiplayer features from their designs. Built on AWS's proven computing environment, Amazon GameLift lets you scale high-performance game servers up and down to meet player demand. You pay only for the capacity you use, so you can get started quickly whether you're working on a new game idea or running a game with millions of players.

Top features include:

- **Server management** – Amazon GameLift provisions instances, deploys your game servers onto running instances, load balances traffic across fleets of game servers, monitors instance and game server health, and replaces unhealthy instances without your intervention.
- **Choice of instance types** – Amazon GameLift offers both On-Demand and Spot Instances. With Spot Instances you get the same high-performance hardware as On-Demand Instances at savings of up to 90%.
- **Worldwide game placement** – Amazon GameLift continually scans for available game servers around the world. If low-latency game servers are not available, you can configure the service to automatically add more capacity nearby your players.
- **Autoscaling** – Amazon GameLift's autoscaling feature can start one, hundreds, or even thousands of instances simultaneously and stop unused instances in just minutes.
- **DDoS protection** – Amazon GameLift is designed to safeguard your game servers from frequently occurring network and transport layer DDoS attacks.

Top benefits include:

- **Preview** – Create your game server in minutes. Building a multiplayer game is now possible for studios of any size and experience. With GameLift Realtime Servers, quickly create and customise your game server with just a few lines of JavaScript.
- **Automatically scale server capacity with player traffic** – Be prepared for unexpected spikes in player traffic, but stop paying for idle server capacity when there's less player traffic.
- **Flexible matchmaking** – Use your own matchmaking service or our

customisable solution, FlexMatch, to create competitive matches based on rules you define.

- **No downtime for updates** – Keep servers live and players playing 24/7 by using Amazon GameLift, even when releasing updates or new content.
- **Support cross-platform play** – Connect your community across devices, and broaden your pool of players for potential matches.
- **Works with your existing engine and workflows** – Whether you use a AAA engine like Unity, Unreal, Lumberyard, or a home-grown C++ solution, the Amazon GameLift SDK easily integrates to get your servers up and running in the cloud.
- **Reliable infrastructure** – Amazon GameLift runs within Amazon’s proven network infrastructure and data centres.
- **Insights and monitoring** – The Amazon GameLift management console includes dashboards for monitoring the performance of your game. You can see real-time data on player demand, server capacity and health, CPU and memory utilisation, and create operational alarms using [Amazon CloudWatch](#).
- **Pay as you go** – Amazon GameLift includes everything you need to run your dedicated game servers in the cloud: a high-performance virtual machine, SSD-based storage, and fast data transfer to the internet—on demand without monthly commitments. You only pay for the resources you actually use.
- **Multi-platform support** – Amazon GameLift supports games built for major game platforms, including Android, FireOS, iOS, Mac, PC, PlayStation 4, and Xbox One.

38.2 Backup/Restore and Disaster Recovery

Amazon GameLift stores server builds within the build catalogue. These builds are then used to provision GameLift instances into a fleet. Individual instances are each provisioned with a 50 GB Amazon EBS volume. These volumes can utilise snapshots to provide point-in-time recovery of the data. Snapshots can be taken in real time while the volume is attached and in use.

38.3 Pricing Overview

Please visit <https://aws.amazon.com/pricing/> for further details on pricing AWS services.

38.4 Service Constraints

Please see https://aws.amazon.com/gamelift/faq/#Operational_Limits for more information.

38.5 Technical Requirements

Please refer to <https://aws.amazon.com/gamelift/faq/> for comprehensive technical documentation regarding Amazon GameLift.

39.0 Service Definition – Amazon Glacier

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

39.1 Service Overview

Amazon Glacier is a secure, durable, and extremely low-cost storage service for data archiving and online backup. Customers can reliably store large or small amounts of data. To keep costs low yet suitable for varying retrieval needs, Amazon Glacier provides three options for access to archives, from a few minutes to several hours.

Top benefits include:

- **Low Cost** – Amazon Glacier allows you to archive large amounts of data at a very low cost.
- **Secure** – Amazon Glacier supports data transfer over SSL and automatically encrypts your data at rest.
- **Durable** – Amazon Glacier provides a highly durable storage infrastructure designed for online backup and archival. Your data is redundantly stored across multiple facilities and multiple devices in each facility.
- **Simple** – Amazon Glacier allows you to offload the administrative burden of operating storage infrastructure to AWS.
- **Flexible** – Amazon Glacier scales to meet your storage needs. There is no limit to how much data you can store, and you can choose to store your data in the AWS Region that supports your regulatory and business criteria.
- **Integrated** – Through Amazon S3 life cycle policies, you can optimise your storage costs by moving infrequently accessed objects from Amazon S3 to Amazon Glacier (or vice versa).

39.2 Backup/Restore and Disaster Recovery

Amazon S3 enables you to use Amazon Glacier's extremely low-cost storage service as a storage option for data archival. Amazon Glacier is optimised for data that is infrequently accessed and for which retrieval times of several hours are suitable. Examples include digital media archives, financial and healthcare records, raw genomic sequence data, long-term database backups, and data that must be retained for regulatory compliance.

For additional information beyond what is described herein, please refer to <http://aws.amazon.com/documentation>.

39.3 Pricing Overview

Please visit <https://aws.amazon.com/pricing/> for further details on pricing AWS services.

39.4 Service Constraints

Please see <http://aws.amazon.com/documentation/glacier/> for more information.

39.5 Technical Requirements

Please refer to <http://aws.amazon.com/documentation/glacier/> and the following links for comprehensive technical documentation regarding Amazon Glacier.

- **Developer Guide** – Provides detailed information about setting up and working with Amazon Glacier using the REST API and the AWS SDKs for Java and Microsoft .NET. Available in [HTML](#), [PDF](#), and [Kindle](#) formats.

40.0 Service Definition – Amazon GuardDuty

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

40.1 Service Overview

Amazon GuardDuty is a managed threat detection service that continuously monitors for malicious or unauthorized behaviour to help you protect your AWS accounts and workloads. It monitors for activity such as unusual API calls or potentially unauthorized deployments that indicate a possible account compromise. GuardDuty also detects potentially compromised instances or reconnaissance by attackers.

Enabled with a few clicks in the AWS Management Console, Amazon GuardDuty can immediately begin analysing billions of events across your AWS accounts for signs of risk. GuardDuty identifies suspected attackers through integrated threat intelligence feeds and uses machine learning to detect anomalies in account and workload activity. When a potential threat is detected, the service delivers a detailed security alert to the GuardDuty console and AWS CloudWatch Events. This makes alerts actionable and easy to integrate into existing event management and workflow systems.

- Amazon GuardDuty is cost effective and easy. It does not require you to deploy and maintain software or security infrastructure, meaning it can be enabled quickly with no risk of negatively impacting existing application workloads. There are no upfront costs with GuardDuty, no software to deploy, and no threat intelligence feeds required. Customers pay for the events analysed by GuardDuty and there is a 30-day free trial available for every new account to the service..

Top benefits include:

- **Intelligent threat detection** - Amazon GuardDuty gives you intelligent threat detection by collecting, analysing, and correlating billions of events from AWS CloudTrail, Amazon VPC Flow Logs, and DNS Logs across all of your associated AWS accounts. GuardDuty detections are made more accurate by incorporating threat intelligence (such as lists of known malicious IP addresses provided by AWS Security and 3rd party threat intelligence partners). GuardDuty also uses

machine learning to detect anomalous account and network activities. For example, GuardDuty will alert you if it detects remote API calls from a known malicious IP address indicating potentially compromised AWS credentials. GuardDuty also detects direct threats to your AWS environment indicating a compromised instance, such as an Amazon EC2 instance sending encoded data within DNS queries.

- **Centralize threat detection across accounts** - Many organisations use multiple AWS accounts to help provide proper cost allocation, agility, and security. With a few clicks in the AWS Management Console, you can centralize your threat detection by enabling Amazon GuardDuty across any of your AWS accounts. With GuardDuty, there is no need to install additional security software or infrastructure to analyse your account and workload activity data. Your security operations centre team can easily manage and triage threats from a single console view and automate security responses using a single security account.
- **Strengthens security through automation** - in addition to detecting threats, Amazon GuardDuty also makes it easy to automate how you respond to these threats, reducing your remediation and recovery time. You can set up your remediation scripts or AWS Lambda functions to trigger based on GuardDuty findings. GuardDuty security findings include the affected resource's details, such as tags, security groups, or credentials. GuardDuty findings also include attacker information, such as IP address and geo-location. This makes GuardDuty security findings informative and actionable. For example, account compromise can be difficult to detect quickly if you are not continuously monitoring account activities in near real-time. With GuardDuty, when an instance is suspected of having data stolen the service will alert you to be able to automatically create an access control entry restricting outbound access for that instance..

40.2 Backup/Restore and Disaster Recovery

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

40.3 Pricing Overview

Please visit <https://aws.amazon.com/pricing/> for further details on pricing AWS services.

40.4 Service Constraints

Please see <https://aws.amazon.com/guardduty/resources/> for more information.

40.5 Technical Requirements

- **User Guide** – Walks through how to set up Amazon GuardDuty and evaluate the security of your AWS environment. [HTML](#) | [PDF](#)
- **API Reference** – Describes all of the API operations for Amazon GuardDuty.

Also provides sample requests, responses, and errors for the supported web services protocols. [HTML](#) | [PDF](#)

41.0 Service Definition – Amazon Inspector

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

41.1 Service Overview

Amazon Inspector is an automated security assessment service that helps improve the security and compliance of applications deployed on AWS. Amazon Inspector automatically assesses applications for vulnerabilities or deviations from best practices. After performing an assessment, Amazon Inspector produces a detailed list of security findings prioritised by level of severity.

To help you get started quickly, Amazon Inspector includes a knowledge base of hundreds of rules mapped to common security best practices and vulnerability definitions. Examples of built-in rules include checking for remote root login being enabled or vulnerable software versions installed. These rules are regularly updated by AWS security researchers.

Top benefits include:

- **Identify Application Security Issues** – Amazon Inspector helps you to identify security vulnerabilities as well as deviations from security best practices in applications, both before they are deployed and while they are running in a production environment. This helps improve the overall security posture of your applications deployed on AWS.
- **Integrate Security into DevOps** – Amazon Inspector is agent-based, API-driven, and delivered as a service. This makes it easy for you to build right into your existing DevOps process, decentralizing and automating vulnerability assessments and empowering your development and operations teams to make security assessment an integral part of the deployment process.
- **Increase Development Agility** – Amazon Inspector helps you reduce the risk of introducing security issues during development and deployment by automating the security assessment of your applications and proactively identifying vulnerabilities. This allows you to develop and iterate on new applications quickly and assess compliance with best practices and policies.
- **Leverage AWS Security Expertise** – The AWS security organisation is continuously assessing the AWS environment and updating a knowledge base of security best practices and rules. Amazon Inspector makes this expertise available to you in the form of a service that simplifies the process of establishing and enforcing best practices within your AWS environment.
- **Streamline Security Compliance** – Amazon Inspector gives security teams and auditors visibility into the security testing that is being performed during

development of applications on AWS. This streamlines the process of validating and demonstrating that security and compliance standards and best practices are being followed throughout the development process.

- **Enforce Security Standards** – Amazon Inspector allows you to define standards and best practices for your applications and validate adherence to these standards. This simplifies enforcement of your organisation's security standards and best practices and helps to proactively manage security issues before they impact your production application.

41.2 Backup/Restore and Disaster Recovery

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

41.3 Pricing Overview

Please visit <https://aws.amazon.com/pricing/> for further details on pricing AWS services.

41.4 Service Constraints

Please see <http://aws.amazon.com/inspector/> for more information.

41.5 Technical Requirements

Amazon Inspector is a security vulnerability assessment service that helps improve the security and compliance of your AWS resources. Amazon Inspector automatically assesses resources for vulnerabilities or deviations from best practices and then produces a detailed list of security findings prioritised by level of severity. Amazon Inspector includes a knowledge base of hundreds of rules mapped to common security standards and vulnerability definitions that are regularly updated by AWS security researchers.

- **User Guide** – Walks through how to set up Amazon Inspector and evaluate your security configuration. [HTML](#) | [PDF](#)
- **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. [HTML](#) | [PDF](#)
- **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. [HTML](#)

42.0 Service Definition – Amazon Kinesis

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

This service listing includes the following AWS services:

- Amazon Kinesis
- Amazon Kinesis Data Analytics
- Amazon Kinesis Data Firehose
- Amazon Kinesis Data Streams
- Amazon Kinesis Video Streams

42.1 Service Overview

Amazon Kinesis is a platform for [streaming data](#) on AWS, offering powerful services to make it easy to load and analyse streaming data and also providing the ability for you to build custom streaming data applications for specialised needs. Web applications, mobile devices, wearables, industrial sensors, and many software applications and services can generate staggering amounts of streaming data—sometimes TBs per hour—that need to be collected, stored, and processed continuously. Amazon Kinesis services enable you to do that simply and at a low cost. Amazon Kinesis consists of Amazon Kinesis Firehose and Amazon Kinesis Streams.

Amazon Kinesis Firehose is the easiest way to load [streaming data](#) into AWS. It can capture and automatically load streaming data into [Amazon S3](#), [Amazon Redshift](#), and [Amazon Elasticsearch Service](#), 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. You can easily create a Firehose delivery stream from the AWS Management Console, configure it with a few clicks, and start sending data to the stream from hundreds of thousands of data sources to be loaded continuously to AWS—all in just a few minutes. With Amazon Kinesis Firehose, you only pay for the amount of data you transmit through the service. There is no minimum fee or setup cost.

Amazon Kinesis Streams enables you to build custom applications that process or analyse [streaming data](#) for specialised needs. Amazon Kinesis Streams can continuously capture and store terabytes of data per hour from hundreds of thousands of sources such as website clickstreams, financial transactions, social media feeds, IT logs, and location-tracking events. With Amazon Kinesis Client Library (KCL), you can build Amazon Kinesis Applications and use streaming data to power real-time dashboards, generate alerts, implement dynamic pricing and advertising, and more. You can also emit data from Amazon Kinesis Streams to other AWS Cloud services such as Amazon S3, Amazon Redshift, Amazon EMR, and AWS Lambda.

42.2 Backup/Restore and Disaster Recovery

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

42.3 Pricing Overview

Please visit <https://aws.amazon.com/pricing/> for further details on pricing AWS services.

42.4 Service Constraints

Please see <https://aws.amazon.com/kinesis/> for more information.

42.5 Technical Requirements

Amazon Kinesis is a platform for streaming data on AWS, making it easy to load and analyse streaming data, and also providing the ability for you to build custom streaming data applications for specialised needs. Use Amazon Kinesis Streams to collect and process large streams of data records in real time. Use Amazon Kinesis Firehose to deliver real-time streaming data to destinations such as Amazon S3 and Amazon Redshift.

Amazon Kinesis Streams

- **User Guide** – Provides a conceptual overview of Amazon Kinesis Streams and includes detailed development instructions for using the various features. [HTML](#) | [PDF](#)
- **API Reference** – Describes all the API operations for Amazon Kinesis Streams in detail. Also provides sample requests, responses, and errors for the supported web services protocols. [HTML](#) | [PDF](#)

Amazon Kinesis Firehose

- **Developer Guide** – Provides a conceptual overview of Amazon Kinesis Firehose and includes detailed instructions for using the service. [HTML](#) | [PDF](#) | [Kindle](#)
- **API Reference** – Describes all the API operations for Amazon Kinesis Firehose in detail. Also provides sample requests, responses, and errors for the supported web services protocols. [HTML](#) | [PDF](#)

42.6 Service Overview - Amazon Kinesis Data Analytics

Amazon Kinesis Data Analytics enables the processing and analysis of streaming data. The service enables you to quickly author and run powerful SQL or Java code against streaming sources to perform time series analytics, feed real-time dashboards, and create real-time metrics.

The SQL service uses standard SQL extended with new language elements to analyse the incoming stream. To do this you use an interactive editor to author SQL with new temporal predicates such as “LEADS” or “STRICTLY LAGS.” The Java service leverages the [Apache Flink](#) framework to perform the analysis with code. Both support

ingesting data from Amazon Kinesis Data Streams and Amazon Kinesis Data Firehose streaming sources. You can test with live streaming data and configure destinations where you want to send the results.

Top features include:

- Fully managed AWS service
- Serverless processing where you pay only for what you use. You are charged an hourly rate based on the average number of KPIs that are used to run your stream processing application. A single KPI provides you with 1 vCPU and 4 GB of memory and 50 GB of running application storage or 1GB of SQL reference data.
- The ability to specify up to three output destinations

Top benefits include:

- Rapid development of applications using SQL
- Ability to leverage existing open source tooling via Apache Flink

42.6.1 Backup/Restore and Disaster Recovery

We recommend that you configure your Kinesis streams to back up to Amazon S3. This can be done directly via configuration or by pipelining through additional Kinesis Firehose streams. When using a Java application, you can create up to 1,000 “snapshots,” which can then be used to update and/or restart processing.

42.6.2 Pricing Overview

Please visit <https://aws.amazon.com/pricing/> for further details on pricing AWS services.

42.6.3 Service Constraints

There are different constraints for [the SQL service](#) and [the Java service](#). You should click on the previous links for detailed information; however, we have provided the main constraints per application below:

- The size of a row in a stream is limited to 511 KB (512 KB less 1 KB for metadata).
- The SQL code is limited to 100 KB.
- Maximum streaming throughput is 100 MB/sec.

42.6.4 Technical Requirements

Please refer to the [Kinesis Data Analytics for SQL Applications](#) or the [Kinesis Data Analytics for Java Applications](#) guides for comprehensive technical documentation.

42.7 Service Overview – Amazon Kinesis Data Firehose

Amazon Kinesis Data Firehose is a fully managed service for delivering real-time streaming data to destinations such as Amazon S3, Amazon Redshift, Amazon

Elasticsearch Service (Amazon ES), and Splunk. Kinesis Data Firehose is part of the Kinesis streaming data platform, along with Kinesis Data Streams, Kinesis Video Streams, and Kinesis Data Analytics. With Kinesis Data Firehose, you don't need to write applications or manage resources. You configure your data producers to send data to Kinesis Data Firehose, and it automatically delivers the data to the destination that you specified. You can also configure Kinesis Data Firehose to transform your data before delivering it.

Top features include:

- Simple consumption of data from common sources, including CloudWatch (both Logs and Events) and AWS IoT.
- An agent is available for ingestion from disk files.
- No code is necessary to store data to common destinations, including Amazon Redshift, Elasticsearch clusters, and Splunk.

Simple data transformations are possible by using Lambda functions.

Top benefits include:

- Rapid creation of a data processing pipeline with little to no coding
- Highly scalable, fully managed service
- Data compression formats are optional and may be GZIP, Snappy, or Zip.

42.7.1 Backup/Restore and Disaster Recovery

Kinesis Data Firehose can be configured to back up data to an Amazon S3 bucket for disaster recovery. It is a regional service and automatically uses multiple Availability Zones for resilience.

42.7.2 Pricing Overview

Please visit <https://aws.amazon.com/pricing/> for further details on pricing AWS services.

42.7.3 Service Constraints

Records are limited to 1000 kB.

Snappy or Zip compression is not available for delivery streams with Amazon Redshift as the destination.

42.7.4 Technical Requirements

Please refer to the [Developer documentation](#) for comprehensive technical documentation regarding Kinesis Data Firehose.

42.8 Service Overview – Amazon Kinesis Data Streams

- With Amazon Kinesis Data Streams, you can create custom data processing applications that collect, process, and analyse large streams of data records in real time using popular stream-processing frameworks. A typical Kinesis Data

Streams application reads data from a data stream as data records. These applications can use the Kinesis Client Library, and they can run on Amazon EC2 instances. The records can be processed and sent to dashboards and/or used to generate alerts. They can also send data to a variety of other AWS Cloud services.

Top features include:

- Data once input is highly durable.
- The capture process is fully fault tolerant.
- The processing capacity of the stream can be scaled up and down elastically.
- Incoming data may be processed by subsequent pipelines with sub-second latency.
- Multiple consumers may process the same stream.

Top benefits include:

- Ability to dynamically create directed acyclic graphs of processing pipelines to handle incoming data in a variety of ways.
- Ease of integration with Amazon EMR for data reduction
- Fully managed service allowing simple scaling
- Ability to specify data retention period to automatically expire old data
- Integration with AWS KMS to allow encryption of data
- Uses private VPC endpoints for increased data security

42.8.1 Backup/Restore and Disaster Recovery

Kinesis Data Streams is a regional construct and uses multiple Availability Zones for durability. Data older than the retention period of the stream must be archived by the application appropriately.

42.8.2 Pricing Overview

Please visit <https://aws.amazon.com/pricing/> for further details on pricing AWS services.

42.8.3 Service Constraints

Please see the [Service Limits](#) for more information.

42.8.4 Technical Requirements

Please refer to [the Developer guide](#) for comprehensive technical documentation regarding Kinesis Data Streams.

42.9 Service Overview - Amazon Kinesis Video Streams

Amazon Kinesis Video Streams makes it easy to securely stream video from connected devices to AWS for analytics, machine learning (ML), playback, and other processing. Kinesis Video Streams automatically provisions and elastically scales all the

infrastructure needed to ingest streaming video data from millions of devices. It durably stores, encrypts, and indexes video data in your streams, and allows you to access your data through easy-to-use APIs.

Top Features include

- Enables you to securely ingest, process, and store video and time-encoded data from devices at any scale.
- Amazon Rekognition Video allows you to specify any of your Amazon Kinesis video streams as an input. This enables you to automatically detect and recognize faces in streaming video.
- Supports WebRTC for low-latency, peer-to-peer, two-way media streaming.
- Automatically indexes the data you store in your video streams based on timestamps generated by the device, or timestamps generated by Kinesis Video Streams when it receives the video.
- Provides SDKs in C++ and Java that you can build and configure for your connected devices.

Top benefits include

- **Stream video from millions of devices** - Amazon Kinesis Video Streams provides SDKs that make it easy for devices to securely stream media to AWS for playback, storage, analytics, machine learning, and other processing. Kinesis Video Streams can ingest data from edge devices, smartphones, security cameras, and other data sources such as RADARs, LIDARs, drones, satellites, dash cams, and depth-sensors.
- **Build real-time vision and video-enabled apps** - Easily build applications with real-time computer vision capabilities through integration with Amazon Rekognition Video, and with real-time video analytics capabilities using popular open source machine learning frameworks.
- **Playback live and recorded video streams** - Easily stream live and recorded media from your Kinesis video streams to your browser or mobile application using the Kinesis Video Streams HTTP Live Streaming (HLS) capability.
- **Build apps with two-way, real-time media streaming** - Amazon Kinesis Video Streams supports the open-source project WebRTC for two-way, real-time media streaming between web browsers, mobile applications, and connected devices. With support for WebRTC, you can use simple APIs to build rich applications like video chat and peer-to-peer data sharing with ultra-low latency and two-way communication between your applications and connected devices.
- **Secure** - Amazon Kinesis Video Streams allows you to control access to your streams using AWS Identity and Access Management (IAM). It helps you protect your data by automatically encrypting the data at rest using AWS Key Management Service (KMS) and in transit using the industry-standard Transport Layer Security (TLS) protocol.
- **Durable, searchable storage** - Amazon Kinesis Video Streams uses Amazon S3 as the underlying data store, which means your data is stored durably and

reliably. Kinesis Video Streams enables you to quickly search and retrieve video fragments based on device and service generated timestamps.

- **No infrastructure to manage** - Amazon Kinesis Video Streams manages all the infrastructure for you. You don't have to worry about configuration, software updates, failures, or scaling infrastructure as the number of streams and consuming applications grows. Kinesis Video Streams handles all the administration and maintenance required to manage your streams, so you can focus your time on building innovative applications.

42.9.1 Backup/Restore and Disaster Recovery

This requirement is not applicable for Amazon Kinesis Video Streams. For additional information beyond what is described herein, please refer to <http://aws.amazon.com/documentation/>.

The AWS global infrastructure is built around AWS Regions and Availability Zones. AWS Regions provide multiple physically separated and isolated 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 Availability Zones without interruption. Availability Zones are more highly available, fault tolerant, and scalable than traditional single or multiple data centre infrastructures.

For more information about AWS Regions and Availability Zones, see AWS Global Infrastructure (<https://aws.amazon.com/about-aws/global-infrastructure/>).

42.9.2 Service Constraints

Please see <https://docs.aws.amazon.com/kinesisvideostreams/latest/dg/limits.html> for more information.

42.9.3 Technical Requirements

Please refer to <https://docs.aws.amazon.com/kinesis/index.html> and the following links for comprehensive technical documentation regarding Amazon Kinesis Video Streams.

- <https://docs.aws.amazon.com/kinesisvideostreams/latest/dg/what-is-kinesis-video.html>
- https://docs.aws.amazon.com/kinesisvideostreams/latest/dg/API_Reference.html
- <https://docs.aws.amazon.com/kinesisvideostreams-webRTC-dg/latest/devguide/what-is-kvswebRTC.html>
- <https://docs.aws.amazon.com/kinesisvideostreams/latest/dg/security.html>

43.0 Service Definition – Amazon Lex

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

43.1 Service Overview

Amazon Lex is a service for building conversational interfaces into any application using voice and text. Amazon Lex provides the advanced deep learning functionalities of automatic speech recognition (ASR) for converting speech to text, and natural language understanding (NLU) to recognize the intent of the text, to enable you to build applications with highly engaging user experiences and lifelike conversational interactions. With Amazon Lex, the same deep learning technologies that power Amazon Alexa are now available to any developer, enabling you to quickly and easily build sophisticated, natural language, conversational bots ("[chatbots](#)").

Speech recognition and natural language understanding are some of the most challenging problems to solve in computer science, requiring sophisticated deep learning algorithms to be trained on massive amounts of data and infrastructure. Amazon Lex democratizes these deep learning technologies by putting the power of Amazon Alexa within reach of all developers. Harnessing these technologies, Amazon Lex enables you to define entirely new categories of products made possible through conversational interfaces.

As a fully managed service, Amazon Lex scales automatically, so you don't need to worry about managing infrastructure. With Amazon Lex, you pay only for what you use.

There are no upfront commitments or minimum fees.

Top benefits include:

- **Easy to Use** - Amazon Lex provides an easy-to-use console to guide you through the process of creating your own chatbot in minutes, building conversational interfaces into your applications. You supply just a few example phrases and Lex builds a complete natural language model through which your user can interact using voice and text, to ask questions, get answers, and complete sophisticated tasks.
- **Seamlessly Deploy and Scale** - With Amazon Lex, you can build, test, and deploy your chatbots directly from the Amazon Lex console. Lex enables you to easily publish your voice or text chatbots to mobile devices, web apps, and chat services such as Facebook Messenger (with Slack and Twilio integration coming soon). Once published, your Amazon Lex bot processes voice or text input in conversation with your end-users. Amazon Lex is a fully managed service so as your user engagement increases, you don't need to worry about provisioning hardware and managing infrastructure to power your bot experience.
- **Built-in Integration with the AWS Platform** - Amazon Lex provides built-in integration with AWS Lambda, AWS MobileHub and Amazon CloudWatch and you can easily integrate with many other services on the AWS platform including Amazon Cognito, and Amazon DynamoDB. You can take advantage of the power of the AWS platform for security, monitoring, user authentication, business logic, storage and mobile app development.
- **Cost Effective** - With Amazon Lex, there are no upfront costs or minimum fees.

You are only charged for the text or speech requests that are made. Amazon Lex' pay-as-you-go pricing and low cost per request make it a cost-effective way to build conversational interfaces anywhere. With the Amazon Lex free tier, you can easily try Lex without any initial investment.

43.2 Backup/Restore and Disaster Recovery

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

43.3 Pricing Overview

Please visit <https://aws.amazon.com/pricing/> for further details on pricing AWS services.

43.4 Service Constraints

Please see <https://aws.amazon.com/lex/> for more information. Prior to using these Services, Buyer agrees to contact Supplier before using these Services in connection with Buyer Personal Data.

43.5 Technical Requirements

Amazon Lex is an AWS service for building conversational interfaces into applications using voice and text. With Amazon Lex, the same deep learning engine that powers Amazon Alexa is now available to any developer, enabling you to build sophisticated, natural language chatbots into your new and existing applications. Amazon Lex provides the deep functionality and flexibility of natural language understanding (NLU) and automatic speech recognition (ASR) to enable you to build highly engaging user experiences with lifelike, conversational interactions and create new categories of products.

- **Developer Guide** - Provides a conceptual overview of Amazon Lex, includes detailed instructions for using the various features, and provides a complete API reference for developers. [HTML](#) | [PDF](#)

44.0 Service Definition – Amazon Lightsail

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

44.1 Service Overview

Amazon Lightsail is designed to be the easiest way to launch and manage a virtual private server with AWS. Amazon Lightsail plans include everything you need to jumpstart your project—a virtual machine, SSD-based storage, data transfer, DNS management, and a static IP—for a low, predictable price. Amazon Lightsail virtual private servers run on the same highly available and reliable AWS Cloud infrastructure used by millions of customers.

Top benefits include:

- **Amazon Lightsail Virtual Private Server (VPS)** – Experience the power and reliability of AWS. Deploy in seconds and manage from the intuitive Lightsail management console or API.
- **Keeping servers simple** – Launch a virtual private server with just a few clicks.
- **Built on the galaxy's leading cloud** – Give your application the power, reliability, and security of AWS.
- **You know code, we know infrastructure** – Let Amazon Lightsail do the heavy lifting. We make server management a breeze.
- **The cloud that grows with you** – Scale applications with access to dozens of AWS Cloud services.
- **Powerful API** – Use the simple and flexible Amazon Lightsail API to extend your application or integrate it with external applications.
- **Highly available storage** – Every Amazon Lightsail server comes with high performing, persistent SSD-based block storage.
- **Speedy & secure networking** – Amazon Lightsail servers run at warp speed on the AWS network. Configure your network simply and securely, including your IP addresses, DNS, firewall, and more.
- **Snapshots** – Protect your data, clone your server, and more with Amazon Lightsail snapshots.
- **Access to AWS Services** – Extend the capabilities of your Amazon Lightsail server by connecting it to popular AWS Cloud services, including managed databases, CDN, and many others.

44.2 Backup/Restore and Disaster Recovery

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

44.3 Pricing Overview

Please visit <https://aws.amazon.com/pricing/> for further details on pricing AWS services.

44.4 Service Constraints

Please see <https://amazonlightsail.com/> for more information.

44.5 Technical Requirements

Amazon Lightsail is the easiest way to get started with AWS for developers who just need virtual private servers. Lightsail includes everything you need to launch your project quickly—a virtual machine, SSD-based storage, data transfer, DNS management, and a static IP—for a low, predictable price.

- **Administration Guide** – Describes how to get started with Amazon Lightsail instances to create a development environment or an application. Quickly configure the domain name and back it up using a snapshot. [HTML](#)
- **API References** – Describes all the API operations for Amazon Lightsail in detail. Also provides sample requests, responses, and errors for the supported web services protocols. [HTML](#) | [PDF](#)

45.0 Service Definition – Amazon Machine Learning

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

45.1 Service Overview

Amazon Machine Learning is a service that makes it easy for developers of all skill levels to use ML technology. Amazon Machine Learning provides visualization tools and wizards that guide you through the process of creating ML models without having to learn complex ML algorithms and technology. Once your models are ready, Amazon Machine Learning makes it easy to obtain predictions for your application using simple APIs, without having to implement custom prediction generation code or manage any infrastructure.

Amazon Machine Learning is based on the same proven, highly scalable, ML technology used for years by Amazon's internal data scientist community. The service uses powerful algorithms to create ML models by finding patterns in your existing data. Then, Amazon Machine Learning uses these models to process new data and generate predictions for your application.

Amazon Machine Learning is highly scalable and can generate billions of predictions daily and serve those predictions in real time and at high throughput. With Amazon Machine Learning, there is no up-front hardware or software investment, and you pay as you go, so you can start small and scale as your application grows.

Top benefits include:

- **Easily Create ML Models** – Amazon Machine Learning APIs and wizards make it easy for any developer to create and fine-tune ML models from data stored in [Amazon S3](#), [Amazon Redshift](#), or MySQL databases in [Amazon RDS](#), and query these models for predictions. The service's built-in data processors, scalable ML algorithms, interactive data and model visualization tools, and quality alerts help you build and refine your models quickly.
- **From Models to Predictions in Seconds** – Amazon Machine Learning is a managed service that provides end-to-end model creation, deployment, and monitoring. Once your model is ready, you can quickly and reliably generate predictions for your applications, eliminating the time and investment needed to build, scale, and maintain ML infrastructure.
- **Scalable, High-Performance Prediction Generation Service** – Amazon

Machine Learning prediction APIs can be used to generate billions of predictions for your applications. You can request predictions for large numbers of data records all at once using the batch prediction API or use the real-time API to obtain predictions for individual data records and use them within interactive web, mobile, or desktop applications.

- **Low Cost and Efficient** – With Amazon Machine Learning there is no setup cost and you pay as you go, so you can start small and scale as your application grows.
- **Leverage Proven Technology** – Amazon Machine Learning is based on the same proven, highly scalable, ML technology used by Amazon to perform critical functions like supply chain management, fraudulent transaction identification, and catalogue organisation.

45.2 Backup/Restore and Disaster Recovery

This requirement is not applicable for Amazon Machine Learning. For additional information beyond what is described herein, please refer to <http://aws.amazon.com/documentation/>.

45.3 Pricing Overview

Please visit <https://aws.amazon.com/pricing/> for further details on pricing AWS services.

45.4 Service Constraints

Please see <https://aws.amazon.com/machine-learning/> for more information.

45.5 Technical Requirements

Amazon Machine Learning makes it easy for developers to build smart applications, including applications for fraud detection, demand forecasting, targeted marketing, and click prediction. The powerful algorithms of Amazon Machine Learning create ML models by finding patterns in your existing data. The service uses these models to process new data and generate predictions for your application.

- **Developer Guide** – Provides a conceptual overview of Amazon Machine Learning and includes detailed instructions for using the service. [HTML](#) | [PDF](#)
- **API Reference** – Describes all the API operations for Amazon Machine Learning in detail. Also provides sample requests and responses for supported web service protocols. [HTML](#) | [PDF](#)

46.0 Service Definition – Amazon Macie

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

46.1 Service Overview

Amazon Macie is a security service that uses machine learning to automatically discover, classify, and protect sensitive data in AWS. Amazon Macie recognizes sensitive data such as personally identifiable information (PII) or intellectual property, and provides you with dashboards and alerts that give visibility into how this data is being accessed or moved. The fully managed service continuously monitors data access activity for anomalies, and generates detailed alerts when it detects risk of unauthorized access or inadvertent data leaks. Today, Amazon Macie is available to protect data stored in Amazon S3, with support for additional AWS data stores coming later this year. Top benefits include:

- **Superior Visibility Of Your Data** - Amazon Macie makes it easy for security administrators to have management visibility into data storage environments, beginning with S3, with additional AWS data stores coming soon.
- **Simple To Set Up, Easy To Manage** - Getting started with Amazon Macie is fast and easy. Simply log into the AWS console, select the Amazon Macie service, and provide the AWS accounts you would like to protect.
- **Data Security Automation Through Machine Learning** - Amazon Macie uses machine learning to automate the process of discovering, classifying, and protecting data stored in AWS. This helps you better understand where sensitive information is stored and how it's being accessed, including user authentications and access patterns.
- **Custom Alert Monitoring With Cloudwatch** - Amazon Macie can send all findings to Amazon CloudWatch Events. This allows you to build custom remediation and alert management for your existing security ticketing systems

46.2 Backup/Restore and Disaster Recovery

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

46.3 Pricing Overview

Please visit <https://aws.amazon.com/pricing/> for further details on pricing AWS services.

46.4 Service Constraints

Please see <http://aws.amazon.com/inspector/> for more information.

Prior to using these Services, Buyer agrees to contact Supplier before using these Services in connection with Buyer Personal Data.

46.5 Technical Requirements

Amazon Macie is a security service that uses machine learning to automatically discover, classify, and protect sensitive data in AWS. Amazon Macie recognizes

sensitive data such as personally identifiable information (PII) or intellectual property, and provides you with dashboards and alerts that give visibility into how this data is being accessed or moved.

- **User Guide** – Describes key concepts for Amazon Macie and provides instructions for using the Amazon Macie console. [HTML](#) | [PDF](#) | [Kindle](#).

47.0 Service Definition – Amazon Managed Blockchain

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

47.1 Service Overview

Amazon Managed Blockchain is a fully managed service that makes it easy to create and manage scalable blockchain networks using the popular open source frameworks Hyperledger Fabric and Ethereum (coming soon).

Top features include:

- **Fully managed** – Getting started with Amazon Managed Blockchain is easy; you can launch a blockchain network in minutes without additional configuration.
- **Choice of Hyperledger Fabric or Ethereum** – With Amazon Managed Blockchain you can choose between two popular blockchain frameworks, Hyperledger Fabric and Ethereum, so you can choose the framework that best fits your needs.
- **Easy to scale** – Managed Blockchain provides APIs that let you quickly create new nodes to meet the changing demands of your application. Also, Managed Blockchain provides a selection of instance families (bc.t3, bc.m5, and bc.c5) that comprise varying combinations of CPU and memory and give you the flexibility to choose the appropriate mix of resources to support your workload.
- **Backed by AWS Key Management Service (AWS KMS)** – Amazon Managed Blockchain uses AWS KMS to secure Hyperledger Fabric's certificate authority, a component that manages user identities and issues enrolment certificates for securely communicating within the blockchain network.
- **Reliability** – Amazon Managed Blockchain's ordering service is built using QLDB technology, which has an immutable change log and maintains the complete history of all uncommitted transactions in the blockchain network, making the ordering service more durable.

Top benefits include:

- **Fully managed** – With Amazon Managed Blockchain, you can quickly create blockchain networks that span multiple AWS accounts, enabling a group of members to execute transactions and share data without a central authority. Unlike self-hosting your blockchain infrastructure, Amazon Managed Blockchain eliminates the need for manually provisioning hardware, configuring software,

and setting up networking and security components.

- **Monitoring** – Amazon Managed Blockchain helps you to continuously monitor the blockchain network to adapt to the changing needs of your applications quickly. With Managed Blockchain's voting API, network participants can vote to add or remove members. Once a new member is added, Managed Blockchain lets that member launch and configure multiple blockchain peer nodes to process transaction requests and store a copy of the ledger. Managed Blockchain also monitors the network and automatically replaces poorly performing nodes.
- **Easily analyse blockchain activity** – Amazon Managed Blockchain can replicate network activity to Amazon QLDB, a fully managed ledger database. By storing network activity data in QLDB, you get a complete, immutable history of your network activity securely stored outside the blockchain network. This eliminates the need for custom development to extract blockchain network activity data for analysis and optimisation.
- **Insights** – Once data is stored in QLDB, you can query it and gain advanced insights, such as who is transacting and their transaction rate.
- **Flexibility** – Amazon Managed Blockchain supports two popular blockchain frameworks: Hyperledger Fabric and Ethereum. Hyperledger Fabric is well suited for applications that require stringent privacy and permission controls with a known set of members, for example, a financial application where certain trade-related data is only shared with select banks. Ethereum is well suited for highly distributed blockchain networks where transparency of data for all members is important, for example, a customer loyalty blockchain network that allows any retailer in the network to independently verify a user's activity across all members to redeem benefits.
- **Reliability** – Amazon Managed Blockchain improves the reliability of the "ordering service," a component in the Hyperledger Fabric framework that ensures delivery of transactions across the blockchain network. Hyperledger Fabric's default ordering service does not store a complete history of transactions, making it hard to keep track of and recover transaction history when needed. Managed Blockchain's ordering service is built using QLDB technology and has an immutable change log that accurately maintains the complete history of all transactions in the blockchain network, ensuring that you durably save this data.
- **Scalable** – Amazon Managed Blockchain can easily scale your blockchain network as the usage of applications on the network grows over time. When a network member requires additional capacity for creating and validating transactions, the member can quickly add a new peer node using Managed Blockchain's APIs. Managed Blockchain provides a selection of instance types that comprise varying combinations of CPU and memory to give you the flexibility to choose the appropriate mix of resources for your workload.
- **Secure** – Managed Blockchain secures your network's certificates with AWS

KMS, eliminating the need for you to set up your own secure key storage.

- **Pay for use** – There is no upfront commitment with Amazon Managed Blockchain. For Hyperledger Fabric on Amazon Managed Blockchain, you simply pay an hourly charge (billed per second) for your network membership, peer nodes, and peer node storage, and you pay for the amount of data you write to the network. Amazon Managed Blockchain offers two editions, the Standard Edition and the Starter Edition, and each edition has a different membership hourly rate. Additionally, you pay standard data transfer rates. To interact with your Amazon Managed Blockchain resources, you will need a VPC PrivateLink endpoint that is billed separately.
- **Familiar interactions** – You can access Amazon Managed Blockchain from the AWS Management Console, AWS Command Line Interface (AWS CLI), or AWS Software Development Kit (SDK). To interact with the Hyperledger Fabric components provisioned and managed by Amazon Managed Blockchain, such as the certificate authority, ordering service, and peer nodes, you can use the open source Hyperledger Fabric CLI and SDK. Amazon Managed Blockchain provides endpoints to access these services, and you create a VPC PrivateLink endpoint for your network to access these endpoints.

47.2 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.

47.3 Pricing Overview

Please visit <https://aws.amazon.com/pricing/> for further details on pricing AWS services.

47.4 Service Constraints

This requirement is not applicable to Amazon Managed Blockchain. For additional information beyond what is described herein, please refer to <https://docs.aws.amazon.com/managed-blockchain/latest/managementguide/known-issues-limits.html>

47.5 Technical Requirements

Please see <https://docs.aws.amazon.com/managed-blockchain/latest/managementguide> for comprehensive technical documentation regarding Amazon Managed Blockchain.

48.0 Service Definition – Amazon MQ

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

48.1 Service Overview

Amazon MQ is a managed message broker service for Apache ActiveMQ that makes it easy to set up and operate message brokers in the cloud. Message brokers allow different software systems—often using different programming languages, and on different platforms—to communicate and exchange information. Messaging is the communications backbone that connects and integrates the components of distributed applications, such as order processing, inventory management, and order fulfilment for e-commerce. Amazon MQ manages the administration and maintenance of ActiveMQ, a popular open-source message broker. The underlying infrastructure is automatically provisioned for high availability and message durability to support the reliability of your applications. With Amazon MQ, you get direct access to the ActiveMQ console and industry standard APIs and protocols for messaging, including JMS, NMS, AMQP, 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.

Top features include:

- **Fully Managed** - Amazon MQ makes it easy to setup and operate message brokers in the cloud. You can use the AWS Management Console, CLI, or API calls to launch a production-ready message broker in minutes. There is no need to provision hardware, and no need to install and maintain ActiveMQ software. Amazon MQ manages the set up and ongoing administrative tasks such as software upgrades, security updates, and failure detection and recovery. It's integrated with Amazon CloudWatch so you can monitor metrics and generate alarms to get alerts about potential issues. For example, you can monitor the depth of a queue, or generate alarms if messages are not getting through.
- **Easy Migration** - Amazon MQ makes it easy to migrate messaging to the cloud while preserving the existing connections between your applications. It supports industry-standard APIs and protocols for messaging, including JMS, NMS, AMQP, STOMP, MQTT, and WebSocket. This enables you to move from any message broker that uses these standards to Amazon MQ without having to rewrite any messaging code in your applications. In most cases, you can simply update the endpoints of your Amazon MQ broker to connect to your existing applications, and start sending messages.
- **Available and Durable** - Amazon MQ provides high availability and message durability. It runs on the same highly reliable infrastructure used by other Amazon Web Services. Amazon MQ always stores messages redundantly across multiple Availability Zones (AZs). Active/standby brokers are designed for high availability. In the event of a failure of the broker, or even a full AZ outage,

Amazon MQ automatically fails over to the standby instance, so you can continue sending and receiving messages.

- **Low Cost** - With Amazon MQ, you pay only for what you use. There are no minimum fees or upfront commitments. You pay for the number of hours your broker instance runs, and the storage you use monthly. Amazon MQ is free to try. The AWS Free Tier includes up to 750 hours of a single-instance mq.t2.micro broker and up to 1GB of storage per month for one year. There is no infrastructure to maintain, so you can repurpose or retire on-premises servers, storage, and load balancers associated with your message broker.

48.2 Backup/Restore and Disaster Recovery

Amazon MQ runs on the same highly reliable infrastructure used by other Amazon Web Services. Amazon MQ stores your messages redundantly across multiple Availability Zones (AZs). Active/standby brokers are designed for high availability. In the event of a failure of the broker, or even a full AZ outage, Amazon MQ automatically fails over to the standby broker so you can continue sending and receiving messages. For more information, please see <https://docs.aws.amazon.com/amazon-mq/latest/developerguide/welcome.html>.

48.3 Pricing Overview

With Amazon MQ, you pay only for what you use. There are no minimum fees or upfront commitments. Please visit <https://aws.amazon.com/pricing/> for further details on pricing AWS services.

48.4 Service Constraints

Please see <https://aws.amazon.com/amazon-mq/> for more information.

48.5 Technical Requirements

Amazon MQ is a managed message broker service for Apache ActiveMQ that makes it easy to set up and operate message brokers in the cloud. Amazon MQ provides interoperability with your existing applications and services. Amazon MQ works with your existing applications and services without the need to manage, operate, or maintain your own messaging system.

- **Developer Guide** - Provides a conceptual overview of Amazon MQ and includes detailed instructions for creating and managing brokers and migrating from on premises brokers. [HTML](#) | [PDF](#) | [Kindle](#)
- **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. [HTML](#) | [PDF](#)

49.0 Service Definition – Amazon Neptune

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

49.1 Service Overview

Amazon Neptune is a fast, reliable, fully managed graph database service that makes it easy to build and run applications that work with highly connected datasets. The core of Amazon Neptune is a purpose-built, high-performance graph database engine optimised for storing billions of relationships and querying the graph with milliseconds latency. Amazon Neptune supports popular graph models Property Graph and W3C's RDF, and their respective query languages Apache TinkerPop Gremlin and SPARQL, allowing you to easily build queries that efficiently navigate highly connected datasets. Neptune powers graph use cases such as recommendation engines, fraud detection, knowledge graphs, drug discovery, and network security

Top features include:

- **High performance and scalability** – High Throughput, Low Latency for Graph Queries, Easy Scaling of Database Compute Resources, Storage that Automatically Scales, Low Latency Read Replicas.
- **High availability and durability** – Automated Instance Monitoring and Repair, Multi-AZ Deployments with Read Replicas, Fault-tolerant and Self-healing Storage, Database Snapshots
- **Open graph APIs** – Supports Property Graph Apache TinkerPop Gremlin, Supports W3C's Resource Description Framework (RDF) 1.1 and SPARQL 1.1
- **Highly secure** – Network Isolation, Resource-Level Permissions, Encryption, Advanced Auditing
- **Fast parallel bulk data loading** – Property Graph Bulk Loading, RDF Bulk Loading

Top benefits include:

- **Supports open graph APIs** – Amazon Neptune supports open graph APIs for both Gremlin and SPARQL and provides high performance for both of these graph models and their query languages. It lets you choose the Property Graph model and its open source query language Apache TinkerPop Gremlin or the W3C standard RDF model and its standard query language SPARQL.
- **High performance and scalability** – Amazon Neptune is a purpose-built, high-performance graph database. It is optimised for processing graph queries. Neptune supports up to 15 low-latency read replicas across three Availability Zones to scale read capacity and execute more than 100,000 graph queries per second. You can easily scale your database deployment up and down from smaller to larger instance types as your needs change.
- **High availability and durability** – Amazon Neptune is highly available, durable,

and ACID (Atomicity, Consistency, Isolation, Durability) compliant. Neptune is designed to provide greater than 99.99% availability. It features fault-tolerant and self-healing storage built for the cloud that replicates six copies of your data across three Availability Zones. Neptune continuously backs up your data to [Amazon S3](#) and transparently recovers from physical storage failures. For high availability, instance failover typically takes less than 30 seconds.

- **Highly secure** – Amazon Neptune provides multiple levels of security for your database, including network isolation using [Amazon VPC](#), support for IAM authentication for endpoint access, HTTPS-encrypted client connections, and encryption at rest using keys you create and control through [AWS KMS](#). On an encrypted Neptune instance, data in the underlying storage is encrypted, as are the automated backups, snapshots, and replicas in the same cluster.
- **Fully managed** – With Amazon Neptune, you don't need to worry about database management tasks such as hardware provisioning, software patching, setup, configuration, or backups. Neptune automatically and continuously monitors and backs up your database to [Amazon S3](#), enabling granular point-in-time recovery. You can monitor database performance using [Amazon CloudWatch](#).

49.2 Backup/Restore and Disaster Recovery

Amazon Neptune's backup capability enables point-in-time recovery for your instance. This allows you to restore your database to any second during your retention period, up until the last five minutes. Your automatic backup retention period can be configured to be up to 35 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.

49.3 Pricing Overview

Please visit <https://aws.amazon.com/pricing/> for further details on pricing AWS services.

49.4 Service Constraints

Please see <https://docs.aws.amazon.com/neptune/latest/userguide/limits.html> for more information.

49.5 Technical Requirements

Please refer to <https://docs.aws.amazon.com/neptune/latest/userguide/intro.html> and the following links for comprehensive technical documentation regarding Amazon Neptune.

50.0 Service Definition – Amazon Personalize

Amazon Personalize is a machine learning service that makes it easy for developers to create individualized recommendations for customers using their applications. With Amazon Personalize, you provide an activity stream from your application – clicks, page

views, signups, purchases, and so forth – as well as an inventory of the items you want to recommend, such as articles, products, videos, or music. Amazon Personalize will process and examine the data, identify what is meaningful, select the right algorithms, and train and optimize a personalization model that is customized for your data. You can start serving personalized recommendations via a simple API call.

50.1 Service overview

Top Features include

- Combine user interaction data with contextual data to generate high-quality recommendations
- Based on the same technology used at Amazon.com on over twenty years of personalization experience. Amazon Personalize provides a range of algorithms suited for user personalization, similar item and reranking use cases.
- Allows you to send user events in real time and generate recommendations which respond to real time user activity. Amazon Personalize learns from every user interaction and continually improves to meet your business objectives.
- Can be easily integrated into websites, mobile apps, or content management and email marketing systems. You simply call the Amazon Personalize APIs and the service will output item recommendations or a reranked item list in a JSON format, which you can use in your application.

Top benefits include

- Create high-quality recommendations - Delivering personalization to individuals at scale requires a combination of the right data and the right technology. The algorithms used by Amazon Personalize are designed to overcome common problems when creating custom recommendations – such as new users with no data, popularity biases, and evolving intent of users – to deliver high-quality recommendations that respond to specific needs, preferences, and behaviour of your users.
- Personalize every touchpoint along the user journey - Amazon Personalize enables companies to provide a cohesive and unique experience for every user across all channels and devices. Personalized recommendations from the model can be easily integrated into websites, mobile apps, or content management and email marketing systems, via a simple API call. Everything from on-site search, product sorting, recommendations and offers and can be tailored to individual users.
- Own the moment with real-time recommendations - Amazon Personalize can blend real-time user activity data with existing user profile and product information to identify the right product recommendations for your users at that moment. With Amazon Personalize you can also easily add real-time personalization to your applications, to surface the most relevant video or article to a user.

- Deliver personalization within days, not months - With Amazon Personalize, you can generate a custom personalization model in just a few clicks. Amazon Personalize automates and accelerates the complex machine learning tasks required to build, train, tune, and deploy a personalization model – so you can start delivering relevant experiences for your users quickly

50.2 Backup/Restore and Disaster Recovery

This requirement is not applicable for Amazon Personalize. For additional information beyond what is described herein, please refer to <https://docs.aws.amazon.com/index.html>.

The AWS global infrastructure is built around AWS Regions and Availability Zones. AWS Regions provide multiple physically separated and isolated 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 centre infrastructures.

For more information about AWS Regions and Availability Zones, see AWS Global Infrastructure (<http://aws.amazon.com/about-aws/global-infrastructure/>).

50.3 Pricing Overview

Please see the AWS UK G Cloud 12 Pricing Document affiliated with this service in the Digital Marketplace.

50.4 Service Constraints

Please see <https://docs.aws.amazon.com/personalize/latest/dg/limits.html> for more information.

50.5 Technical Requirements

Please refer to <https://docs.aws.amazon.com/personalize/index.html> and the following links for comprehensive technical documentation regarding Amazon Personalize.

- <https://docs.aws.amazon.com/personalize/latest/dg/what-is-personalize.html>
- <https://docs.aws.amazon.com/personalize/latest/dg/how-it-works.html>
- <https://docs.aws.amazon.com/personalize/latest/dg/security.html>
- https://docs.aws.amazon.com/personalize/latest/dg/API_Reference.html

51.0 Service Definition – Amazon Pinpoint

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

51.1 Service Overview

Amazon Pinpoint makes it easy to run targeted campaigns to drive user engagement in mobile apps. Amazon Pinpoint helps you understand user behaviour, define which users to target, determine which messages to send, schedule the best time to deliver the messages, and then track the results of your campaign.

Targeted push notifications based on app usage trends and user behaviour have become a popular approach for mobile app user engagement because response rates are often several times higher than traditional email marketing campaigns. By using targeted push notifications, you can increase message relevance and effectiveness, measure engagement, and continually improve your campaigns.

Getting started with Amazon Pinpoint is easy. First, AWS Mobile Hub guides you through the process to integrate the AWS Mobile SDK with your app. Next, you define your target segments, campaign message, and specify the delivery schedule. Once your campaign is running, Amazon Pinpoint provides metrics so you can run analytics and track the impact of your campaign.

With Amazon Pinpoint, there are no upfront setup costs and no fixed monthly cost. You only pay for the number of users your campaign targets, the messages you send, and the events you collect, so you can start small and scale as your application grows.

Top benefits include:

- **Understand User Behaviour** – Amazon Pinpoint gives you a clear view of how users are engaging with your mobile app. Amazon Pinpoint includes real-time analytics with dashboards for analysing user acquisition, user engagement, monetization, user demographics, custom events, and campaign funnels so you can understand how users engage with your application. With Pinpoint, you can filter user data by entire segments, segmentation attributes, or time, and then view and drill down on the data in the Pinpoint dashboard.
- **Create Targeted Campaigns** – Amazon Pinpoint lets you create targeted campaigns that deliver relevant and personalised push notification messages to targeted segments of your app's user base. Segmentation for the audience of your campaign can be defined by common metrics like the number of times your app is opened, or custom metrics like usage of specific features in your app. You can use data from a variety of different sources to define your target segments in Amazon Pinpoint. Pinpoint can automatically collect user data from your mobile app, or import data collected in other AWS Cloud services such as Amazon S3 and Amazon Redshift. Integration with data in third party data sources is easy. Just export your data to Amazon S3 and use the pre-built integration to import your data into Pinpoint. Once you define your target segments you can run push notification campaigns that can be customised and integrated directly into the user experiences of your mobile app.
- **Measure Results** – Amazon Pinpoint provides reporting on app usage activity to track the impact of your campaign on user engagement, including the number of

times the app was opened as a result of the campaign, and revenue generated from campaigns. You can also export the resulting event data and run custom analytics using your existing analytics tools. Pinpoint can also help you A/B test different messages, track results, and then send the best message to your target segment.

51.2 Backup/Restore and Disaster Recovery

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

51.3 Pricing Overview

With Amazon Pinpoint, you pay only for what you use. There is no setup cost and no minimum fee. Amazon Pinpoint charges for the number of users your campaign targets, the number of push notifications sent, and number of events collected. The Amazon Pinpoint free tier includes 5,000 free targeted users per month, 1M free push notifications per month, and 100M events per month. Please refer to the AWS UK G-Cloud 12 Pricing Document affiliated with this service in the Digital Marketplace.

51.4 Service Constraints

Please see <https://aws.amazon.com/pinpoint/> for more information.

51.5 Technical Requirements

Amazon Pinpoint is a campaign management web service for mobile apps. Use Amazon Pinpoint to understand app user behaviour and define which users to engage. Amazon Pinpoint helps you determine which push notifications to send, decide when to deliver the notifications, and track the results of your campaign.

- **User Guide** – Describes key concepts for Amazon Pinpoint and provides instructions for using the Amazon Pinpoint console. [HTML](#) | [PDF](#)
- **Developer Guide** – Describes how to integrate Amazon Pinpoint functionality into your app and includes development instructions for its features. [HTML](#) | [PDF](#)
- **API Reference** – Describes all the Amazon Pinpoint REST API actions. [HTML](#) | [PDF](#)

52.0 Service Definition – Amazon Polly

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

52.1 Service Overview

Amazon Polly is a service that turns text into lifelike speech. Amazon Polly lets you create applications that talk, enabling you to build entirely new categories of speech enabled products. Amazon Polly is an Amazon AI service that uses advanced deep

learning technologies to synthesise speech that sounds like a human voice. Amazon Polly includes 47 lifelike voices spread across 24 languages, so you can select the ideal voice and build speech-enabled applications that work in many different countries.

Amazon Polly delivers the consistently fast response times required to support real-time, interactive dialog. You can cache and save Amazon Polly's speech audio to replay offline or redistribute. And Amazon Polly is easy to use—you simply send the text you want converted into speech to the Amazon Polly API, and Polly immediately returns the audio stream to your application so your application can play it directly or store it in a standard audio file format, such as MP3.

With Amazon Polly, you only pay for the number of characters you convert to speech, and you can save and replay Amazon Polly's generated speech. Amazon Polly's low cost per character converted, and lack of restrictions on storage and reuse of voice output, make it a cost-effective way to enable Text-to-Speech everywhere.

Top benefits include:

- **Natural Sounding Voices** – Amazon Polly provides 47 lifelike voices and supports 24 languages, including a wide range of male and female voices with a variety of accents. Amazon Polly's fluid pronunciation of text in multiple languages enables you to deliver high-quality voice output and create applications for global users.
- **Store and Redistribute Speech** – Unlike other solutions that require a royalty or charge a fee every time you replay previously generated audio, Amazon Polly allows for unlimited replays without any additional fees. These free replays extend to offline use as well. You can create speech files in a variety of standard formats, such as MP3 and OGG, and store these on devices such as a mobile phones or IoT devices for offline playback.
- **Fast Response** – Delivering lifelike voices and conversational user experiences requires consistently fast response times. Voice-enabled applications need to play synthesised speech without delay. Consider apps that provide spoken directions for navigation, eLearning applications that provide verbal instruction to students, and apps that engage the user through real time dialog. These apps are most effective when responses can start without perceived delays in the conversational flow. Even when you send lengthy text to Amazon Polly's API, it returns the audio to your application as a stream so you can play the voices immediately. These kinds of dynamic, spoken responses require access to a much larger quantity of speech audio than is typically available to store on users' devices. Amazon Polly is in the cloud, so you have access to a wide variety of synthesised speech. With Amazon Polly, your application can provide even more valuable responses that include real-time data.
- **Easy Integration** – Amazon Polly makes it easy to add voice to your website, mobile app, or device. With Polly, you just write the text you want converted to speech to the Amazon Polly API and it immediately returns the audio stream. Unlike other solutions that require a lengthy approval process, Polly doesn't

require you to describe how you will use Amazon Polly's speech in your application, and there are no distribution agreements to sign, so you can start right away.

- **Low Cost** – Amazon Polly's pay-as-you-go pricing, low cost per character converted, and unlimited replays make it a cost-effective way to enable speech synthesis in virtually any application.

52.2 Backup/Restore and Disaster Recovery

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

52.3 Pricing Overview

Please visit <https://aws.amazon.com/pricing/> for further details on pricing AWS services.

52.4 Service Constraints

Please see <https://aws.amazon.com/polly/details/> for more information. Prior to using these Services, Buyer agrees to contact Supplier before using these Services in connection with Buyer Personal Data.

52.5 Technical Requirements

Amazon Polly is a Text-to-Speech cloud service that converts text into lifelike speech. You can use Amazon Polly to develop applications that increase engagement and accessibility. Amazon Polly supports multiple languages and includes a variety of lifelike voices, so you can build speech-enabled applications that work in multiple locations and use the ideal voice for your customers.

- **Developer Guide** – Provides a conceptual overview of Amazon Polly, including detailed instructions for using the various features. [HTML](#) | [PDF](#)

53.0 Service Definition – Amazon Quantum Ledger Database (QLDB)

Amazon QLDB is a fully managed ledger database that provides a transparent, immutable, and cryptographically verifiable transaction log owned by a central trusted authority. Amazon QLDB tracks each and every application data change and maintains a complete and verifiable history of changes over time.

Amazon QLDB is a new class of database that eliminates the need to engage in the complex development effort of building your own ledger-like applications. With QLDB, your data's change history is immutable – it cannot be altered or deleted – and using cryptography, you can easily verify that there have been no unintended modifications to your application's data. QLDB uses an immutable transactional log, known as a journal, that tracks each application data change and maintains a complete and verifiable history of changes over time. QLDB is easy to use because it provides developers with

a familiar SQL-like API, a flexible document data model, and full support for transactions. QLDB is also serverless, so it automatically scales to support the demands of your application. There are no servers to manage and no read or write limits to configure. With QLDB, you only pay for what you use.

53.1 Service Overview

Top Features include:

- 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
- Amazon QLDB uses cryptography to create a concise summary of your change history
- Seamless, automatic scaling to meet the demands of your application without the need to provision capacity or configure read and write limits; no servers to manage or capacity to provision
- 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
- Provides atomicity, consistency, isolation, and durability (ACID) properties. Also, QLDB transactions have full serializability- the highest level of isolation

Top benefits include

- Amazon QLDB uses a journal that tracks each application data change and maintains a complete and sequenced history of changes over time. Data on the journal cannot be deleted or modified. The full history of your database can be accessed and you can query and analyse the history to see how your data has changed over time.
- With Amazon QLDB, you can trust that the history of changes to your application data is accurate. QLDB uses a cryptographic hash function (SHA-256) to generate a secure output file of your data's change history, known as a digest. The digest acts as a proof of your data's change history, allowing you to look back and validate the integrity of your data changes.
- Amazon QLDB is highly scalable and can execute 2 – 3X as many transactions than ledgers in common blockchain frameworks. Blockchain frameworks are decentralized so to execute a transaction, they require a majority of members of the network to reach consensus on the validity of the transaction. On the other hand, QLDB has a centralized design, allowing its transactions to execute without the need for multi-party consensus.
- With Amazon QLDB, you don't have to worry about provisioning capacity or configuring read and write limits. You create a ledger, define your tables, and QLDB automatically scales to support the demands of your application. To help

you gain a better understanding of the operational health of your database, QLDB also allows you to monitor operational metrics for your read and write IOs.

- Amazon QLDB's familiar database capabilities make it easy to use. QLDB supports PartiQL - a new, open source, SQL-compatible query language designed to easily work with all data types and structures. With PartiQL, you can easily query, manage, and update your data with SQL operators. QLDB's document-oriented data model is flexible, enabling you to easily store and process both structured and semi-structured data. Also, QLDB transactions are ACID compliant and have full serializability- the highest level of isolation.
- Amazon QLDB is designed for high availability, replicating multiple copies of data within an Availability Zone (AZ) as well as across 3 AZs in an AWS region, without any additional cost or setup. QLDB backs up your data continuously while maintaining consistent performance, allowing it to transparently recover from any instance or physical storage failures.

53.2 Backup/Restore and Disaster Recovery

Amazon QLDB does not support a backup and restore feature as of now. At present, an export to S3 functionality is available. Using this functionality you can export the contents of your QLDB journal to S3.

Amazon QLDB's ledger is deployed across multiple AZs with multiple copies per AZ. We maintain redundancy within the region and ensure full recovery from availability zone failures. A write is acknowledged only after being written to a durable storage in multiple AZs, and hence, QLDB is strongly durable.

Amazon QLDB does not support cross-region replication as of now. QLDB's export to S3 feature enables customers to export the contents of the QLDB journal to a S3 bucket. The S3 buckets can be configured for cross-region replication.

53.3 Pricing Overview

Please see the AWS UK G Cloud 12 Pricing Document affiliated with this service in the Digital Marketplace.

53.4 Service Constraints

Please see

https://docs.aws.amazon.com/general/latest/gr/aws_service_limits.html#limits_qldb for more information.

53.5 Technical Requirements

Please refer to <https://docs.aws.amazon.com/qldb/index.html> and the following links for comprehensive technical documentation regarding Amazon Quantum Ledger Database.

- <https://docs.aws.amazon.com/qldb/latest/developerguide/what-is.html>

- <https://docs.aws.amazon.com/qldb/latest/developerguide/getting-started.html>
- <https://docs.aws.amazon.com/qldb/latest/developerguide/security.html>
- <https://aws.amazon.com/qldb/faqs/>

54.0 Service Definition – Amazon QuickSight

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

54.1 Service Overview

Amazon QuickSight is a very fast, cloud-powered Business Intelligence (BI) service that makes it easy for all employees to build visualizations, perform analysis, and quickly get business insights from their data. Amazon QuickSight uses a new, Super-fast, Parallel, In-memory Calculation Engine (“SPICE”) to perform advanced calculations and render visualizations rapidly. Amazon QuickSight integrates automatically with AWS data services, enables organisations to scale to hundreds of thousands of users, and delivers fast and responsive query performance to them via SPICE’s query engine. At one tenth the cost of traditional solutions, Amazon QuickSight enables you to deliver rich BI functionality to everyone in your organisation.

Top benefits include:

- **Get Started Quickly** – No need to spend months building complex data models and invest in complex and costly software or hardware to generate the first report. Simply log in to Amazon QuickSight, point to a data source or upload a file, and begin visualizing your data with actionable insights in about a minute.
- **Access Data from Multiple Sources** – Easily connect Amazon QuickSight to
- AWS data services, including Amazon Redshift, Amazon RDS, Amazon Aurora,
- Amazon EMR, Amazon DynamoDB, Amazon S3, and Amazon Kinesis; upload CSV, TSV, and spreadsheet files; or connect to third-party data sources such as Salesforce.
- **Easy to Use** – Amazon QuickSight automatically infers data types and relationships and provides suggestions for the best possible visualizations, optimised for your data, to help you get quick, actionable business insights.
- **Get Answers Fast** – Amazon QuickSight uses [SPICE](#) to generate answers on large data sets.
- **Tell a Story with Your Data** – Securely [share your analysis](#) with others in your organisation by building interactive stories for collaboration using the StoryBoard and annotations. Recipients can further explore the data and respond back with their insights and knowledge, making the whole organisation efficient and effective.
- **Use AWS Partner BI Solutions with Amazon QuickSight** – Amazon QuickSight

provides partners a simple SQL-like interface to query the data stored in SPICE so that you can continue using your existing BI tools from AWS BI partners while benefiting from the faster performance delivered by SPICE.

54.2 Backup/Restore and Disaster Recovery

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

54.3 Pricing Overview

Please visit <https://aws.amazon.com/pricing/> for further details on pricing AWS services.

54.4 Service Constraints

Please see <https://aws.amazon.com/quicksight/> for more information.

54.5 Technical Requirements

Amazon QuickSight is a fast business analytics service to build visualizations, perform ad hoc analysis, and quickly get business insights from your data. QuickSight seamlessly discovers AWS data sources, enables organisations to scale to hundreds of thousands of users, and delivers fast and responsive query performance by using a robust in-memory engine (SPICE).

- User Guide - Describes all Amazon QuickSight concepts and features, and provides instructions on using these features in the QuickSight web application.
[HTML](#) | [PDF](#) | [Kindle](#)

55.0 Service Definition – Amazon Redshift

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

55.1 Service Overview

Amazon Redshift is a fast, fully managed, petabyte-scale data warehouse service that makes it simple and cost-effective to efficiently analyse all your data using your existing business intelligence tools. It is optimised for datasets ranging from a few hundred gigabytes to a petabyte or more and costs a tenth the cost of most traditional data warehousing solutions.

Amazon Redshift delivers fast query performance by using columnar storage technology to improve I/O efficiency and parallelizing queries across multiple nodes. Amazon Redshift has custom JDBC and ODBC drivers that you can download from the Connect Client tab of our Console, allowing you to use a wide range of familiar SQL clients. You can also use standard PostgreSQL JDBC and ODBC drivers. Data load speed scales linearly with cluster size, with integrations to Amazon S3, Amazon DynamoDB, Amazon EMR, Amazon Kinesis, or any SSH-enabled host.

Amazon Redshift's data warehouse architecture allows you to automate most of the common administrative tasks associated with provisioning, configuring and monitoring a cloud data warehouse. Backups to Amazon S3 are continuous, incremental, and automatic. Restores are fast—you can start querying in minutes while your data is spooled down in the background. Enabling disaster recovery across regions takes just a few clicks.

Top features and benefits include:

- **Optimised for Data Warehousing** – Amazon Redshift uses a variety of innovations to obtain very high query performance on datasets ranging in size from a hundred gigabytes to a petabyte or more. It uses columnar storage, data compression, and zone maps to reduce the amount of I/O needed to perform queries. Amazon Redshift has a Massively Parallel Processing (MPP) data warehouse architecture, parallelizing and distributing SQL operations to take advantage of all available resources. The underlying hardware is designed for high performance data processing, using local attached storage to maximise throughput between the CPUs and drives, and a 10GigE mesh network to maximise throughput between nodes.
- **Scalable** – With a few clicks of the AWS Management Console or a simple API call, you can easily change the number or type of nodes in your cloud data warehouse as your performance or capacity needs change. Dense Storage (DS) nodes allow you to create very large data warehouses using HDDs for a very low price point. Dense Compute (DC) nodes allow you to create very high performance data warehouses using fast CPUs, large amounts of RAM and SSDs.
- **No Up-Front Costs** – You pay only for the resources you provision. You can choose On-Demand pricing with no up-front costs or long-term commitments, or obtain significantly discounted rates with Reserved Instance pricing.
- **Get Started in Minutes** – With a few clicks in the AWS Management Console or simple API calls, you can create a cluster, specifying its size, underlying node type, and security profile. Amazon Redshift will provision your nodes, configure the connections between them, and secure the cluster. Your data warehouse should be up and running in minutes.
- **Fully Managed** – Amazon Redshift handles all of the work needed to manage, monitor, and scale your data warehouse, from monitoring cluster health and taking backups to applying patches and upgrades. You can easily resize your cluster as your performance and capacity needs change. By handling all these time-consuming, labour-intensive tasks, Amazon Redshift frees you up to focus on your data and business.
- **Fault Tolerant** – Amazon Redshift has multiple features that enhance the reliability of your data warehouse cluster. All data written to a node in your cluster is automatically replicated to other nodes within the cluster and all data is continuously backed up to Amazon S3. Amazon Redshift continuously monitors

the health of the cluster and automatically re-replicates data from failed drives and replaces nodes as necessary.

- **Encryption** – With just a couple of 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. By default, Amazon Redshift takes care of key management but you can choose to manage your keys using your own HSMs, AWS CloudHSM, or AWS KMS.
- **Network Isolation** – Amazon Redshift enables you to configure firewall rules to control network access to your data warehouse cluster. You can run Amazon Redshift inside Amazon VPC to isolate your data warehouse cluster in your own virtual network and connect it to your existing IT infrastructure using industry standard encrypted IPsec VPN.
- **Audit and Compliance** – Amazon Redshift integrates with AWS CloudTrail to enable you to audit all Amazon Redshift API calls. Amazon Redshift also logs all SQL operations, including connection attempts, queries and changes to your database. You can access these logs using SQL queries against system tables or choose to have them downloaded to a secure location on Amazon S3. Amazon Redshift is compliant with Service Organisation Controls (SOC) 1, SOC 2, SOC 3, and Payment Card Industry Data Security Standard (PCI DSS) Level 1 requirements. For more details, please visit the [AWS Compliance Center](#).
- **SQL** – Amazon Redshift is a SQL data warehouse solution and uses industry standard ODBC and JDBC connections. You can download our custom JDBC and ODBC drivers from the Connect Client tab of the service console. Many popular software vendors have certified Amazon Redshift with their offerings to enable you to continue to use the tools you do today.
- **Integrated** – Amazon Redshift is integrated with other AWS Cloud services and has built in commands to load data in parallel to each node from Amazon S3 or your Amazon EC2 and on premise servers using SSH.

55.2 Backup/Restore and Disaster Recovery

Amazon Redshift's automated snapshot feature continuously backs up new data on the cluster to Amazon S3. Snapshots are continuous, incremental and automatic. Amazon Redshift stores your snapshots for a user-defined period, which can be from one to thirty-five days. You can take your own snapshots at any time, which leverage all existing system snapshots and are retained until you explicitly delete them. Redshift can also asynchronously replicate your snapshots to Amazon S3 in another region for disaster recovery. Once you delete a cluster, your system snapshots are removed but your user snapshots are available until you explicitly delete them.

You can use any system or user snapshot to restore your cluster using the AWS Management Console or the Amazon Redshift APIs. Your cluster is available as soon

as the system metadata has been restored and you can start running queries while user data is spooled down in the background.

55.3 Pricing Overview

Please visit <https://aws.amazon.com/pricing/> for further details on pricing AWS services.

55.4 Service Constraints

Please see <https://aws.amazon.com/redshift/> for more information.

55.5 Technical Requirements

Amazon Redshift is a fast, fully managed, petabyte-scale data warehouse service that makes it simple and cost-effective to efficiently analyse all your data using your existing business intelligence tools. It is optimised for datasets ranging from a few hundred gigabytes to a petabyte or more.

- **Getting Started** – Introduces you to Amazon Redshift, helps you set up an account, and walks you through a simple example to use Amazon Redshift for the first time. Also provides tips and links to advanced product features and resources. [HTML](#) | [PDF](#) | [Kindle](#)
- **Cluster Management Guide** – Shows you how to create and manage Amazon Redshift clusters. [HTML](#) | [PDF](#) | [Kindle](#)
- **API Reference** – Describes all the API operations for Amazon Redshift in detail. Also provides sample requests, responses, and errors for the supported web services protocols. [HTML](#) | [PDF](#)
- **Database Developer Guide** – Explains how to design, build, query, and maintain the databases that make up your data warehouse, and includes syntax for Amazon Redshift SQL commands and functions. [HTML](#) | [PDF](#)
- **CLI Guide** – Describes the CLI for Amazon Redshift in detail. Provides basic syntax, options, and usage examples for each command. [HTML](#)

56.0 Service Definition – Amazon Relational Database Service (Amazon RDS)

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

56.1 Service Overview

Amazon RDS is a web 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 management tasks, freeing you up to focus on your applications and business.

Top benefits include:

- **Easy to Administer** – Amazon RDS makes it easy to go from project conception to deployment. Use the [AWS Management Console](#), the [Amazon RDS CLI](#), or simple API calls to access the capabilities of a production-ready relational database in minutes. There is no need for infrastructure provisioning and no need for installing and maintaining database software.
- **Scalable** – You can scale your database's compute and storage resources with only a few mouse clicks or an API call, often with no downtime.
- **Available and Durable** – Amazon RDS runs on the same highly reliable infrastructure used by other AWS Cloud services. When you provision a MultiAZ Database Instance (DB Instance), Amazon RDS synchronously replicates the data to a standby instance in a different Availability Zone. Amazon RDS has many other features that enhance reliability for critical production databases, including automated backups, Database Snapshots (DB Snapshots), and automatic host replacement.
- **Fast** – Amazon RDS offers database server sizing choices up to 32 vCPUs and 244 GiB, as well as storage choices for a wide range of application performance requirements.
- **Secure** – Amazon RDS makes it easy to control network access to your database. Amazon RDS also lets you run your database instances in Amazon VPC, which enables you to isolate your database instances and to connect to your existing IT infrastructure through an industry-standard encrypted IPsec VPN. Many Amazon RDS engine types offer encryption at rest and encryption in transit.
- **Inexpensive** – You pay very low rates and only for the resources you actually consume. In addition, you benefit from the option of on-demand pricing with no up-front or long-term commitments or even lower hourly rates via our reserved pricing option.

56.2 Backup/Restore and Disaster Recovery

The following sections provide information regarding backup/restore and disaster recovery for Amazon RDS.

56.2.1 Database Instance Backups

Amazon RDS provides two different methods for backing up and restoring your Amazon database instances: automated backups and DB Snapshots. Automated backups automatically back up your DB instance during a specific, user-definable backup window and keeps the backups for a limited, user-specified period of time (called the backup retention period); you can later recover your database to any point in time during that retention period. DB Snapshots are user-initiated backups that enable you to back up your DB instance to a known state and to restore to that specific state at any time. Amazon RDS keeps all DB Snapshots until you delete them. A brief I/O freeze, typically lasting a few seconds, occurs during both automated backups and DB Snapshot operations on Single-AZ DB instances.

56.2.2 Automated Backup

Automated backup is an Amazon RDS feature that automatically creates a backup of your database. Automated backups are enabled by default for a new DB instance. An automated backup occurs during a daily user-configurable period of time known as the preferred backup window. Backups created during the preferred backup window are retained for a user-configurable number of days (the backup retention period).

The preferred backup window is the user-defined period of time during which your DB instance is backed up. Amazon RDS uses these periodic data backups in conjunction with your transaction logs to enable you to restore your DB instance to any second during your retention period, up to the LatestRestorableTime (typically up to the last five minutes). During the backup window, storage I/O may be suspended while your data is being backed up and you may experience elevated latency. This I/O suspension typically lasts for the duration of the snapshot. This period of I/O suspension is shorter for Multi-AZ database deployments because the backup is taken from the standby, but latency can still occur during the backup process.

When the backup retention changes to a non-zero value, the first backup occurs immediately. Changing the backup retention period to 0 turns off automatic backups for the DB instance and deletes all existing automated backups for the instance. If you don't specify a preferred backup window when you create the DB instance, Amazon RDS assigns a default 30-minute backup window that is selected at random from an 8-hour block of time per region. Changes to the backup window take effect immediately. The backup window cannot overlap with the weekly maintenance window for the DB instance.

When you delete a DB instance, you can create a final DB Snapshot upon deletion; if you take this action, you can use that DB Snapshot to restore the deleted DB instance at a later date. Amazon RDS retains this final, user-created DB Snapshot along with all other manually created DB Snapshots after the DB instance is deleted. All automated backups are deleted and cannot be recovered when you delete a DB instance.

For more information on working with automated backups, see the [Working With Automated Backups](#) user guide.

56.2.3 Point-In-Time Recovery

In addition to the daily automated backup, Amazon RDS archives database change logs. This enables you to recover your database to any point in time during the backup retention period, up to the last five minutes of database usage.

Amazon RDS stores multiple copies of your data, but for Single-AZ DB instances, these copies are stored in a single Availability Zone. If for any reason a Single-AZ DB instance becomes unusable, you can use point-in-time recovery to launch a new DB instance with the latest restorable data. For more information on working with point-in-time recovery, see the [Restoring a DB Instance to a Specified Time](#) user guide. Multi-AZ deployments store copies of your data in different Availability Zones for greater levels of data durability. For more information on Multi-AZ deployments, see the [High Availability \(Multi-AZ\)](#) user guide.

56.2.4 Automated Backups with Unsupported MySQL Storage Engines

Amazon RDS automated backups and DB Snapshots are currently supported for all database engines. For the MySQL database engine, only the InnoDB storage engine is supported; use of these features with other MySQL storage engines, including MyISAM, may lead to unreliable behaviour while restoring from backups. Specifically, since storage engines like MyISAM do not support reliable crash recovery, your tables can be corrupted in the event of a crash. For this reason, we encourage you to use the InnoDB storage engine.

If you choose to use MyISAM, you can attempt to manually repair tables that become damaged after a crash by using the REPAIR command (refer to <http://dev.mysql.com/doc/refman/5.5/en/repair-table.html> for detailed instructions). However, as noted in the MySQL documentation, there is a good chance that you will not be able to recover all of your data.

56.2.5 DB Snapshots

DB Snapshots are user-initiated and enable you to back up your DB instance in a known state as frequently as you wish and then restore to that specific state at any time. DB Snapshots can be created with the Amazon RDS console or the CreateDBSnapshot action in the Amazon RDS API. DB Snapshots are kept until you explicitly delete them with the Amazon RDS console or the DeleteDBSnapshot action in the Amazon RDS API. For more information on working with DB Snapshots, see the [Creating a DB Snapshot](#) and [Restoring From a DB Snapshot](#) user guides.

Please refer to the [Amazon RDS Backing Up and Restoring](#) user guide for additional information.

56.3 Pricing Overview

Please visit <https://aws.amazon.com/pricing/> for further details on pricing AWS services.

56.4 Service Constraints

Please see <http://aws.amazon.com/documentation/rds/> for more information.

56.5 Technical Requirements

Please refer to <http://aws.amazon.com/documentation/rds/> and the following links for comprehensive technical documentation regarding Amazon RDS.

- **User Guide** – Describes all Amazon RDS concepts and provides instructions on using the various features with both the console and the CLI. Available in HTML, PDF, and Kindle formats.
- **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. Available in HTML and PDF formats.
- **CLI Reference** – Describes all the API operations for Amazon RDS in detail.

Also provides sample requests, responses, and errors for the supported web services protocols. Available in HTML and PDF formats.

- **Quick Reference Card** – Briefly covers the essential commands for using [Amazon RDS from the command line. Available in PDF format.](#)

57.0 Service Definition – Amazon Route 53

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

This service listing includes the following AWS Services:

- Amazon Route 53
- Amazon Route 53 Resolver
- Amazon Route 53 Private DNS

57.1 Service Overview

Amazon Route 53 is a highly available and scalable cloud DNS web service. It is designed to give developers and businesses an extremely reliable and cost-effective way to route end users to Internet applications by translating names like `www.example.com` into the numeric IP addresses like `192.0.2.1` that computers use to connect to each other.

Amazon Route 53 effectively connects user requests to infrastructure running in AWS—such as Amazon EC2 instances, Elastic Load Balancing load balancers, or Amazon S3 buckets—and can also be used to route users to infrastructure outside of AWS. You can use Amazon Route 53 to configure DNS health checks to route traffic to healthy endpoints or to independently monitor the health of your application and its endpoints. Amazon Route 53 Traffic Flow makes it easy for you to manage traffic globally through a variety of routing types, including Latency Based Routing, Geo DNS, and Weighted Round Robin—all of which can be combined with DNS Failover in order to enable a variety of low-latency, fault-tolerant architectures. Using Amazon Route 53 Traffic Flow's simple visual editor, you can easily manage how your end users are routed to your application's endpoints—whether in a single AWS Region or distributed around the globe. Amazon Route 53 also offers Domain Name Registration—you can purchase and manage domain names such as `example.com` and Amazon Route 53 will automatically configure DNS settings for your domains.

Top benefits include:

- **Highly Available and Reliable** – Amazon Route 53 is built using AWS's highly available and reliable infrastructure. The distributed nature of our DNS servers helps ensure a consistent ability to route your end users to your application. Features such as Amazon Route 53 Traffic Flow help you improve reliability with

easy configuration of failover to re-route your users to an alternate location if your primary application endpoint becomes unavailable. Amazon Route 53 is designed to provide the level of dependability required by important applications. Amazon Route 53 is backed by the [Amazon Route 53 Service Level Agreement](#).

- **Scalable** – Amazon Route 53 is designed to automatically scale to handle very large query volumes without any intervention from you.
- **Designed for Use with Other AWS Cloud Services** – Amazon Route 53 is designed to work well with other AWS features and offerings. You can use Amazon Route 53 to map domain names to your Amazon EC2 instances, Amazon S3 buckets, Amazon CloudFront distributions, and other AWS resources. By using the AWS Identity and Access Management (IAM) service with Amazon Route 53, you get fine-grained control over who can update your DNS data. You can use Amazon Route 53 to map your zone apex (example.com versus www.example.com) to your Elastic Load Balancing instance, Amazon CloudFront distribution, AWS Elastic Beanstalk environment, or Amazon S3 website bucket using a feature called Alias record.
- **Simple** – With self-service sign-up, Amazon Route 53 can start to answer your DNS queries within minutes. You can configure your DNS settings with the AWS Management Console or our easy-to-use API. You can also programmatically integrate the Amazon Route 53 API into your overall web application. For instance, you can use Amazon Route 53's API to create a new DNS record whenever you create a new Amazon EC2 instance. Amazon Route 53 Traffic Flow makes it easy to set up sophisticated routing logic for your applications by using the simple visual policy editor.
- **Fast** – Using a global anycast network of DNS servers around the world, Amazon Route 53 is designed to automatically route your users to the optimal location depending on network conditions. As a result, the service offers low query latency for your end users, as well as low update latency for your DNS record management needs. Amazon Route 53 Traffic Flow lets you further improve your customers' experience by running your application in multiple locations around the world and using traffic policies to ensure your end users are routed to the closest healthy endpoint for your application.
- **Cost Effective** – Amazon Route 53 passes on the benefits of AWS's scale to you. You pay only for the resources you use, such as the number of queries that the service answers for each of your domains, hosted zones for managing
- domains through the service, and optional features such as traffic policies and health checks, all at a low cost and without minimum usage commitments or any up-front fees.
- **Secure** – By integrating Amazon Route 53 with AWS IAM, you can grant unique credentials and manage permissions for every user within your AWS account and specify who has access to which parts of the Amazon Route 53 service.
- **Flexible** – Amazon Route 53 Traffic Flow routes traffic based on multiple criteria,

such as endpoint health, geographic location, and latency. You can configure multiple traffic policies and decide which policies are active at any given time. You can create and edit traffic policies using the simple visual editor in the

- Amazon Route 53 console, AWS Software Development Kits (SDKs), or the Amazon Route 53 API. Traffic Flow's versioning feature maintains a history of changes to your traffic policies, so you can easily roll back to a previous version using the console or API.

57.2 Backup/Restore and Disaster Recovery

This requirement is not applicable for Amazon Route 53. For additional information beyond what is described herein, please refer to <http://aws.amazon.com/documentation/>.

57.3 Pricing Overview

Please visit <https://aws.amazon.com/pricing/> for further details on pricing AWS services.

57.4 Service Constraints

Please see <https://aws.amazon.com/route53/> for more information.

57.5 Technical Requirements

Amazon Route 53 is a highly available and scalable DNS web service.

- **Developer Guide** – Provides an overview of Amazon Route 53, detailed feature descriptions, procedures for using the console, and an explanation of how to use the API. [HTML](#) | [PDF](#) | [Kindle](#)
- **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. [HTML](#) | [PDF](#)

57.6 Service Overview – Amazon Route 53 Resolver

- Amazon Route 53 Resolver is a fully managed DNS recursive service. It eliminates the undifferentiated heavy lifting of running your own DNS resolvers in the cloud to provide a unified view of DNS across hybrid networks and gives you seamless query resolution for your internal DNS zones, irrespective of whether these are managed in AWS or on premises. Customers can get bi-directional query resolution—a unified view of DNS across their hybrid network—as a managed service.
- The top features include:
- **Resolver endpoints** – Inbound query capability is provided by Route 53 Resolver Endpoints, allowing DNS queries that originate on premises to resolve AWS-hosted domains.

- **Conditional forwarding rules** – Outbound DNS queries are enabled through the use of conditional forwarding rules. Domains hosted within your on-premises DNS infrastructure can be configured as forwarding rules in Route 53 Resolver.
- **Cross-account rules sharing** – Resource Access Management (RAM) uses AWS Organisations to share rules with other AWS accounts that you nominate.
- **Managed service** – With Route 53 Resolver there are no servers to provision, patch, or manage and no software to install, maintain, or operate.

Top benefits include:

- **Hybrid cloud DNS resolution** – Route 53 Resolver Endpoints provide inbound query capability, allowing DNS queries that originate on premises to resolve AWS-hosted domains. Connectivity needs to be established between your on-premises DNS infrastructure and AWS through AWS Direct Connect or a VPN. Outbound DNS queries are enabled through the use of conditional forwarding rules. Domains hosted within your on-premises DNS infrastructure can be configured as forwarding rules in Route 53 Resolver. Rules will trigger when a query is made to one of those domains and will attempt to forward DNS requests to your DNS servers that were configured along with the rules. Like the inbound queries, this requires a private connection over AWS Direct Connect or VPN.
- **Reduced complexity** – With Route 53 Resolver, you no longer have to worry about migrating complex DNS infrastructure into the cloud or building your own DNS servers to bridge the gap between the cloud and on premises. Instead, you can manage the DNS naming of your AWS resources using Amazon Route 53, you don't have to make any changes to your DNS infrastructure on premises, and your resources can query any DNS name and get the correct answer, anywhere inside your hybrid environment.
- **High availability** – Unlike other solutions that require you to run software yourself and manage launching and scaling of instances, Route 53 Resolver is a fully managed, native AWS service. It provides higher availability than third-party or custom solutions, it scales automatically to accommodate any workload, and it never requires maintenance, patching, or upgrades.
- **VPC-specific DNS resolution** – Route 53 Resolver integrates with Amazon VPC, Amazon EC2, Route 53 Private DNS, and AWS Direct Connect to provide a seamless, unified view of DNS across your network, spanning both AWS and on-premises environments. It makes it possible to resolve DNS names managed anywhere in your network, from custom names in Route 53 Private DNS, to the default instance names provided by Amazon EC2, to names for on-premises resources managed by appliances in your data centre. Route 53 Resolver gives you the flexibility to define how DNS queries should be answered based on the domain name and where in your network the query is coming from.
- **Security** – A Route 53 Resolver Endpoint includes one or more elastic network interfaces that attach to your Amazon VPC. Each elastic network interface is assigned an IP address from the subnet space of the VPC where it is located.

Each elastic network interface is assigned to a VPC and has a security group associated with it.

- **Metrics and monitoring** – You can use Amazon CloudWatch to monitor the number of DNS queries that are forwarded by Route 53 Resolver Endpoints. CloudWatch collects and processes raw data into readable, near-real-time metrics. These statistics are recorded for a period of two weeks so that you can access historical information and gain a better perspective on how your resources are performing. By default, metric data for Route 53 Resolver Endpoints is automatically sent to CloudWatch at five-minute intervals.

57.6.1 Backup/Restore and Disaster Recovery

This requirement is not applicable to Route 53 Resolver. For additional information beyond what is described herein, please refer to <http://aws.amazon.com/documentation/>.

57.6.2 Pricing Overview

Please visit <https://aws.amazon.com/pricing/> for further details on pricing AWS services.

57.6.3 Service Constraints

Please see <https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/DNSLimitations.html> for more information.

57.6.4 Technical Requirements

Please see <https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/resolver-getting-started.html> for comprehensive technical documentation regarding Route 53 Resolver.

58.0 Service Definition – Amazon SageMaker

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

This service listing includes the following AWS services:

- Amazon SageMaker Neo
- Amazon SageMaker Ground Truth

58.1 Service Overview

Amazon SageMaker is a fully-managed platform that enables developers and data scientists to quickly and easily build, train, and deploy machine learning models at any scale. Amazon SageMaker removes all the barriers that typically slow down developers who want to use machine learning.

Machine learning often feels a lot harder than it should be to most developers because the process to build and train models, and then deploy them into production is too complicated and too slow. First, you need to collect and prepare your training data to discover which elements of your data set are important. Then, you need to select which algorithm and framework you'll use. After deciding on your approach, you need to teach the model how to make predictions by training, which requires a lot of compute. Then, you need to tune the model so it delivers the best possible predictions, which is often a tedious and manual effort. After you've developed a fully trained model, you need to integrate the model with your application and deploy this application on infrastructure that will scale. All of this takes a lot of specialized expertise, access to large amounts of compute and storage, and a lot of time to experiment and optimize every part of the process. In the end, it's not a surprise that the whole thing feels out of reach for most developers.

Amazon SageMaker removes the complexity that holds back developer success with each of these steps. Amazon SageMaker includes modules that can be used together or independently to build, train, and deploy your machine learning models.

Top benefits include:

- **Get to Production with Machine Learning Quickly** - Amazon SageMaker significantly reduces the amount of time needed to train, tune, and deploy machine learning models. Amazon SageMaker manages and automates all the sophisticated training and tuning techniques so you can get models into production quickly.
- **Choose Any Framework or Algorithm** - Amazon SageMaker supports all machine algorithms and frameworks so you can use the technology you are already familiar with. Apache MXNet and TensorFlow are pre-installed, and Amazon SageMaker offers a range of built-in, high performance machine learning algorithms. If you want to train with an alternative framework or algorithm, you can bring your own in a Docker container.
- **One-Click Training and Deployment** - Amazon SageMaker lets you begin training your model with a single click in the console or with a simple API call. When the training is complete, and you're ready to deploy your model, you can launch it with a single click in the Amazon SageMaker console.
- **Easily Integrate With Your Existing Workflow** - Amazon SageMaker is designed in three modules that can be used together or independently as part of any existing ML workflow you might already have in place.
- **Easy Access to Trained Models** - Amazon SageMaker makes it easy to integrate machine learning models into your applications by providing an HTTPS endpoint that can be called from any application.
- **Optimized for Speed** - Amazon SageMaker is pre-configured with the latest versions of TensorFlow and Apache MXNet, with CUDA9 library support for maximum performance with NVIDIA GPUs. With Amazon SageMaker P3 instances running NVIDIA Volta V100 GPUs, Amazon SageMaker lets you train

deep learning models with unparalleled speed.

58.2 Backup/Restore and Disaster Recovery

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

58.3 Pricing Overview

Please visit <https://aws.amazon.com/pricing/> for further details on pricing AWS services.

58.4 Service Constraints

Please see <https://aws.amazon.com/sagemaker/> for more information.

58.5 Technical Requirements

Amazon SageMaker is a fully managed machine learning service. With Amazon SageMaker, data scientists and developers can quickly and easily build and train machine learning models, and then directly deploy them into a production-ready hosted environment. It provides an integrated Jupyter authoring notebook instance for easy access to your data sources for exploration and analysis, so you don't have to manage servers. It also provides common machine learning algorithms that are optimized to run efficiently against extremely large data in a distributed environment. With native support for bring-your-own-algorithms and frameworks, Amazon SageMaker provides flexible distributed training options that adjust to your specific workflows.

- **Developer Guide** - Provides a conceptual overview of Amazon SageMaker, offers step-by-step instructions for training and deploying models, and includes a complete API reference for developers. [HTML](#) | [PDF](#)

58.6 Service Overview – Amazon SageMaker Ground Truth

Amazon SageMaker Ground Truth helps you quickly build highly accurate training datasets for machine learning. SageMaker Ground Truth offers easy access to public and private human labellers and provides them with built-in workflows and interfaces for common labelling tasks. Additionally, SageMaker Ground Truth can lower your labelling costs by up to 70% using automatic labelling, which works by training SageMaker Ground Truth from data labelled by humans so that the service learns to label data independently

The top features include:

- **Automated data labelling** – Amazon SageMaker Ground Truth provides automated data labelling using machine learning. SageMaker Ground Truth will first select a random sample of data and send it to humans to be labelled.
- **Flexibility** – Amazon SageMaker Ground Truth supports multiple choices for human labelling directly in the SageMaker Ground Truth console. You can use

your private team of labellers for in-house labelling jobs, especially for handling data that needs to stay within your organisation.

- **Easy instructions for human labelling** – With Amazon SageMaker Ground Truth, you provide labelling guidance to human labellers to help ensure consistency. These detailed instructions are available to labellers within their labelling interface.
- **Workflows** – Amazon SageMaker Ground Truth provides built-in labelling workflows that take human labellers step by step through tasks and provide tools to help them produce good results. Built-in workflows are currently available for object detection, image classification, text classification, and semantic segmentation labelling jobs.
- **Seamless integration into Amazon SageMaker** – Training datasets created with SageMaker Ground Truth can be easily imported into Amazon SageMaker for use in model development and training.

Top benefits include:

- **Reduce labelling costs by up to 70%** – SageMaker Ground Truth uses a machine learning model to automatically label raw data to produce high-quality training datasets quickly at a fraction of the cost of manual labelling. Data is only routed to humans if the active learning model cannot confidently label it. The human-labelled data is then used to train the model to improve its capabilities. Less data is then sent to humans in the next round of labelling, lowering your costs.
- **Work with public and private human labellers** – You can choose to use your team of labellers and route labelling requests directly to them. Alternatively, if you need to scale up, options are provided directly in the Amazon SageMaker Ground Truth console to work with labellers outside of your organisation. You can access a public workforce of over 500,000 labellers via integration with Amazon Mechanical Turk. Alternatively, if your data requires confidentiality or special skills, you can use professional labelling companies pre-screened by Amazon.
- **Automated labelling** – SageMaker Ground Truth will first select a random sample of data and send it to humans to be labelled. The results are then used to train a labelling model that attempts to label a new sample of raw data automatically. The labels are committed when the model can label the data with a confidence score that meets or exceeds a threshold you set. Where the confidence score falls below your threshold, the data is sent to human labellers. Some of the data labelled by humans is used to generate a new training dataset for the labelling model, and the model is automatically retrained to improve its accuracy. This process repeats with each sample of raw data to be labelled. The labelling model becomes more capable of automatically labelling raw data with each iteration, and less data is routed to humans.
- **Achieve accurate results quickly** – Amazon SageMaker Ground Truth helps

quickly build high-quality and accurate training datasets. Machine-generated labels provide consistent results with a confidence score for each label so that you can easily understand how certain the service is that the label is correct. Human-labelled results are automatically scored against criteria you provide to help ensure that more data is sent to high-quality labellers, and low-quality labellers are de-emphasised.

- **Scale** – If you want to scale up to a large number of labellers, and your data does not contain confidential or personally identifiable information, you have access to an on-demand, 24x7 workforce of over 500,000 independent contractors worldwide, powered by [Amazon Mechanical Turk](#). Mechanical Turk is a crowdsourcing marketplace that connects your labelling jobs with a distributed workforce that can perform these tasks virtually.
- **Third-party vendors** – Alternatively, you can use a third-party vendor that specialises in data labelling. These vendors have been screened by Amazon to provide high-quality labels and to follow security processes. Labelling services from these vendors are provided through [AWS Marketplace](#). All relevant details are provided, including pricing and customer reviews to help you select the best vendor for your needs.
- **Custom workflows** – In addition to the built-in workflows, SageMaker Ground Truth gives you the option to upload custom workflows. A custom workflow consists of an HTML interface and an accuracy improvement algorithm, both provided by you. The HTML interface provides the human labellers with all of the instructions and tools they need to complete the labelling task. The accuracy improvement algorithm is a function you write to tell SageMaker Ground Truth how it should assess the quality of labels that humans provide. The algorithm is used to find consensus on what is “right” when the same data is provided to multiple human labellers and to identify and de-emphasise labellers who tend to provide poor-quality data. You upload both the HTML interface and the accuracy improvement algorithm using the SageMaker Ground Truth console.
- **Simple setup** – Amazon SageMaker Ground Truth provides a managed experience where you can set up an entire data labelling job with just a few steps. To get started with Amazon SageMaker Ground Truth, you sign in to the AWS Management Console and navigate to the SageMaker console. From there, select Labelling jobs under Ground Truth. Here you can create a labelling job. First, as part of the labelling job creation flow, you provide a pointer to the Amazon S3 bucket that contains your dataset to be labelled. SageMaker Ground Truth offers templates for common labelling tasks. Using these templates, you only need to select a few options and provide minimal instructions on how to get the data labelled. Alternatively, you can create your own custom template. As the last step of creating a labelling job, you select one of the three human workforce options: 1) a public crowdsourced workforce, 2) a curated set of vendors who provide data labelling services, and 3) bring your own workers. You also have the option to enable automated data labelling.

- **Simple data management** – Amazon SageMaker Ground Truth manages the metadata, associated labels, and a taxonomy of your labels and datasets. You can easily use the AWS SDK through a SageMaker Notebook or the SageMaker Ground Truth console within the Amazon SageMaker console to query and manage your datasets and labels. Visit the Amazon SageMaker Ground Truth documentation for more information.
- **Secure** – By default, Amazon SageMaker Ground Truth encrypts your data at rest and in transit. In addition, access to your data can be controlled using IAM. SageMaker Ground Truth does not store or make copies of your data outside of your AWS environment, and your data remains in your control. Further, SageMaker Ground Truth supports compliance standards such as GDPR and provides comprehensive logging and auditing capabilities using Amazon CloudWatch and AWS CloudTrail.

58.6.1 Backup/Restore and Disaster Recovery

This requirement is not applicable to Amazon SageMaker Ground Truth. For additional information beyond what is described herein, please refer to <http://aws.amazon.com/documentation/>.

58.6.2 Pricing Overview

Please visit <https://aws.amazon.com/pricing/> for further details on pricing AWS services.

58.6.3 Service Constraints

Please see

https://docs.aws.amazon.com/general/latest/gr/aws_service_limits.html#limits_appsync for more information.

58.6.4 Technical Requirements

Please see <https://docs.aws.amazon.com/appsync/index.html> for comprehensive technical documentation regarding Amazon SageMaker Ground Truth.

59.0 Service Definition – Amazon Simple Email Service (Amazon SES)

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

59.1 Service Overview

Amazon SES is a cost-effective email service built on the reliable and scalable infrastructure that Amazon.com developed to serve its own customer base. With Amazon SES, you can send transactional email, marketing messages, or any other type of high-quality content to your customers. You can also use Amazon SES to receive messages and deliver them to an Amazon S3 bucket, call your custom code via an AWS Lambda function, or publish notifications to Amazon SNS. With Amazon SES, you

have no required minimum commitments—you pay as you go, and you only pay for what you use.

Top benefits include:

- **Inexpensive** – There are no up-front fees or fixed expenses with Amazon SES, and you benefit from the efficiencies of Amazon’s scale. You pay low charges for the number of emails sent, number of emails received, data transfer fees, and attachments.
- **Reliable** – Amazon SES runs within Amazon’s cloud network infrastructure and data centres. Multiple servers and data centres provide high availability and data durability.
- **Scalable** – Amazon SES is based on the scalable cloud-based email technology used by Amazon websites around the world to send billions of messages a year.
- **Highly Deliverable** – Amazon SES takes proactive steps to maximise the percentage of your emails that arrive in your recipients’ inboxes.
- **Highly Configurable** – When you use Amazon SES to receive your email, you have fine-grained control over what happens to your incoming messages. You can set up rules that apply to your entire domain or to specific email addresses.
- **Designed for Use with Other AWS Cloud Services** – Amazon SES integrates with [Amazon SNS](#), [Amazon Route 53](#), [Amazon S3](#), [Amazon WorkMail](#), [AWS Lambda](#), [AWS KMS](#), and [Amazon CloudWatch](#). Additionally, emails you send from [Amazon EC2](#) and [AWS Elastic Beanstalk](#) qualify for the free usage tier.

59.2 Backup/Restore and Disaster Recovery

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

59.3 Pricing Overview

Please visit <https://aws.amazon.com/pricing/> for further details on pricing AWS services.

59.4 Service Constraints

Please see <https://aws.amazon.com/ses/> for more information.

59.5 Technical Requirements

Amazon SES is an email sending and receiving service that provides an easy, cost-effective way for you to send email.

- **Developer Guide** – Provides a conceptual overview of Amazon SES and includes detailed instructions for using its various features. [HTML](#) | [PDF](#) | [Kindle](#)
- **API Reference** – Describes all the Amazon SES API operations in detail. Also

provides sample requests, responses, and errors for the supported web services protocols. [HTML](#) | [PDF](#)

60.0 Service Definition – Amazon Simple Notification Service (Amazon SNS)

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

60.1 Service Overview

Amazon SNS is a fast, flexible, fully managed push notification service that lets you send individual messages or to fan-out messages to large numbers of recipients. Amazon SNS makes it simple and cost effective to send push notifications to mobile device users, email recipients or even send messages to other distributed services.

With Amazon SNS, you can send notifications to Apple, Google, Fire OS, and Windows devices, as well as to Android devices in China with Baidu Cloud Push. You can use Amazon SNS to send SMS messages to mobile device users in the US or to email recipients worldwide.

Beyond these endpoints, Amazon SNS can also deliver messages to Amazon SQS, AWS Lambda functions, or to any HTTP endpoint.

Top features include:

Amazon SNS lets you push messages to mobile devices or distributed services via API or an easy-to-use management console. You can seamlessly scale from a handful of messages per day to millions of messages or higher.

With Amazon SNS you can publish a message once and deliver it one or more times. So you can choose to direct unique messages to individual Apple, Google, or Amazon devices or broadcast deliveries to many mobile devices with a single publish request.

Amazon SNS allows you to group multiple recipients using topics. A topic is an “access point” for allowing recipients to dynamically subscribe for identical copies of the same notification. One topic can support deliveries to multiple endpoint types—for example, you can group together iOS, Android and SMS recipients. When you publish once to a topic, Amazon SNS delivers appropriately formatted copies of your message to each subscriber.

Amazon SNS has no upfront costs and you can pay as you go.

60.2 Backup/Restore and Disaster Recovery

Amazon SNS runs within Amazon’s proven network infrastructure and datacentres, so topics will be available whenever applications need them. To prevent messages from being lost, all messages published to Amazon SNS are stored redundantly across multiple servers and data centres.

60.3 Pricing Overview

Please visit <https://aws.amazon.com/pricing/> for further details on pricing AWS services.

60.4 Service Constraints

Please see <https://aws.amazon.com/sns/> for more information.

60.5 Technical Requirements

Amazon SNS is a web service that enables applications, end-users, and devices to instantly send and receive notifications from the cloud.

- **Developer Guide** – Provides a conceptual overview of Amazon SNS and includes detailed development instructions for using the various features. [HTML](#) | [PDF](#) | [Kindle](#)
- **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. [HTML](#) | [PDF](#)
- **Quick Reference** – Briefly covers the essential commands for using Amazon SNS from the command line. [PDF](#)

61.0 Service Definition – Amazon Simple Queue Service (Amazon SQS)

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

61.1 Service Overview

Amazon SQS is a fast, reliable, scalable, fully managed message queuing service. Amazon SQS makes it simple and cost effective to decouple the components of a cloud application. You can use Amazon SQS to transmit any volume of data, at any level of throughput, without losing messages or requiring other services to be always available. With Amazon SQS, you can offload the administrative burden of operating and scaling a highly available messaging cluster, while paying a low price for only what you use.

Top benefits include:

- **Reliable** – Amazon SQS runs within Amazon's high-availability data centres, so queues will be available whenever applications need them. To prevent messages from being lost or becoming unavailable, all messages are stored redundantly across multiple servers and data centres.
- **Simple** – Developers can get started with Amazon SQS by using only three APIs: SendMessage, ReceiveMessage, and DeleteMessage. Additional APIs are available to provide advanced functionality.
- **Scalable** – Amazon SQS was designed to enable an unlimited number of

messaging services to read and write an unlimited number of messages at any time.

- **Secure** – Authentication mechanisms are provided to ensure that messages stored in Amazon SQS queues are secured against unauthorised access.
- **Inexpensive** – No up-front or fixed expenses. The only costs of sending messages through Amazon SQS are small per-request handling fees and data transfer fees.

61.2 Backup/Restore and Disaster Recovery

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

61.3 Pricing Overview

Please visit <https://aws.amazon.com/pricing/> for further details on pricing AWS services.

61.4 Service Constraints

Please see <https://aws.amazon.com/sqs/> for more information.

61.5 Technical Requirements

Amazon SQS is a messaging queue service that handles message or workflows between other components in a system.

- **Getting Started Guide** – Introduces you to Amazon SQS, helps you set up an account, and walks you through a simple example to use Amazon SQS for the first time. Also provides tips and links to advanced product features and resources. [HTML](#) | [PDF](#) | [Kindle](#)
- **Developer Guide** – Provides a conceptual overview of Amazon SQS and includes detailed development instructions for using the various features. [HTML](#) | [PDF](#) | [Kindle](#)
- **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. [HTML](#) | [PDF](#)
- **Amazon SQS section of AWS CLI Reference** – Describes the AWS CLI commands that you can use to automate queues. [HTML](#)

62.0 Service Definition – Amazon Simple Storage Service (Amazon S3)

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

62.1 Service Overview

Amazon S3 provides developers and IT teams with secure, durable, and highly scalable object storage. Amazon S3 is easy to use, with a simple interface to store and retrieve any amount of data from anywhere on the web. With Amazon S3, you pay only for the storage you actually use. There is no minimum fee and no setup cost.

Amazon S3 can be used alone or together with other AWS Cloud services such as Amazon EC2, Amazon EBS, and Amazon Glacier, as well as third-party storage repositories and gateways. Amazon S3 provides cost-effective object storage for a wide variety of use cases, including cloud applications, content distribution, backup and archiving, disaster recovery, and big data analytics.

Top benefits include:

- **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.
- **Scalable** – With Amazon S3, you can store as much data as you want and access it when you need it.
- **Send Event Notifications** – Amazon S3 can send event notifications when objects are uploaded to Amazon S3.
- **High Performance** – Amazon S3 supports multi-part uploads to help maximise network throughput and resiliency and lets you choose the AWS Region in which to store your data, minimising network latency.
- **Integrated** – Amazon S3 is integrated with other AWS Cloud services to simplify uploading and downloading data from Amazon S3 and to make it easier to build solutions that use a range of AWS Cloud services.
- **Easy to Use** – Amazon S3 is easy to use with a web-based management console and mobile app and full Representational State Transfer (REST) APIs and SDKs for easy integration with third-party technologies.
- **Access management** - To protect your data in Amazon S3, by default, users only have access to the S3 resources they create. You can grant access to other users by using one or a combination of the following access management features: AWS Identity and Access Management (IAM) to create users and manage their respective access; Access Control Lists (ACLs) to make individual objects accessible to authorized users; bucket policies to configure permissions for all objects within a single S3 bucket; S3 Access Points to simplify managing

data access to shared data sets by creating access points with names and permissions specific to each application or sets of applications; and Query String Authentication to grant time-limited access to others with temporary URLs. Amazon S3 also supports Audit Logs that list the requests made against your S3 resources for complete visibility into who is accessing what data.

- **Write Once, Read Many** - You can also enforce write-once-read-many (WORM) policies with S3 Object Lock. This S3 management feature blocks object version deletion during a customer-defined retention period so that you can enforce retention policies as an added layer of data protection or to meet compliance obligations.
- **Transferring large amounts of data** - AWS has a suite of data migration services that make transferring data into the AWS Cloud simple, fast, and secure. S3 Transfer Acceleration is designed to maximize transfer speeds to S3 buckets over long distances. For very large data transfers, consider using AWS Snowball, AWS Snowball Edge, and AWS Snowmobile to move petabytes to exabytes of data to the AWS Cloud for as little as one-fifth the cost of high-speed Internet.
- **Query in place** - Amazon S3 has a built-in feature and complimentary services that query data without needing to copy and load it into a separate analytics platform or data warehouse. This means you can run big data analytics directly on your data stored in Amazon S3. S3 Select is an S3 feature designed to increase query performance by up to 400%, and reduce querying costs as much as 80%. It works by retrieving a subset of an object's data (using simple SQL expressions) instead of the entire object, which can be up to 5 terabytes in size.

62.2 Backup/Restore and Disaster Recovery

Amazon S3 offers a highly durable, scalable, and secure solution for backing up and archiving your critical data. You can use Amazon S3's versioning capability to provide even further protection for your stored data. You can also define rules to archive sets of Amazon S3 objects to Amazon Glacier's extremely low-cost storage service based on object lifetimes. As your data ages, these rules enable you to ensure that it is automatically stored on the storage option that is most cost effective for your needs.

62.3 Pricing Overview

Please visit <https://aws.amazon.com/pricing/> for further details on pricing AWS services.

62.4 Service Constraints

Please see <http://aws.amazon.com/documentation/s3/> for more information.

62.5 Technical Requirements

Please refer to <http://aws.amazon.com/documentation/s3/> and the following links for comprehensive technical documentation regarding Amazon S3.

- **Getting Started Guide** – Introduces you to Amazon S3, helps you set up an account, and walks you through a simple example to help you use Amazon S3 for the first time. Also provides tips and links to advanced product features and resources. Available in [HTML](#), [PDF](#), and [Kindle](#) formats.
- **API Reference** – Describes all the Amazon S3 API operations in detail. Also provides sample requests, responses, and errors for the supported web services protocols. Available in [HTML](#) and [PDF](#) formats.
- **Quick Reference Card** – Briefly covers the essential commands for using Amazon S3 from the CLI. Available in [PDF](#) format.
- **Developer Guide** – Provides a conceptual overview of Amazon S3 and includes detailed instructions for using the various features. Available in [HTML](#), [PDF](#), and [Kindle](#) formats.
- **Console User Guide** – Provides information to help you use Amazon S3 with the AWS Management Console. Available in [HTML](#), [PDF](#), and [Kindle](#) formats.

63.0 Service Definition – Amazon Simple Workflow Service (Amazon SWF)

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

63.1 Service Overview

Amazon SWF helps developers build, run, and scale background jobs that have parallel or sequential steps. You can think of Amazon SWF as a fully managed state tracker and task coordinator in the cloud. If your application's steps take more than 500 milliseconds to complete, you need to track the state of processing, and you need to recover or retry if a task fails. Amazon SWF can help you.

Top benefits include:

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 centres, so the state tracking and task processing engine is available whenever applications need them. Amazon SWF redundantly stores the tasks, reliably dispatches them to application components, tracks their progress, and keeps their latest state.
- **Simple** – Amazon SWF replaces the complexity of custom-coded workflow solutions and process automation software with a fully managed cloud workflow

web service. This eliminates the need for developers to manage the infrastructure plumbing of process automation so they can focus their energy on the unique functionality of their application.

- **Scalable** – Amazon SWF seamlessly scales with your application’s usage. No manual administration of the workflow service is required as you add more cloud workflows to your application or increase the complexity of your workflows.
- **Flexible** – Amazon SWF lets you write your application components and coordination logic in any programming language and run them in the cloud or on premises.

63.2 Backup/Restore and Disaster Recovery

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

63.3 Pricing Overview

Please visit <https://aws.amazon.com/pricing/> for further details on pricing AWS services.

63.4 Service Constraints

Please see <https://aws.amazon.com/swf/> for more information.

63.5 Technical Requirements

Amazon SWF makes it easy to build applications that coordinate work across distributed components. In Amazon SWF, a task represents a logical unit of work that is performed by a component of your application. Coordinating tasks across the application involves managing intertask dependencies, scheduling, and concurrency in accordance with the logical flow of the application. Amazon SWF gives you full control over implementing tasks and coordinating them without worrying about underlying complexities such as tracking their progress and maintaining their state.

- **Developer Guide** – Provides a conceptual overview of Amazon SWF and includes detailed development instructions for using the various features. [HTML](#) | [PDF](#) | [Kindle](#)
- **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. [HTML](#) | [PDF](#)

64.0 Service Definition – Amazon Sumerian

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

64.1 Service Overview

Amazon Sumerian lets you create and run virtual reality (VR), augmented reality (AR), and 3D applications quickly and easily without requiring any specialised programming or 3D graphics expertise. With Sumerian, you can build highly immersive and interactive scenes that run on popular hardware such as Oculus Go, Oculus Rift, HTC Vive, HTC Vive Pro, Google Daydream, and Lenovo Mirage as well as Android and iOS mobile devices. For example, you can build a virtual classroom that lets you train new employees around the world, or you can build a virtual environment that enables people to tour a building remotely. Sumerian makes it easy to create all the building blocks needed to build highly immersive and interactive 3D experiences, including adding objects (e.g., characters, furniture, and landscape) and designing, animating, and scripting environments. Sumerian does not require specialised expertise, and you can design scenes directly from your browser.

The top features include:

- **Sumerian editor** – Sumerian has a simple-to-use editor based on WebGL and WebVR so you can easily author scenes from a browser. You can import 3D objects, select the surroundings for your scene, add characters, and script the logic that controls your scene. Sumerian's text editor lets you script the logic behind your scene so you can decide how objects behave or how your scene is narrated.
- **Sumerian Hosts** – Sumerian Hosts makes it easy to create and add animated 3D characters called Hosts. Hosts can guide your users through a scene by narrating scripts or answering questions. You can select from three Hosts and can customise the Host's appearance, clothing, voice, and language. Using Sumerian's integration with Amazon Lex and Amazon Polly, you can enable engaging spoken interactions between Hosts and your users.
- **Asset management** – You can import and place 3D objects in your scenes, such as furniture, natural objects, buildings, surroundings (e.g., stadium, classroom), and more. Sumerian also comes with templates with pre-populated scenes. Sumerian supports importing FBX and OBJ. You can choose from Sumerian's library of prebuilt objects, download and import objects from third-party 3D asset repositories, such as Sketchfab or TurboSquid, or create and import your own objects.
- **Script the logic in your scenes** – Sumerian lets you build dynamic scenes that respond to user interactions. Using Sumerian's visual workflow or JavaScript editors, you can script the logic that controls how objects behave or respond to actions. For example, you can start by building a simple behaviour that makes a flag billow or a Host perform a gesture. You can also build a trigger that turns on a light bulb whenever a user toggles a switch. You could even build a trigger that calls a web API to retrieve external data for your scene, such as social media photos that are shown within the scene.
- **Multi-platform compatibility** – You can build VR/AR applications for platforms

such as Oculus Go, Oculus Rift, HTC Vive, HTC Vive Pro, Google Daydream, Lenovo Mirage, and AR-compatible iOS and Android mobile devices. Your Sumerian scenes will run in browsers that support WebGL or WebVR. Sumerian hosts your completed scenes in the cloud and serves them from a publicly accessible endpoint, enabling your users to access them from anywhere. To learn more, please visit the [tutorials](#) on creating and publishing scenes on various devices.

Top benefits include:

- **No VR/AR expertise needed** – Amazon Sumerian lets anyone create VR/AR applications. You don't need specialised programming or 3D graphics expertise to get started.
- **Lifelike characters** – Sumerian allows you to easily create 3D characters, called Hosts, which narrate and guide users through your scenes. You can create lifelike, animated Hosts that speak a variety of languages and can converse with your users.
- **Design immersive environments** – Sumerian makes it easy to design scenes that immerse your users in 3D, lifelike surroundings that are populated with 3D objects and animated characters. You can import 3D objects and select from preloaded scene templates. You can also script the logic that controls how the objects in your scenes behave.
- **Run on multiple platforms** – Sumerian is platform agnostic, so that once you're done building your scene, you can easily run it on popular hardware, including Oculus Go, Oculus Rift, HTC Vive, HTC Vive Pro, Google Daydream, Lenovo Mirage, and mobile devices. Sumerian is based on the WebGL JavaScript API for graphic rendering and WebVR JavaScript API for virtual reality devices. It also supports the ARKit framework for augmented reality on iOS devices as well as ARCore for Android devices.

64.2 Backup/Restore and Disaster Recovery

This requirement is not applicable to Amazon Sumerian, as it collates other services rather than being a primary data storage service.

64.3 Pricing Overview

Please visit <https://aws.amazon.com/pricing/> for further details on pricing AWS services.

64.4 Service Constraints

Please see <https://aws.amazon.com/sumerian/> for more information.

64.5 Technical Requirements

Please refer to <https://aws.amazon.com/sumerian/> for comprehensive technical documentation regarding Amazon Sumerian.

65.0 Service Definition – Amazon Textract

Amazon Textract is a service that automatically extracts text and data from scanned documents. Amazon Textract goes beyond simple optical character recognition (OCR) to also identify the contents of fields in forms and information stored in tables.

65.1 Service Overview

Top features include:

- Uses Optical Character Recognition (OCR) technology to automatically detect printed text and numbers in a scan or rendering of a document, such as a legal document or a scan of a book.
- Detect key-value pairs (set of linked data items) in document images automatically so that you can retain the inherent context of the document without any manual intervention and import the extracted data into a database or provide it as a variable into an application.
- Preserves the composition of structured data stored in tables during extraction, enabling the ability to automatically load the extracted data into a database using a pre-defined schema.
- When information is extracted from documents, Amazon Textract returns a confidence score for everything it identifies so that you can make informed decisions about how you want to use the results.
- Amazon Textract is directly integrated with Amazon Augmented AI (Amazon A2I) so you can easily implement human review of text extracted from documents.

Top benefits include:

- Amazon Textract makes it easy to quickly and accurately extract data from documents, forms, and tables. Amazon Textract automatically detects a document's layout and the key elements on the page, understands the data relationships in any embedded forms or tables, and extracts everything with its context intact. This means you can instantly use the extracted data in an application or store it in a database without a lot of complicated code in between.
- Extracted data is returned with bounding box coordinates that encompasses each piece of identified data, enabling the ability to audit where a word or number came from in the source document or to guide the user in document search systems.
- Pre-trained machine learning models eliminate the need to write code for data extraction, because they have already been trained on tens of millions of documents from virtually every industry, including contracts, tax documents, sales orders, enrolment forms, benefit applications, insurance claims, policy documents and many more. You no longer need to maintain code for every document or form you might receive or worry about how page layouts change over time.

- Amazon Textract provides OCR and structured data extraction (forms and tables) at very low cost, and you only pay for what you use. There are no upfront commitments or long-term contracts. You can easily process millions of documents using Amazon Textract's text extraction APIs.

65.2 Backup/Restore and Disaster Recovery

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

The AWS global infrastructure is built around AWS Regions and Availability Zones. AWS Regions provide multiple physically separated and isolated 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 centre infrastructures.

For more information about AWS Regions and Availability Zones, see AWS Global Infrastructure (<http://aws.amazon.com/about-aws/global-infrastructure/>).

65.3 Pricing Overview

Please see the AWS UK G Cloud 12 Pricing Document affiliated with this service in the Digital Marketplace.

65.4 Service Constraints

Please see <https://docs.aws.amazon.com/textract/latest/dg/limits.html> for more information regarding the constraints in place when using Amazon Textract.

65.5 Technical Requirements

Please refer to <https://docs.aws.amazon.com/textract/index.html> and the following links for comprehensive technical documentation regarding Amazon Textract:

- <https://docs.aws.amazon.com/textract/latest/dg/what-is.html>

66.0 Service Definition – Amazon Transcribe

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

66.1 Service Overview

Amazon Transcribe is an automatic speech recognition (ASR) service that makes it easy for developers to add speech-to-text capability to their applications. Using the

Amazon Transcribe API, you can analyse audio files stored in Amazon S3 and have the service return a text file of the transcribed speech.

Amazon Transcribe can be used for lots of common applications, including the transcription of customer service calls and generating subtitles on audio and video content. The service can transcribe audio files stored in common formats, like WAV and MP3, with time stamps for every word so that you can easily locate the audio in the original source by searching for the text. Amazon Transcribe is continually learning and improving to keep pace with the evolution of language.

Top benefits include:

- **Simple-to-Use API** - The Amazon Transcribe API makes it easy to convert speech to text. No complicated programming is required. Just call the API with a few lines of code, and Amazon Transcribe will return the text from your audio file stored in Amazon S3.
- **Support for a Wide Range of Use Cases** - Amazon Transcribe is designed to provide accurate and automated transcripts for a wide range of audio quality. You can generate subtitles for any video or audio files, and even transcribe low quality telephony recordings such as customer service calls.
- **Easy-to-Read Transcriptions** - Most speech recognition systems output a string of text without punctuation. Amazon Transcribe uses deep learning to add punctuation and formatting automatically, so that the output is more intelligible and can be used without any further editing.
- **Custom Vocabulary** - Amazon Transcribe gives you the ability to expand and customize the speech recognition vocabulary. You can add new words to the base vocabulary and generate highly-accurate transcriptions specific to your use case, such as product names, domain-specific terminology, or names of individuals.
- **Timestamp Generation** - Amazon Transcribe returns a timestamp for each word, so that you can easily locate the audio in the original recording by searching for the text.
- **Recognize Multiple Speakers** - Amazon Transcribe is able to recognize when the speaker changes and attribute the transcribed text appropriately. This can significantly reduce the amount of work needed to transcribe audio with multiple speakers like telephone calls, meetings, and television shows.

66.2 Backup/Restore and Disaster Recovery

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

66.3 Pricing Overview

Please visit <https://aws.amazon.com/pricing/> for further details on pricing AWS services.

66.4 Service Constraints

Please see <https://aws.amazon.com/transcribe/> for more information. Prior to using these Services, Buyer agrees to contact Supplier before using these Services in connection with Buyer Personal Data.

66.5 Technical Requirements

For additional resources and information, please see <https://aws.amazon.com/transcribe/resources/>.

- **Developer Guide** – Provides a conceptual overview of Amazon Transcribe, includes detailed instructions for using the various features, and provides a complete API reference for developers. [HTML](#) | [PDF](#)

67.0 Service Definition – Amazon Transcribe Medical

Amazon Transcribe Medical is a machine learning service that makes it easy to quickly create accurate transcriptions from medical consultations between patients and physicians. With Transcribe Medical, the medical and pharmacological terms used in physician dictated notes, practitioner/patient consultations, and tele-medicine are automatically converted from speech to text for use in clinical documentation applications.

Amazon Transcribe Medical uses machine learning to provide highly accurate automatic speech recognition (ASR) for the medical industry. You can use Transcribe Medical to quickly and efficiently capture physician-patient conversations in text for later analysis using natural language processing or for entry into electronic health record (EHR) systems, because the service is trained to understand the terminology and style of clinical language. With Transcribe Medical, physicians are able to better focus on their patient and provide a more attentive experience instead of interrupting the conversation for note taking.

67.1 Service Overview

Top features include:

- **Highly accurate.** Transcribe Medical provides accurate real-time speech-to-text that's built for medical language using state of the art machine learning models. This means that a statement like “patient suffered a plantar fibroma” will be captured accurately. This cuts down on the time that physicians have to take and edit notes manually, reducing physician burnout that’s typically caused by tedious hours of medical note taking.
- **Natural speech.** Transcribe Medical supports conversations between multiple people or dictation from a single-speaker. It also provides auto-punctuation so that physicians can speak naturally, instead of explicitly speaking out punctuation marks. The transcription will also automatically include normalized numbers (e.g.

50 beats per minute instead of “fifty”) and capitalization. Because Transcribe Medical captures natural conversations, the physician doesn't have to pause for notes, providing a more attentive and improved patient experience.

- **Lower medical documentation costs.** Because Transcribe Medical reduces reliance on expensive human scribes and doesn't require licensing fees like other traditional transcription services, it reduces the overall cost of medical documentation. Clinical documentation applications using Transcribe Medical can be deployed at scale across thousands of medical centers, providing affordable, consistent, and accurate note taking for clinical staff and facilities no matter where they're located.

Top benefits include:

- **Dictation.** Using Transcribe Medical, physicians can dictate their notes using a mobile device after the patient interaction. This gives physicians the ability to dictate medical notes quickly, instead of manual note-taking or waiting for third party services, reducing time and cost and improving the patient experience.
- **Conversational transcription.** Healthcare and life science customers can build transcription applications that capture physician-patient conversations in real time, without disrupting the interaction. Transcripts can be used to gather insights like medication, dosage, strength, and frequency, before final entry into electronic health record (EHR) systems. You can use Comprehend Medical, a natural language processing service for medical text, to extract those insights.
- **Drug safety monitoring:** Pharmaceutical companies can use Transcribe Medical to transcribe patient calls with physicians or physician calls with pharmaceutical companies, accurately capturing names of medicines and key terms that describe health side effects. With this information in text format, pharmaceutical companies can more efficiently analyze and detect drug-related safety issues, for example an unexpected side effect for a new medication.

67.2 Backup/Restore and Disaster Recovery

This requirement is not applicable for Amazon Transcribe Medical. For additional information beyond what is described herein, please refer to <http://aws.amazon.com/documentation/>.

The AWS global infrastructure is built around AWS Regions and Availability Zones. AWS Regions provide multiple physically separated and isolated 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 centre infrastructures.

For more information about AWS Regions and Availability Zones, see AWS Global Infrastructure (<http://aws.amazon.com/about-aws/global-infrastructure/>).

67.3 Pricing Overview

Please see the AWS UK G Cloud 12 Pricing Document affiliated with this service in the Digital Marketplace.

67.4 Service Constraints

Please see <https://docs.aws.amazon.com/transcribe/latest/dg/transcribe-medical.html> for more information regarding the constraints in place when using Amazon Transcribe Medical.

67.5 Technical Requirements

Please refer to <https://docs.aws.amazon.com/transcribe/latest/dg/transcribe-medical.html> for comprehensive technical documentation regarding Amazon Transcribe Medical

68.0 Service Definition – Amazon Translate

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

68.1 Service Overview

Amazon Translate is a neural machine translation service that delivers fast, high-quality, and affordable language translation. Neural machine translation is a form of language translation automation that uses deep learning models to deliver more accurate and more natural sounding translation than traditional statistical and rule-based translation algorithms. Amazon Translate allows you to localize content - such as websites and applications - for international users, and to easily translate large volumes of text efficiently.

Top benefits include:

- **Highly Accurate & Continuously Learning** - Amazon Translate is a neural machine translation service. The translation engines are always learning from new and expanded datasets to produce more accurate and fluent translations for a wide range of use cases.
- **Easy to Integrate into Your Applications** - Amazon Translate removes the complexity of building real-time and batch translation capabilities into your applications with a simple API call. This makes it easy to localize an application or web site, or process multilingual data within your existing workflows.
- **Scalable** - Whether it's just a few words or large volumes of text, Amazon Translate scales easily as your translation needs grow. The service provides consistently fast and reliable translations, regardless of the volume of translation requests you make.

68.2 Backup/Restore and Disaster Recovery

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

68.3 Pricing Overview

Please visit <https://aws.amazon.com/pricing/> for further details on pricing AWS services.

68.4 Service Constraints

Please see <https://aws.amazon.com/translate/resources/> for more information. Prior to using these Services, Buyer agrees to contact Supplier before using these Services in connection with Buyer Personal Data.

68.5 Technical Requirements

Amazon Translate uses advanced machine learning technologies to provide high-quality translation on demand. Use it to translate unstructured text documents or to build applications that work in multiple languages. For additional resources and information, please see <https://aws.amazon.com/translate/resources/>.

- **Developer Guide** – Provides a conceptual overview of Amazon Translate, includes detailed instructions for using the various features, and provides a complete API reference for developers. [HTML](#) | [PDF](#)

69.0 Service Definition – Amazon Virtual Private Cloud (Amazon VPC)

The following subsections provide service definition information.

69.1 Service Overview

Amazon VPC lets you provision a logically isolated section of the 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 range, creation of subnets, and configuration of route tables and network gateways.

You can easily customise the network configuration for your Amazon VPC. For example, you can create a public-facing subnet for your web servers that has access to the Internet and place your backend systems such as databases or application servers in a private-facing subnet with no Internet access. You can leverage multiple layers of security, including security groups and network access control lists, to help control access to Amazon EC2 instances in each subnet.

Additionally, you can create a hardware Virtual Private Network (VPN) connection between your corporate data centre and your VPC and leverage the AWS Cloud as an extension of your corporate data centre.

Top benefits include:

- **Multiple Connectivity Options** – A variety of connectivity options exist for your Amazon VPC. You can connect your VPC to the Internet, to your data centre, or to other VPCs, based on the AWS resources that you want to expose publicly and those that you want to keep private.
- **Secure** – Amazon VPC provides advanced security features such as security groups and network access control lists to enable inbound and outbound filtering
- **Simple** – You can create a VPC quickly and easily using the AWS Management Console.
- **All the Scalability and Reliability of AWS** – Amazon VPC provides all the same benefits as the rest of the AWS platform. You can instantly scale your resources up or down, select Amazon EC2 instance types and sizes that are right for your applications, and pay only for the resources you use—all within Amazon's proven infrastructure.

69.2 Backup/Restore and Disaster Recovery

This requirement is not applicable for Amazon VPC. 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 12 Pricing Document affiliated with this service in the Digital Marketplace.

69.4 Service Constraints

Please see <http://aws.amazon.com/documentation/vpc/> for more information.

69.5 Technical Requirements

Please refer to <http://aws.amazon.com/documentation/vpc/> and the following links for comprehensive technical documentation regarding Amazon VPC.

- **Getting Started Guide** – Provides instructions to create a VPC and launch an instance into the Amazon VPC. Available in [HTML](#), [PDF](#), and [Kindle](#) formats.
- **CLI Reference** – Documents the Amazon VPC CLI. Available in [HTML](#) and [PDF](#) formats.
- **API Reference** – Documents the Amazon VPC Query API. Available in [HTML](#) and [PDF](#) formats.
- **User Guide** – Describes key concepts for Amazon VPC and provides

instructions for using the features of Amazon VPC. Available in [HTML](#), [PDF](#), and [Kindle](#) formats.

- **Network Administrator Guide** – Describes customer gateways and helps network administrators configure them. Available in [HTML](#), [PDF](#), and [Kindle](#) formats.
- **Quick Reference Card** – Briefly covers the essential commands for using Amazon VPC from the CLI. Available in [PDF](#) format.

70.0 Service Definition – Amazon WorkDocs

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

70.1 Service Overview

Amazon WorkDocs is a fully managed, secure enterprise storage and sharing service with strong administrative controls and feedback capabilities that improve user productivity. Users can comment on files, send them to others for feedback, and upload new versions without having to resort to emailing multiple versions of their files as attachments. Users can take advantage of these capabilities wherever they are, using the device of their choice, including PCs, Macs, tablets and phones. Amazon WorkDocs offers IT administrators the option of integrating with existing corporate directories, flexible sharing policies and control of the location where data is stored. Customers can get started using Amazon WorkDocs with a 30-day free trial providing 1 TB of storage per user for up to 50 users.

Amazon WorkDocs offers an Administrative SDK. The Administrative SDK allows you to integrate your applications with Amazon WorkDocs by performing content and permissions updates, and managing users, programmatically.

Top benefits include:

- **Simple Document Feedback** – Users can comment on files, send them to others for feedback, and upload new versions without having to resort to emailing multiple versions of their files as attachments. Teammates can leave detailed comments on sections of a file as well as overall comments. Amazon WorkDocs notifies contributors about review activities and deadlines via email and offers versioning capabilities.
- **Central Hub** – Amazon WorkDocs provides users with a central location for both the documents and files they are reviewing as well as those they own and are soliciting feedback on. With all these files in one location, reviewers have access to all of the related feedback in a single web view, making reading or contributing comments as simple as a few clicks.
- **Access and Sync from Any Device** – Users can access files stored in Amazon WorkDocs and view and leave feedback on files anywhere, anytime, from the device of their choice. Users can leverage the Amazon WorkDocs apps for iPad,

Kindle Fire, and Android tablets. The Amazon WorkDocs Sync client lets users have files saved on their computer automatically uploaded to Amazon WorkDocs over an encrypted connection and synced across their other devices.

- **Secure** – Amazon WorkDocs offers flexible security settings and access controls to manage data storage and sharing. You can use policies to control users' sharing behaviour and choose the AWS Region where your users' data is stored. All data is encrypted in transit and at rest, and users don't have to send documents as email attachments.
- **Integrate Your Corporate Directory** – Amazon WorkDocs can integrate with your existing Active Directory. This means that your users can easily access Amazon WorkDocs using their existing Active Directory credentials and allows you to control which users in your organisation are permitted to access the service.

70.2 Backup/Restore and Disaster Recovery

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

70.3 Pricing Overview

Please visit <https://aws.amazon.com/pricing/> for further details on pricing AWS services.

70.4 Service Constraints

Please see <https://aws.amazon.com/workdocs/> for more information.

70.5 Technical Requirements

Amazon WorkDocs is a fully managed, secure enterprise storage and sharing service with strong administrative controls and feedback capabilities that improve user productivity.

- **Administration Guide** – Helps you use Amazon WorkDocs to perform several administrative tasks, such as creating a new directory in the cloud, connecting to your on-premises directory, or setting user privileges and defaults. [HTML](#) | [PDF](#)
- **User Guide** – Helps you use the Amazon WorkDocs collaboration and utility applications. File collaboration applications are available for desktop web browsers, as well as several tablets. [HTML](#) | [PDF](#)

71.0 Service Definition – Amazon WorkLink

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

71.1 Service Overview

Amazon WorkLink is a fully managed service that lets you provide your employees with secure, easy access to your internal corporate websites and web apps using their mobile phones. Traditional solutions such as VPNs and device management software are inconvenient to use on the go and often require the use of custom browsers that have a poor user experience. As a result, employees often forgo using them altogether.

With Amazon WorkLink, employees can access internal web content as easily as they access any public website, without the hassle of connecting to their corporate network. When a user accesses an internal website, the page is first rendered in a browser running in a secure container in AWS. Amazon WorkLink then sends the contents of that page to employee phones as vector graphics while preserving the functionality and interactivity of the page. This approach is more secure than traditional solutions because internal content is never stored or cached by the browser on employee phones, and employee devices never connect directly to your corporate network. With Amazon WorkLink, there are no minimum fees or long-term commitments. You pay only for users that connect to the service each month, and there is no additional charge for bandwidth consumption.

The top features include:

- **Easy setup and administration** – Amazon WorkLink can be easily set up from the AWS Management Console. To get started, link your existing identity provider to Amazon WorkLink, and use that to configure access permissions for your employees. Next, add your web domains that will be accessed using Amazon WorkLink. To enable access to these added web domains, use your existing on-premises VPN hardware to create a point-to-point connection with your Amazon VPC, or simply use AWS Direct Connect if you have it set up already. Once you have completed these steps, you can use the provided email template to invite employees to download the [Amazon WorkLink app](#) from their device app store, log in with their corporate credentials, and start accessing internal websites using Safari.
- **Secure content isolation** – Amazon WorkLink performs on-device DNS resolution to identify internal website and web app requests and then loads this content in a secure container running in AWS. Amazon WorkLink processes all HTML, JavaScript, and CSS in AWS and transforms it into vector graphics. It then delivers the content to employee phones as vector graphics while preserving native interactions. As a result, internal web pages are never directly rendered on these devices, and content isn't stored or downloaded to local browser caches. Amazon WorkLink also isolates browsing sessions by providing a dedicated pool of EC2 instances to each customer and a dedicated container to each active user.
- **Managed service** – Amazon WorkLink owns and manages the deployment, provisioning, and scaling of the resources you need and automatically keeps these resources up to date. Amazon WorkLink-managed resources dynamically connect with your Amazon VPC to access the internal websites you specify. You

can leverage AWS Direct Connect installations to route traffic from AWS to company websites and deprecate the use of VPN gateway hardware and software on premises. Alternatively, you can reuse existing VPN installations to set up a site-to-site VPN tunnel between AWS and the on-premises network. This allows you to reduce your on-premises management overhead, since you no longer need to maintain complex client-to-site VPN gateways that need to be secured to allow direct access from employee phones.

- **Granular access control** – Amazon WorkLink lets you specify which of your internal websites and web apps should be available to your employees, contractors, and partners. You can whitelist the sites that you want to make accessible externally in the Amazon WorkLink console and set permissions for your users through your existing identity provider, including SAML 2.0 and Active Directory. This lets you control level of access for users and makes it easier to protect your information.

Top benefits include:

- **Seamless mobile access** – Amazon WorkLink helps your employees instantly access internal websites without requiring them to log in to a VPN client or open an app.
- **Enhanced security** – Amazon WorkLink renders the content of your websites in a browser running in a secure container in AWS before sending the contents of that page to employee phones as vector graphics.
- **Easy for IT** – Amazon WorkLink enables secure, one-click access to internal web content, regardless of where it is hosted. Amazon WorkLink does not require you to migrate your content to AWS and works with your existing infrastructure.

71.2 Backup/Restore and Disaster Recovery

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

71.3 Pricing Overview

Please visit <https://aws.amazon.com/pricing/> for further details on pricing AWS services.

71.4 Service Constraints

Please see <https://aws.amazon.com/worklink/> for more information.

71.5 Technical Requirements

Amazon WorkLink is a fully managed service that lets you provide your employees with secure, easy access to your internal corporate websites and web apps using their mobile phones.

- **Administration Guide** – Helps you use Amazon WorkDocs to perform several

administrative tasks, such as creating a new directory in the cloud, connecting to your on-premises directory, or setting user privileges and defaults. [HTML](#) | [PDF](#)

- **User Guide** – Helps you use the Amazon WorkDocs collaboration and utility applications. File collaboration applications are available for desktop web browsers, as well as several tablets. [HTML](#) | [PDF](#)

72.0 Service Definition – Amazon WorkMail

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

72.1 Service Overview

Amazon WorkMail is a secure, managed business email and calendar service with support for existing desktop and mobile email clients. Amazon WorkMail gives users the ability to seamlessly access their email, contacts, and calendars using Microsoft Outlook, their web browser, or their native iOS and Android email applications. You can integrate Amazon WorkMail with your existing corporate directory and control both the keys that encrypt your data and the location in which your data is stored. You can get started for free with a 30-day trial of Amazon WorkMail for up to 25 users.

Top benefits include:

- **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 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 synchronise their mailboxes with iOS, Android, Amazon Fire, and Windows Phone devices (using the Microsoft Exchange ActiveSync protocol). 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, which includes 50GB of storage per user.

72.2 Backup/Restore and Disaster Recovery

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

72.3 Pricing Overview

Please visit <https://aws.amazon.com/pricing/> for further details on pricing AWS services.

72.4 Service Constraints

Please see <https://aws.amazon.com/workmail/> for more information.

72.5 Technical Requirements

Amazon WorkMail is a managed email and calendaring service that offers strong security controls and support for existing desktop and mobile clients.

- **User Guide** – Helps you configure your desktop email client or mobile device for Amazon WorkMail and use the Amazon WorkMail web application. [HTML](#) | [PDF](#)
- **Administrator Guide** – Helps you perform administrative tasks for Amazon WorkMail, such as setting up Amazon WorkMail for your organisation, adding a domain, managing users, groups, and mobile devices. [HTML](#) | [PDF](#)

73.0 Service Definition – Amazon WorkSpaces

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

73.1 Service Overview

Amazon WorkSpaces is a fully managed, secure desktop computing service that runs on the AWS Cloud. Amazon WorkSpaces allows you to easily provision cloud-based virtual desktops and provide your users access to the documents, applications, and resources they need from any supported device, including Windows and Mac computers, Chromebooks, iPads, Kindle Fire tablets, and Android tablets. With just a few clicks in the AWS Management Console, you can deploy high-quality cloud desktops for any number of users at a cost that is competitive with traditional desktops and half the cost of most VDI solutions.

Top benefits include:

- **Simple to Use** – Amazon WorkSpaces makes it easy to manage your desktop computing infrastructure by eliminating the need for up-front investments and avoiding the complexity of maintaining, patching, and managing a large physical desktop environment or a complex VDI solution.
- **Keep Your Data Secure** – Amazon WorkSpaces provides each user with access to persistent storage in the AWS Cloud. When users access their desktops using Amazon WorkSpaces, a remote display protocol is used to compress, encrypt, and encode data so that only images are transmitted. Amazon WorkSpaces also integrates with the [AWS KMS](#), providing the ability to encrypt WorkSpaces storage volumes.
- **Deploy and Manage Applications** – Amazon WorkSpaces Application Manager (Amazon WAM) offers a fast, flexible, and secure way for you to package, deploy, and update your organisation's desktop applications for Amazon WorkSpaces.
- **Choose the Hardware and Software You Need** – Amazon WorkSpaces offers a choice of bundles providing different amounts of CPU, memory, and storage so you can match your Amazon WorkSpaces to your requirements. Amazon WorkSpaces offers preinstalled OS and applications (including Microsoft Office), or you can bring your own Windows desktop licenses or other licensed software.
- **Support Multiple Devices** – Because the desktops are in the cloud, users can access their Amazon WorkSpaces from any supported device including Windows and Mac computers, Chromebooks, iPads, Kindle Fire tablets, and Android tablets.
- **Integrate Your Corporate Directory** – Amazon WorkSpaces securely integrates with your corporate AD so that your users can continue using their existing enterprise credentials to seamlessly access company resources. This also makes it easy to manage your WorkSpaces using familiar systems management tools.

73.2 Backup/Restore and Disaster Recovery

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

73.3 Pricing Overview

Please visit <https://aws.amazon.com/pricing/> for further details on pricing AWS services.

73.4 Service Constraints

Please see <https://aws.amazon.com/workspaces/> for more information.

73.5 Technical Requirements

Amazon WorkSpaces offers you an easy way to provide a cloud-based desktop experience to your end users. You simply select from a choice of bundles that offer a range of different amounts of CPU, memory, storage, and a choice of applications. Users can connect from a PC, Mac desktop computer, iPad, Kindle, or Android tablet.

- **Administration Guide** – Helps you get started using Amazon WorkSpaces. You will learn how to quickly and easily provision and maintain one or more WorkSpaces. [HTML](#) | [PDF](#) | [Kindle](#)
- **Developer Guide** – Describes the API operations for Amazon WorkSpaces. [HTML](#) | [PDF](#)

74.0 Service Definition – Amazon WorkSpaces Application Manager (Amazon WAM)

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

74.1 Service Overview

Amazon WorkSpaces Application Manager (Amazon WAM) offers a fast, flexible, and secure way for you to deploy and manage applications for Amazon WorkSpaces. Amazon WAM accelerates software deployment, upgrades, patching, and retirement by packaging Microsoft Windows desktop applications into virtualised application containers. These applications run on the end user's Amazon WorkSpaces instance as though they are natively installed. Amazon WAM allows you to build and curate a selection of desktop applications for your users and control access to those applications. You can build an application catalogue of third-party applications that you already own the license for, internally developed custom applications, and applications purchased through the AWS Marketplace for Desktop Apps. Simple, pay-as-you-go, user-based pricing ensures that you only pay for the applications that people use.

The top features include:

- **Build your catalogue** – There are two ways to add applications to your catalogue and make them available for your users: upload applications or shop for applications from the AWS Marketplace for Desktop Apps. Regardless of the method, you can manage user access and monitor real-time usage of applications and user fulfilment.
- **Upload an application** – You can package and deliver your own line-of-business applications or applications for which you already own a license in order to build an application catalogue. You can use Studio and Player to package Microsoft Windows desktop applications into virtualised application containers, validate them, and import them to your application catalogue.
- **Subscribe to an application from AWS Marketplace for Desktop Apps** – In

addition to uploading applications, you can select applications from the AWS Marketplace for Desktop Apps. Applications are available in a variety of categories, including application development, illustration and design, and productivity.

- **Use shared packages** – You can use packages shared with your AWS account to create an application in your catalogue. You can apply customisations to the application you create using Configurable AppEvents set up by the owner of the package.

Top benefits include:

- **Easily manage applications at scale** – Provision desktop applications to Amazon WorkSpaces as virtualised app containers while maintaining centralised control, security, and compliance from an easy-to use management console. Amazon WAM handles all of the upgrades, patching, and retirement of your applications.
- **Curate a wide range of applications** – Amazon WAM enables Microsoft Windows applications to be packaged and deployed to users in your organisation. You can build and manage your application catalogue of third-party applications that you already own the license for, internally developed custom applications, and applications purchased through the AWS Marketplace for Desktop Apps.
- **Visibility and control** – Track real-time usage of applications so you only pay for the applications that are used. You can manage license usage by applications, versions, users, or groups.
- **Empower your users** – Reduce the burden on IT and increase user productivity by allowing users to select and install the applications they need, as they need them.
- **Low cost** – Amazon WAM features simple, pay-as-you-go, monthly, per-user pricing. Eliminate the need for upfront investments in software and hardware used to manage and deploy applications, and reduce operational expenses for managing applications.

74.2 Backup/Restore and Disaster Recovery

This requirement is not applicable to Amazon WorkSpaces Application Manager, as it collates other services rather than being a primary data storage service.

74.3 Pricing Overview

Please visit <https://aws.amazon.com/pricing/> for further details on pricing AWS services.

74.4 Service Constraints

Please see <https://aws.amazon.com/workspaces/applicationmanager/> for more information.

74.5 Technical Requirements

Please refer to <https://aws.amazon.com/workspaces/applicationmanager/> for comprehensive technical documentation regarding Amazon WorkSpaces Application Manager.

75.0 Service Definition - AWS Amplify Console

The AWS Amplify Console provides a Git-based workflow for deploying and hosting fullstack serverless web applications. A fullstack serverless app consists of a backend built with cloud resources such as GraphQL or REST APIs, file and data storage, and a frontend built with single page application frameworks such as React, Angular, Vue, or Gatsby. The Amplify Console accelerates your application release cycle by providing a simple workflow for deploying full-stack serverless applications. You just connect your application's code repository to Amplify Console, and changes to your frontend and backend are deployed in a single workflow on every code commit.

75.1 Service Overview

Top features include

- Manage production and staging environments for your frontend and backend by connecting new branches.
- Set up custom domains managed in Amazon Route 53 with a single click plus get a free HTTPS certificate.
- Preview changes during code reviews by setting up Pull-Request Previews (<https://docs.aws.amazon.com/amplify/latest/userguide/pr-previews.html>).
- Improve your app quality with end to end tests.
- Password protect your web app so you can work on new features without making them publicly accessible.
- Set up rewrites and redirects to maintain SEO rankings and route traffic based on your client app requirements.

Top benefits include

- Use feature based deployments to work on new features without impacting production. Create branch deployments linked to each feature branch.
- Connect your custom domain. If you manage your domain in Amazon Route 53, the Amplify Console automatically connects the root (yourdomain.com), www subdomains (www.yourdomain.com), and branch (<https://dev.yourdomain.com>) subdomains. See custom domains (<https://docs.aws.amazon.com/amplify/latest/userguide/custom-domains.html>).
- Simplify continuous workflows by connecting your repository to 'git push' changes to your frontend and backend in a single workflow.

- Enables your app to go global as it is served via Amazon's reliable content delivery network with hundreds of points of presence globally.
- Achieve Instant+Atomic deployments as all successful deployments are available immediately after an update. Failed deployments do not impact production traffic.
- Share yet-to-be released features with internal stakeholders by setting a username and password.

75.2 Backup/Restore and Disaster Recovery

The AWS Amplify Console connects to your private or public code repositories. You should consider how your code repositories meet your organisational disaster recovery policy. For deployment of an application, the AWS Amplify Console leverages the Amazon CloudFront Global Edge Network to distribute your web app globally.

75.3 Pricing Overview

Please see the AWS UK G Cloud 12 Pricing Document affiliated with this service in the Digital Marketplace.

75.4 Service Constraints

Please see <https://aws.amazon.com/amplify/console/> for more information.

75.5 Technical Requirements

Please refer to <https://aws.amazon.com/amplify/console/> and the following links for comprehensive technical documentation regarding the AWS Amplify Console:

- <https://docs.aws.amazon.com/amplify/latest/userguide/welcome.html>
- <https://aws-amplify.github.io/docs/>

76.0 Service Definition - AWS App Mesh

AWS App Mesh is a service mesh that provides application-level networking to make it easy for your services to communicate with each other across multiple types of compute infrastructure. App Mesh standardizes how your services communicate, giving you end-to-end visibility and ensuring high-availability for your applications.

76.1 Service Overview

Top Feature include

- Separates the logic needed for monitoring and controlling communications into a proxy that manages all network traffic for each service.
- Uses the open source Envoy proxy to manage all traffic into and out of a service's containers.
- Lets you configure services to connect directly to each other instead of requiring code within the application or using a load balancer.

- Works with services managed by Amazon ECS, Amazon EKS, AWS Fargate, Kubernetes running on EC2, and services running directly on EC2.
- AWS App Mesh is a managed and highly available service.

Top benefits include

- App Mesh captures metrics, logs, and traces from all of your applications. You can combine and export this data to Amazon CloudWatch, AWS X-Ray, and compatible AWS partner and community tools for monitoring and tracing. This lets you quickly identify and isolate issues with any service to optimize your entire application.
- App Mesh gives you controls to configure how traffic flows between your services. You can easily implement custom traffic routing rules to ensure every service is highly available during deployments, after failures, and as your application scales.
- App Mesh deploys and configures a proxy that manages all communications traffic to and from your services. This removes the need to configure communication protocols for each service, write custom code, or implement libraries to operate your application.
- You can use App Mesh with services running on any compute service such as AWS Fargate, Amazon ECS, Amazon EKS, and Amazon EC2. App Mesh can monitor and control communications for monoliths running on EC2, teams running containerized applications, orchestration systems, or VPCs as a single application without any code changes.
- You can use AWS App Mesh on AWS Outposts (<https://aws.amazon.com/outposts/>) to configure a service mesh for applications deployed on-premises. AWS Outposts is a fully managed service that extends AWS infrastructure, AWS services, APIs, and tools to virtually any connected site. With AWS App Mesh on Outposts you can provide consistent communication control logic for services across AWS Outposts and AWS cloud to simplify hybrid application networking.
-

76.2 Backup/Restore and Disaster Recovery

This requirement is not applicable for AWS App Mesh. For additional information beyond what is described herein, please refer to <http://aws.amazon.com/documentation/>.

76.3 Pricing Overview

Please see the AWS UK G Cloud 12 Pricing Document affiliated with this service in the Digital Marketplace.

76.4 Service Constraints

Please see <https://docs.aws.amazon.com/app-mesh/latest/userguide/service-quotas.html> for more information.

76.5 Technical Requirements

Please refer to <https://aws.amazon.com/app-mesh/> and the following links for comprehensive technical documentation regarding AWS App Mesh.

- <https://aws.amazon.com/app-mesh/getting-started/>
- <https://docs.aws.amazon.com/app-mesh/latest/userguide/disaster-recovery-resiliency.html>
- <https://docs.aws.amazon.com/app-mesh/latest/userguide/security.html>

77.0 Service Definition – AWS Application Discovery Service

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

77.1 Service Overview

AWS Application Discovery Service helps customers quickly and reliably plan application migration projects by automatically identifying applications running in on premises data centres, their associated dependencies, and their performance profile. Planning data centre migrations can involve thousands of workloads that are often deeply interdependent. Application discovery and dependency mapping are important early first steps in the migration process, but difficult to perform at scale due to the lack of automated tools.

AWS Application Discovery Service automatically collects configuration and usage data from servers, storage and networking equipment to develop a list of applications, how they perform, and how they are interdependent. This information is retained in encrypted format in an AWS Application Discovery Service database which you can export as a CSV or XML file into your preferred visualization tool or cloud migration solution to help reduce the complexity and time in planning your cloud migration.

Top benefits include:

- **Simplify the Cloud Migration Planning Process** – AWS Application Discovery Service automatically identifies installed applications, network dependencies, and system performance data. Data from the service can be used to find multi-tier applications that should be migrated as a group. Start and manage discovery process with a few clicks via the AWS Management Console.
- **Easily Integrate with Cloud Migration Tools** – AWS Application Discovery Service provides a set of public APIs and an open data format. This approach allows you to collect data from other discovery solutions in the Application Discovery database, and use this data in a cloud migration solution or within a

cloud migration framework provided by System Integration partners.

- **Protect Data with End-to-End Encryption** – AWS Application Discovery Service helps ensure that collected application configuration data is protected end-to-end by encrypting it on-premises, maintaining encryption in transit to AWS, and storing in an encrypted format in the Application Discovery database.
- **Consult with AWS Professionals and APN Partners** – AWS Application Discovery Service is available through engagements with AWS Professional Services or through one of our APN Partners. These trained professionals have helped hundreds of enterprise customers successfully complete their migration to the cloud.

77.2 Backup/Restore and Disaster Recovery

This requirement is not applicable for AWS Application Discovery Service. For additional information beyond what is described herein, please refer to <http://aws.amazon.com/documentation/>.

77.3 Pricing Overview

Please visit <https://aws.amazon.com/pricing/> for further details on pricing AWS services.

77.4 Service Constraints

Please see <https://aws.amazon.com/documentation/application-discovery/> for more information.

77.5 Technical Requirements

The AWS Application Discovery Service helps customers quickly and reliably plan application migration projects by automatically identifying applications running in on premises data centres, their associated dependencies, and their performance profile.

- **User Guide** – Walks through how to set up the AWS Application Discovery Service and integrate it with other services. [HTML](#) | [PDF](#)
- **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. [HTML](#) | [PDF](#)

78.0 Service Definition – AWS AppSync

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

78.1 Service Overview

AWS AppSync automatically updates the data in web and mobile applications in real time, and updates data for offline users as soon as they reconnect. AWS AppSync

makes it easy to build collaborative mobile and web applications that deliver responsive, collaborative user experiences.

You can use AWS AppSync to build native mobile and web apps with iOS, Android, JavaScript and React Native. Get started by going to the AWS AppSync console, specify the data for your app with simple code statements, and AWS AppSync will manage everything needed to store, process, and retrieve the data for your application.

Top benefits include:

- **Rapid prototyping and development with GraphQL** - With AWS AppSync you can specify the data requirements of the application with simple code statements and iterate quickly during the prototyping and development process. AWS AppSync uses GraphQL, an open standard query language that makes it easy for applications to request data from the cloud.
- **Build real-time, collaborative mobile and web apps** - AWS AppSync makes it easy to build collaborative apps that keep shared data updated in real time. You just specify the data for your app with simple code statements and AWS AppSync manages everything needed to keep the app data updated in real time. AWS AppSync also automatically updates the data for offline users as soon as they reconnect. Only the data needed is fetched, reducing message traffic and improving application performance and battery life.
- **Automatically manages offline users' data** - AWS AppSync automatically manages all the data operations for offline users. The service supports an offline programming model where application data is not only available for users who are offline, but users can also add and update app data locally as well. This makes it easy to build apps that cache important data locally for offline use, and then synchronize with the cloud when the device reconnects.
- **Helps keep data secure** - AWS AppSync helps you keep users and app data secure. The service integrates with Amazon Cognito and AWS Identity and Access Management, so you can set fine-grained permissions on GraphQL operations that put strict controls on who can access the data.
- **Build apps with data from multiple sources** - AWS AppSync makes it easy to combine data from different sources. For example, you could build an app that uses AppSync to combine a user's activity data from their phone and web usage, fitness tracker activity, and social media postings and combine it with location data to summarize all the user's activities and locations for the day. With AppSync, your app can access data in Amazon DynamoDB, trigger AWS Lambda functions, or run Amazon Elasticsearch queries and combine data from these services to provide the exact data you need for your app.
- **Automatically resolves data conflicts** - AWS AppSync can automatically resolve conflicts on data that is shared between multiple users and devices. Conflicts can happen when two or more users make updates at the same time. The service has default rules to detect conflicts and ensure they are resolved, but

also allows you to add your own business logic into the conflict resolution process with AWS Lambda functions.

78.2 Backup/Restore and Disaster Recovery

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

78.3 Pricing Overview

Please visit <https://aws.amazon.com/pricing/> for further details on pricing AWS services.

78.4 Service Constraints

Please see <https://aws.amazon.com/appsync/> for more information.

78.5 Technical Requirements

AWS AppSync is an enterprise level, fully managed GraphQL service with real-time data synchronization and offline programming features.

- **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. [HTML](#) | [PDF](#)
- **API Reference** – Describes all the API operations for AWS AppSync in detail. Also provides sample requests, responses, and errors for the supported web services protocols. [HTML](#) | [PDF](#)

79.0 Service Definition – AWS Artifact

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

79.1 Service Overview

An audit and compliance portal for on-demand access to download AWS' compliance reports and manage select agreements.

AWS Artifact provides on-demand access to AWS' security and compliance reports and select online agreements. Reports available in AWS Artifact include our Service Organisation Control (SOC) reports, Payment Card Industry (PCI) reports, and certifications from accreditation bodies across geographies and compliance verticals that validate the implementation and operating effectiveness of AWS security controls. Agreements available in AWS Artifact include the Business Associate Addendum (BAA) and the Nondisclosure Agreement (NDA).

Top benefits include:

- **Generate Compliance Reports with Ease and On-Demand** - Save time with on-demand access to AWS' auditor issued reports, certifications, accreditations

and other third-party attestations

- **Globally Available** Save time with on-demand access to AWS' auditor issued reports, certifications, accreditations and other third-party attestations
- **Easy Identification** - Validate the implementation and operating effectiveness of the AWS security control environment from any geography or vertical.
- **Quick Assessments** - Easily identify the scope of each of the audit artefacts, including services, regions and applicable audit dates
- **Continuous monitoring** - Continuously monitor the security and compliance of AWS with immediate access when new reports are released
- **Enhanced Transparency** - Perform due-diligence anytime with enhanced transparency into the AWS control environment.

79.2 Backup/Restore and Disaster Recovery

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

79.3 Pricing Overview

Not applicable as AWS Artifact is a no cost, self-service portal for on-demand access to AWS' compliance reports.

79.4 Service Constraints

Please see <https://aws.amazon.com/artifact/> for more information.

79.5 Technical Requirements

AWS Artifact is a web service that enables you to download AWS security and compliance documents such as ISO certifications and SOC reports.

- **User Guide** - Describes key concepts of AWS Artifact and provides instructions for using the features of AWS Artifact. [HTML](#) | [PDF](#)

80.0 Service Definition – AWS Auto Scaling

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

80.1 Service Overview

AWS Auto Scaling is a web service designed to launch or terminate Amazon EC2 instances automatically based on user-defined policies, schedules, and health checks
Top benefits include:

- **Setup Scaling Quickly** - AWS Auto Scaling lets you set target utilization levels

for multiple resources in a single, intuitive interface. You can quickly see the average utilization of all of your scalable resources without having to navigate to other consoles. For example, if your application uses Amazon EC2 and Amazon DynamoDB, you can use AWS Auto Scaling to manage resource provisioning for all of the EC2 Auto Scaling groups and database tables in your application.

- **Make Smart Scaling Decisions** - AWS Auto Scaling lets you build scaling plans that automate how groups of different resources respond to changes in demand. You can optimize availability, costs, or a balance of both. AWS Auto Scaling automatically creates all of the scaling policies and sets targets for you based on your preference. AWS Auto Scaling monitors your application and automatically adds or removes capacity from your resource groups in real-time as demands change.
- **Automatically Maintain Performance** - Using AWS Auto Scaling, you maintain optimal application performance and availability, even when workloads are periodic, unpredictable, or continuously changing. AWS Auto Scaling continually monitors your applications to make sure that they are operating at your desired performance levels. When demand spikes, AWS Auto Scaling automatically increases the capacity of constrained resources so you maintain a high quality of service.
- **Pay Only For What You Need** - AWS Auto Scaling can help you optimize your utilization and cost efficiencies when consuming AWS services so you only pay for the resources you actually need. When demand drops, AWS Auto Scaling will automatically remove any excess resource capacity so you avoid overspending. AWS Auto Scaling is free to use, and allows you to optimize the costs of your AWS environment.

80.2 Backup/Restore and Disaster Recovery

This requirement is not applicable for AWS Auto Scaling. For additional information beyond what is described herein, please refer to <http://aws.amazon.com/documentation/>

80.3 Pricing Overview

There is no additional charge for AWS Auto Scaling. You pay only for the AWS resources needed to run your applications and [Amazon CloudWatch](#) monitoring fees. For additional information on pricing and pricing resources, please refer to <https://aws.amazon.com/autoscaling/pricing/>.

80.4 Service Constraints

AWS Auto Scaling features and default service limits are found at <https://aws.amazon.com/autoscaling/faqs/>

80.5 Technical Requirements

With AWS Auto Scaling, you can quickly discover the scalable AWS resources for your application and set up dynamic scaling. It uses [Amazon EC2 Auto Scaling](#) to scale your EC2 instances and Application Auto Scaling to scale resources from other services.

- **Administrator Guide** - Provides a conceptual overview of AWS Auto Scaling and includes detailed instructions for using the various features. [HTML](#) | [PDF](#) | [Kindle](#)
- **API Reference** - Describes all the API operations for AWS Auto Scaling in detail. Also provides sample requests, responses, and errors for the supported web services protocols. [HTML](#) | [PDF](#)
- **AWS Auto Scaling section of the AWS CLI Reference** - Describes the AWS CLI commands that you can use to administer AWS Auto Scaling. Provides syntax, options, and usage examples for each command. [HTML](#)

81.0 Service Definition – AWS Backup

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

81.1 Service Overview

AWS Backup is a fully managed backup service that makes it easy to centralise and automate the backup of data across AWS Cloud services in the cloud as well as on premises using AWS Storage Gateway. Using AWS Backup, you can centrally configure backup policies and monitor backup activity for AWS resources, such as Amazon EBS volumes, Amazon RDS databases, Amazon DynamoDB tables, Amazon EFS file systems, and AWS Storage Gateway volumes. AWS Backup automates and consolidates backup tasks previously performed service by service, removing the need to create custom scripts and manual processes. With just a few clicks in the AWS Backup console, you can create backup policies that automate backup schedules and retention management. AWS Backup provides a fully managed, policy-based backup solution, simplifying your backup management and enabling you to meet your business and regulatory backup compliance requirements.

The top features include:

- **Centralised backup management** – AWS Backup provides a centralised backup console, a set of backup APIs, and a CLI to manage backups across the AWS Cloud services that your applications run on, including [Amazon EBS](#), [Amazon RDS](#), [Amazon DynamoDB](#), [Amazon EFS](#), and [AWS Storage Gateway](#).
- **Policy-based backup solution** – With AWS Backup, you can create backup policies called backup plans that enable you to define your backup requirements and then apply them to the AWS resources you want backed up. You can create separate backup plans that meet specific business and regulatory compliance requirements, helping to ensure that each of your AWS resources are backed up

and protected.

- **Automated retention management** – With AWS Backup, you can set backup retention policies that will automatically retain and expire backups according to your business and regulatory backup compliance requirements. Automated backup retention management makes it easy to minimise backup storage costs by retaining backups for only as long as they are needed.
- **Lifecycle management policies** – AWS Backup enables you to meet compliance requirements while minimising backup storage costs by storing backups in a low-cost cold storage tier. You can configure lifecycle policies that will automatically transition backups from warm storage to cold storage according to a schedule that you define. For more information about lifecycle policies, [click here](#).
- **Backup data 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. AWS Backup encrypts your backup data using encryption keys managed by AWS KMS, eliminating the need to build and maintain a key management infrastructure.

Top benefits include:

- **Centrally manage backups** – Configure backup policies from a central backup console, simplifying backup management and making it easy to ensure that your application data across AWS Cloud services is backed up and protected. Use AWS Backup's central console, APIs, or CLI to back up, restore, and set backup retention policies across AWS Cloud services in the cloud and on premises using AWS Storage Gateway.
- **Automate backup processes** – Save time and money with AWS Backup's fully managed, policy-based solution. AWS Backup provides automated backup schedules, retention management, and lifecycle management, removing the need for custom scripts and manual processes. With AWS Backup, you can apply backup policies to your AWS resources by simply tagging them, making it easy to implement your backup strategy across all your AWS resources and ensuring that all your application data is appropriately backed up.
- **Improve backup compliance** – Enforce your backup policies, encrypt your backups, and audit backup activity from a centralised console to help meet your backup compliance requirements. Backup policies make it simple to align your backup strategy with your internal or regulatory requirements. AWS Backup secures your backups by encrypting your data in transit and at rest. Consolidated backup activity logs across AWS Cloud services makes it easier to perform compliance audits. AWS Backup is PCI and ISO compliant as well as HIPAA eligible.

81.2 Backup/Restore and Disaster Recovery

AWS Backup stores data resiliently in Amazon S3. Amazon S3 has a durability of 99.9999999%. However, data can also be copied between AWS Regions using Amazon S3 cross-region replication. more information can be found at <https://docs.aws.amazon.com/AmazonS3/latest/dev/crr.html>.

81.3 Pricing Overview

Please visit <https://aws.amazon.com/pricing/> for further details on pricing AWS services.

81.4 Service Constraints

Please see <https://docs.aws.amazon.com/aws-backup/latest/devguide/aws-backup-limits.html> for more information.

81.5 Technical Requirements

Please refer to <https://docs.aws.amazon.com/aws-backup/latest/devguide/whatisbackup.html> for comprehensive technical documentation regarding AWS Backup.

82.0 Service Definition – AWS Batch

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

82.1 Service Overview

AWS Batch enables developers, scientists, and engineers to easily and efficiently run hundreds of thousands of batch computing jobs on AWS. AWS Batch dynamically provisions the optimal quantity and type of compute resources (e.g., CPU or memory optimised instances) based on the volume and specific resource requirements of the batch jobs submitted. With AWS Batch, there is no need to install and manage batch computing software or server clusters that you use to run your jobs, allowing you to focus on analysing results and solving problems. AWS Batch plans, schedules, and executes your batch computing workloads across the full range of AWS compute services and features, such as Amazon EC2 and Spot Instances.

Top benefits include:

- **Fully Managed** – AWS Batch eliminates the need to operate third-party commercial or open source batch processing solutions. There is no batch software or servers to install or manage. AWS Batch manages all the infrastructure for you, avoiding the complexities of provisioning, managing, monitoring, and scaling your batch computing jobs.
- **Integrated with AWS** – AWS Batch is natively integrated with the AWS platform, allowing you to leverage the scaling, networking, and access management capabilities of AWS. This makes it easy to run jobs that safely and securely

retrieve and write data to and from AWS data stores such as Amazon S3 or Amazon DynamoDB.

- **Cost-Optimised Resource Provisioning** – AWS Batch provisions compute resources and optimises the job distribution based on the volume and resource requirements of the submitted batch jobs. AWS Batch dynamically scales compute resources to any quantity required to run your batch jobs, freeing you from the constraints of fixed-capacity clusters. AWS Batch also dynamically bids for Spot Instances on your behalf, reducing the cost of running your batch jobs further.

82.2 Backup/Restore and Disaster Recovery

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

82.1 Pricing Overview

Please visit <https://aws.amazon.com/pricing/> for further details on pricing AWS services.

82.2 Technical Requirements

AWS Snowball is a service for customers who want to transport terabytes or petabytes of data to and from AWS, or who want to access the storage and compute power of the AWS Cloud locally and cost effectively in places where connecting to the internet may not be an option.

User Guide – Describes key concepts of AWS Batch and provides instructions for using the features of AWS Batch. [HTML](#) | [PDF](#)

API References – Describes all the API operations for AWS Batch in detail. [HTML](#) | [PDF](#)

83.0 Service Definition – AWS Certificate Manager

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

This service listing includes the following AWS Services:

- AWS Certificate Manager
- AWS Certificate Manager Private CA

83.1 Service Overview

AWS Certificate Manager is a service that lets you easily provision, manage, and deploy public and private Secure Sockets Layer/Transport Layer Security (SSL/TLS) certificates for use with AWS services and your internal connected resources. SSL/TLS

certificates are used to secure network communications and establish the identity of websites over the Internet as well as resources on private networks. AWS Certificate Manager removes the time-consuming manual process of purchasing, uploading, and renewing SSL/TLS certificates. With AWS Certificate Manager, you can quickly request a certificate, deploy it on ACM-integrated AWS resources, such as Elastic Load Balancers, Amazon CloudFront distributions, and APIs on API Gateway, and let AWS Certificate Manager handle certificate renewals. It also enables you to create private certificates for your internal resources and manage the certificate lifecycle centrally. Public and private certificates provisioned through AWS Certificate Manager for use with ACM-integrated services are free. You pay only for the AWS resources you create to run your application. For private certificates, you pay monthly for the operation of the private CA and for the private certificates you issue.

Top benefits include:

- **Protect and Secure Your Website** - SSL, and its successor TLS, are industry standard protocols for encrypting network communications and establishing the identity of websites over the Internet. SSL/TLS provides encryption for sensitive data in transit and authentication using SSL/TLS certificates to establish the identity of your site and secure connections between browsers and applications and your site. AWS Certificate Manager provides an easy way to provision and manage these certificates so you can configure a website or application to use the SSL/TLS protocol.
- **Protect and Secure Your Internal Resources** - Private certificates are used for identifying and securing communication between connected resources on private networks, such as servers, mobile and IoT devices, and applications. AWS Certificate Manager (ACM) Private Certificate Authority (CA) is a managed private CA service that helps you easily and securely manage the lifecycle of your private certificates. ACM Private CA provides you a highly-available private CA service without the upfront investment and ongoing maintenance costs of operating your own private CA. ACM Private CA extends ACM's certificate management capabilities to private certificates, enabling you to create and
- manage public and private certificates centrally. ACM Private CA allows developers to be more agile by providing them APIs to create and deploy private certificates programmatically. You also have the flexibility to create private certificates for applications that require custom certificate lifetimes or resource names. Learn more about ACM Private Certificate Authority.
- **Get Certificates Easily** - AWS Certificate Manager removes many of the time-consuming and error-prone steps to acquire an SSL/TLS certificate for your website or application. There is no need to generate a key pair or certificate signing request (CSR), submit a CSR to a Certificate Authority, or upload and install the certificate once received. With a few clicks in the AWS Management Console, you can request a trusted SSL/TLS certificate from AWS. Once the certificate is created, AWS Certificate Manager takes care of deploying certificates to help you enable SSL/TLS for your website or application.

- **Free Public Certificates for ACM-integrated Services** - With AWS Certificate Manager, there is no additional charge for provisioning public or private SSL/TLS certificates you use with ACM-integrated services, such as Elastic Load Balancing and API Gateway. You pay for the AWS resources you create to run your application. For private certificates, ACM Private CA provides you the ability to pay monthly for the service and certificates you create. You pay less per certificate as you create more private certificates.
- **Managed Certificate Renewal** - AWS Certificate Manager manages the renewal process for the certificates managed in ACM and used with ACM-integrated services, such as Elastic Load Balancing and API Gateway. ACM can automate renewal and deployment of these certificates. With ACM Private CA APIs, ACM enables you to automate creation and renewal of private certificates for on premises resources, EC2 instances, and IoT devices.
- **Secure Key Management** - AWS Certificate Manager is designed to protect and manage the private keys used with SSL/TLS certificates. Strong encryption and key management best practices are used when protecting and storing private keys.
- **Centrally Manage Certificates on the AWS Cloud** - You will find it easy to centrally manage AWS Certificate Manager SSL/TLS certificates provided by AWS Certificate Manager in an AWS Region from the AWS Management Console, AWS CLI, or AWS Certificate Manager APIs. You can also audit the use of each certificate by reviewing your Amazon CloudTrail logs.
- **Integrated with Other AWS Cloud Services** - AWS Certificate Manager is integrated with other AWS services, so you can provision an SSL/TLS certificate and deploy it with your Elastic Load Balancer, Amazon CloudFront distribution or API in Amazon API Gateway. AWS Certificate Manager also works with AWS Elastic Beanstalk and AWS CloudFormation for public email-validated certificates to help you manage public certificates and use them with your applications in the AWS Cloud. To deploy a certificate with an AWS resource, you simply select the certificate you want from a drop-down list in the AWS Management Console. Alternatively, you can call an AWS API or CLI to associate the certificate with your resource. AWS Certificate Manager then deploys the certificate to the selected resource for you.
- **Import Third-Party Certificates** - AWS Certificate Manager makes it easy to import SSL/TLS certificates issued by third-party Certificate Authorities (CAs) and deploy them with your Elastic Load Balancers, Amazon CloudFront distributions and APIs on Amazon API Gateway. You can monitor the expiration date of an imported certificate, and import a replacement when the existing certificate is nearing expiration. Alternatively, you can request a free certificate from AWS Certificate Manager and let AWS manage future renewals for you. Importing certificates doesn't cost anything.

83.2 Backup/Restore and Disaster Recovery

This requirement is not applicable for AWS Certificate Manager. For additional information beyond what is described herein, please refer to <http://aws.amazon.com/documentation/>.

83.3 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. ACM Private Certificate Authority (CA) is priced along two dimensions. You pay a monthly fee for the operation of each private CA until you delete it and you pay for the private certificates you issue each month. Please visit <https://aws.amazon.com/pricing/> for further details on pricing AWS services.

83.4 Service Constraints

Please see <https://aws.amazon.com/certificate-manager/> for more information.

83.5 Technical Requirements

AWS Certificate Manager (ACM) makes it easy to provision, manage, and deploy SSL/TLS certificates on AWS-managed resources.

- **User Guide** – Provides conceptual overviews and explains how to provision ACM Certificates on AWS-based websites. [HTML](#) | [PDF](#) | [Kindle](#).
- **API Reference** - Describes the API operations available for ACM along with sample requests, responses, and errors for the supported web services protocols. [HTML](#) | [PDF](#)

83.6 Service Overview – AWS Certificate Manager (ACM) Private Certificate Authority (CA)

AWS Certificate Manager (ACM) Private Certificate Authority (CA) is a managed private CA service that helps you easily and securely manage the lifecycle of your private certificates. ACM Private CA provides you a highly available private CA service without the upfront investment and ongoing maintenance costs of operating your own private CA. ACM Private CA extends ACM's certificate management capabilities to private certificates, enabling you to manage public and private certificates centrally. ACM Private CA allows developers to be more agile by providing them APIs to create and deploy private certificates programmatically. You also have the flexibility to create private certificates for applications that require custom certificate lifetimes or resource names. With ACM Private CA, you can create and manage private certificates for your connected resources in one place, with a secure, pay-as-you-go, managed private CA service.

The top features include:

- **AWS-managed certificate authority** – ACM Private CA is a managed service

that automates time-consuming administrative tasks, such as hardware provisioning, software patching, high availability, and backups. ACM Private CA provides security, configuration, management, and monitoring of a highly available private CA. ACM Private CA allows you to choose from several CA key algorithms and key sizes, including RSA 2048 or 4096 and ECDSA P256 or P384. ACM Private CA operates as a subordinate CA that “chains up” to your existing root CA, allowing you to issue certificates that are trusted within your organisation. You need to have your own root CA to use this service.

- **Certificate lifecycle management** – ACM Private CA is integrated with ACM to allow you to manage both public and private certificates from a single console interface. When you use ACM to request certificates from your private CA, ACM generates and manages the private keys, renews certificates automatically, and deploys certificates to resources on ACM-integrated services, including Elastic Load Balancing load balancers and Amazon API Gateway endpoints. ACM also makes it easy for you to export and deploy private certificates anywhere using API-based automation.
- **Auditing and logging** – ACM Private CA provides you and your auditors with visibility into the activity of your private CAs. You can create audit reports that include the status of all of the certificates issued from the CA. ACM Private CA is integrated with AWS CloudTrail. CloudTrail captures API calls from the ACM Private CA console, from the CLI, or from your code, and delivers the log files to your Amazon S3 bucket. Using the information collected by CloudTrail, you can determine the request that was made, the IP address from which the request came, when it was made, and so on.
- **Secure HSM-backed key storage for CA keys** – The keys used by a certificate authority to sign certificates are highly sensitive. ACM Private CA secures CA keys with AWS-managed HSMs. These HSMs adhere to FIPS 140-2 Level 3 security standards to help protect your private CA against key compromises.
- **IAM integration** – You can control access to ACM Private CA with IAM policies. For example, you can create a policy to grant IT administrators who are responsible for CA management full access to create and configure private CAs while granting limited access to developers and users who need only to issue and revoke certificates.
- **Certificate revocation list (CRL) generation** – ACM Private CA automatically publishes and updates certificate revocation lists to your Amazon S3 bucket. Applications, services, and devices use CRLs to evaluate the status of a certificate each time a connection is made between two resources. For instance, an IoT application can check if the private certificate for a sensor is valid before accepting data from the sensor.
- **API-based automation** – You can write code to automate certificate management in the programming language of your choice using the ACM Private CA and ACM APIs. The AWS SDKs make authentication simpler and integrate efficiently with your development environment. You can also write scripts or one-

off commands using command line tools to interact with the service.

- **Customisation** – ACM Private CA can be used as a standalone service to issue certificates directly without using ACM for certificate and private key management. When used this way, you can create certificates with any subject name; any of the supported key algorithms, key sizes, and signing algorithms; and any validity period, including days, months, or years from the present time or a specific end date.

Top benefits include:

- **Secure and managed private CA** – ACM Private CA provides you an easier and more secure way to create a private CA and use it to create and manage your private certificates. ACM Private CA is secured with AWS-managed HSMs. These HSMs adhere to FIPS 140-2 Level 3 security standards to securely store the keys for your private CA. Private CA administrators can control access to the service using IAM policies. ACM Private CA provides you visibility into private certificate activity and allows you to create reports. You can audit private CA activity using AWS CloudTrail. ACM Private CA also publishes and updates CRLs to Amazon S3 automatically to help prevent the use of revoked certificates. For instance, an IoT application can check if the private certificate for a sensor is valid before accepting data from the sensor.
- **Manage certificates centrally** – ACM Private CA enables you to manage the lifecycle of your private and public certificates. With ACM Private CA you can choose to delegate certificate management to ACM for certificates used with ACM-integrated services, such as Elastic Load Balancing and Amazon API Gateway. You can easily create and deploy private certificates using the AWS Management Console or the AWS APIs. ACM can automate renewal and deployment of these certificates. ACM Private CA also provides you with APIs to automate creation and renewal of private certificates for on-premises resources, EC2 instances, and IoT devices. ACM Private CA gives you the flexibility to manage private certificates on your own without ACM certificate management.
- **Enable developer agility** – ACM Private CA provides you the agility to create and deploy certificates with just a few API calls. ACM Private CA allows you to delegate management of private certificates to developers by allowing them to request certificates from private CAs linked to their AWS accounts. You can also automate certificate creation for use cases that require a high volume of short-lived certificates. For instance, you can automatically create and deploy certificates to identify new EC2 instances and containers in automatically scaling environments or to authenticate event notification messages sent from AWS Lambda functions.
- **Flexibility to customise private certificates** – ACM Private CA can be used as a standalone service, without ACM certificate management, to create and deploy customised private certificates, such as certificates with custom resource names or lifetimes. This flexibility is helpful when you need to identify resources by a specific name (e.g., identifying a device by its serial number or when certificates

cannot be rotated easily, such as certificates embedded into hardware devices during the manufacturing process).

- **Pay-as-you-go pricing** – ACM Private CA is more cost effective compared to the traditional, commercially available options. ACM Private CA provides you the ability to pay monthly for the service and certificates you create and deploy. You pay less as you use more certificates.

83.6.1 Backup/Restore and Disaster Recovery

This requirement is not applicable to ACM Private CA, as it collates other services rather than being a primary data storage service.

83.6.2 Pricing Overview

Please visit <https://aws.amazon.com/pricing/> for further details on pricing AWS services.

83.6.3 Service Constraints

Please see <https://aws.amazon.com/certificate-manager/private-certificate-authority/> for more information.

83.6.4 Technical Requirements

Please refer to <https://aws.amazon.com/certificate-manager/private-certificate-authority/> for comprehensive technical documentation regarding ACM Private CA.

84.0 Service Definition – AWS Cloud Map

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

84.1 Service Overview

AWS Cloud Map is a cloud resource discovery service. With AWS Cloud Map, you can define custom names for your application resources, and the service 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.

The top features include:

- **Discover resources via API calls or DNS queries** – AWS Cloud Map allows your applications to discover any web-based service via AWS SDKs, API calls, or DNS queries.
- **Simplified service naming** – AWS Cloud Map lets you define simple custom names for services in your application.
- **Assign custom attributes** – AWS Cloud Map lets you define custom attributes for each resource, such as location and deployment stage.
- **Access control** – AWS Cloud Map is integrated with IAM to ensure that only

authenticated services can discover resources within the registry and retrieve the locations and credentials for those resources.

- **Fully managed** – AWS Cloud Map eliminates the need to set up, update, and manage your own service discovery tools and software.

Top benefits include:

- AWS Cloud Map allows you to register any application resources, such as databases, queues, microservices, and other cloud resources with custom names. The service then constantly checks the health of resources to make sure the location is up to date.
- AWS Cloud Map provides a single registry for all of 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.
- AWS Cloud Map can reliably obtain up-to-date health statuses for all of your application resources.
- AWS Cloud Map provides a single up-to-date registry of service names and locations that enables microservices to easily locate one another.

84.2 Backup/Restore and Disaster Recovery

This requirement is not applicable to AWS Cloud Map, as it isn't a data storage service.

84.3 Pricing Overview

Please visit <https://aws.amazon.com/pricing/> for further details on pricing AWS services.

84.4 Service Constraints

Please see <https://aws.amazon.com/cloud-map/> for more information.

84.5 Technical Requirements

Please refer to <https://aws.amazon.com/cloud-map/faqs/> for comprehensive technical documentation regarding AWS Cloud Map.

85.0 Service Definition – AWS Cloud9

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

85.1 Service Overview

AWS Cloud9 is a cloud-based integrated development environment (IDE) that lets you write, run, and debug your code with just a browser. It includes a code editor, debugger, and terminal. AWS Cloud9 comes pre-packaged with essential tools for popular

programming languages, including JavaScript, Python, PHP, and more, so you don't need to install files or configure your development machine to start new projects.

The top features include:

- **Fully featured editor** – AWS Cloud9 includes a browser-based editor that makes it easy to write, run, and debug your projects.
- **Broad selection of run configurations** – AWS Cloud9 supports over 40 programming languages and application types, including JavaScript, Python, PHP, Ruby, Go, and C++.
- **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, or C/C++ app.
- **Integrated tools for serverless development** – AWS Cloud9 allows you to easily build serverless applications by writing and debugging serverless application code.
- **Collaborative editing and chat** – AWS Cloud9 makes it easy for multiple developers on your team to actively see each other type and pair program together on the same file. AWS Cloud9 allows you to use the built-in chat capability to communicate with your team without having to leave the IDE.

Top benefits include:

- **Code with just a browser** – AWS Cloud9 lets you write, run, and debug applications with just a browser, without needing to install or maintain a local IDE.
- **Code together in real time** – You can share your development environment with your team in just a few clicks and pair program together.
- **Build serverless applications with ease** – AWS Cloud9 preconfigures the development environment with all the SDKs, libraries, and plugins needed for serverless development.
- **Locally test Lambda functions** – AWS Cloud9 provides an environment for locally testing and debugging AWS Lambda functions.
- **Direct terminal access to AWS** – AWS Cloud9 comes with a terminal that includes sudo privileges to the managed Amazon EC2 instance that is hosting your development environment and a preauthenticated AWS CLI.
- **Preconfigured languages** – AWS Cloud9's development environment comes pre-packaged with tooling for over 40 programming languages, including Node.js, JavaScript, Python, PHP, Ruby, Go, and C++.
- **Start projects quickly** – Start writing code for popular application stacks within minutes by eliminating the need to install or configure files, SDKs, and plugins for your development machine.
- **Cloud based** – Because AWS Cloud9 is cloud based, you can easily maintain

multiple development environments to isolate your project's resources.

- **Integrated tools** – The AWS Cloud9 code editor and integrated debugger include helpful, time-saving features such as code hinting, code completion, and step-through debugging.
- **Connectivity to any Linux server platform** – You have the flexibility to run AWS Cloud9 development environments on a managed Amazon EC2 Linux instance or any Linux server that you are using today.

85.2 Backup/Restore and Disaster Recovery

Any data that you store in your AWS Cloud9 environment, such as code files, packages, or dependencies, is always stored in your resources. If you use an Amazon EC2 environment, your data is stored in the associated Amazon EBS volume that exists in your AWS account. These volumes can utilise snapshots to provide a point-in-time recovery of the data. Snapshots can be done in real time while the volume is attached and in use.

If you use an SSH environment, your data is stored in local storage on your Linux server and would utilise your existing data backup procedures.

85.3 Pricing Overview

Please visit <https://aws.amazon.com/pricing/> for further details on pricing AWS services.

85.4 Service Constraints

Please see <https://aws.amazon.com/cloud9/faqs/> for more information.

85.5 Technical Requirements

Please refer to <https://aws.amazon.com/cloud9/faqs/> for comprehensive technical documentation regarding AWS Cloud9.

86.0 Service Definition – AWS CloudFormation

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

86.1 Service Overview

AWS CloudFormation gives developers and systems administrators an easy way to create and manage a collection of related AWS resources, provisioning and updating them in an orderly and predictable fashion.

You can use AWS CloudFormation's [sample templates](#) or create your own templates to describe the AWS resources, and any associated dependencies or runtime parameters, required to run your application. You don't need to figure out the order for provisioning AWS Cloud services or the subtleties of making those dependencies work. AWS CloudFormation takes care of this for you. After the AWS resources are deployed, you

can modify and update them in a controlled and predictable way, in effect applying version control to your AWS infrastructure the same way you do with your software. You can also visualise your templates as diagrams and edit them using a drag-and-drop interface with the [AWS CloudFormation Designer](#).

You can deploy and update a template and its associated collection of resources (called a stack) by using the AWS Management Console, AWS CLI, or APIs. AWS CloudFormation is available at no additional charge, and you pay only for the AWS resources needed to run your applications.

Top benefits include:

- **Supports a Wide Range of AWS Resources** - AWS CloudFormation supports a wide range of AWS resources, allowing you to build a highly available, reliable, and scalable AWS infrastructure for your application needs.
- **Easy to Use** – AWS CloudFormation makes it easy to organise and deploy a collection of AWS resources and lets you describe any dependencies or special parameters to pass in at runtime. You can use one of the many CloudFormation [sample templates](#)—either verbatim or as a starting point.
- **Declarative and Flexible** – To create the infrastructure you want, you enumerate what AWS resources, configuration values, and interconnections you need in a template and then let AWS CloudFormation do the rest with a few simple clicks in the AWS Management Console, one command by using the AWS CLI, or a single request by calling the APIs. You won't have to recall the details of how to create and interconnect the respective AWS resources via their service APIs; AWS CloudFormation does this for you. You also don't need to write a template from scratch if you start with one of the many sample templates that come with AWS CloudFormation.
- **Infrastructure as Code** – A template can be used repeatedly to create identical copies of the same stack (or to use as a foundation to start a new stack). You can capture and control region-specific infrastructure variations such as Amazon EC2 AMIs, as well as Amazon EBS and Amazon RDS snapshot names.
- Templates are simple JSON-formatted text files that can be placed under your normal source control mechanisms, stored in private or public locations such as Amazon S3, and exchanged via email. With AWS CloudFormation, you can “open the hood,” to see exactly which AWS resources make up a stack. You retain full control and have the ability to modify any of the AWS resources created as part of a stack.
- **Customised Via Parameters** – You can use parameters to customise aspects of your template at run time when the stack is built. For example, you can pass the Amazon RDS database size, Amazon EC2 instance types, database, and web server port numbers to AWS CloudFormation when you create a stack. You can also use a parameterised template to create multiple stacks that may differ in a controlled way. For example, your Amazon EC2 instance types, Amazon CloudWatch alarm thresholds, and Amazon RDS read-replica settings may differ

among AWS Regions if you receive more customer traffic in the US than in Europe. You can use template parameters to tune the settings and thresholds in each region separately and still be sure that the application is deployed consistently across the regions.

- **Visualise and Edit with Drag-and-Drop Interface** – [AWS CloudFormation Designer](#) provides a visual diagram of your template with icons representing your AWS resources and arrows showing their relationships. You can build and edit templates using the drag-and-drop interface, then edit the template details using the integrated JSON text editor. AWS CloudFormation Designer allows you to spend more time designing your AWS infrastructure and less time manually coding your templates.
- **Integration Ready** – You can integrate AWS CloudFormation with the development and management tools of your choice. AWS CloudFormation publishes progress events through the Amazon SNS. With Amazon SNS you can track stack creation and deletion progress via email and integrate with other processes programmatically.

86.2 Backup/Restore and Disaster Recovery

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

86.3 Pricing Overview

Please visit <https://aws.amazon.com/pricing/> for further details on pricing AWS services.

86.4 Service Constraints

Please see <https://aws.amazon.com/cloudformation/> for more information.

86.5 Technical Requirements

AWS CloudFormation enables you to create and provision AWS infrastructure deployments predictably and repeatedly. It helps you leverage AWS products such as Amazon EC2, Amazon EBS, Amazon SNS, Elastic Load Balancing, and Auto Scaling to build highly reliable, highly scalable, cost-effective applications in the cloud without worrying about creating and configuring the underlying AWS infrastructure. AWS CloudFormation enables you to use a template file to create and delete a collection of resources together as a single unit (a stack).

- **User Guide** – Provides a conceptual overview of AWS CloudFormation and includes instructions on using the various features with the AWS CLI. [HTML](#) | [PDF](#) | [Kindle](#)
- **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. [HTML](#) | [PDF](#)

- **AWS CloudFormation in the AWS CLI Reference** – Describes the AWS CloudFormation commands that are available in the AWS CLI. [HTML](#)

87.0 Service Definition – AWS CloudHSM

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

87.1 Service Overview

The AWS CloudHSM service helps you meet corporate, contractual and regulatory compliance requirements for data security by using dedicated HSM appliances within the AWS Cloud. With CloudHSM, you control the encryption keys and cryptographic operations performed by the HSM.

The AWS CloudHSM service allows you to protect your encryption keys within HSMs designed and validated to government standards for secure key management. You can securely generate, store, and manage the cryptographic keys used for data encryption such that they are accessible only by you. AWS CloudHSM helps you comply with strict key management requirements without sacrificing application performance.

The AWS CloudHSM service works with Amazon VPC. CloudHSM instances are provisioned inside your Amazon VPC with an IP address that you specify, providing simple and private network connectivity to your Amazon EC2 instances. Placing AWS CloudHSM instances near your Amazon EC2 instances decreases network latency, which can improve application performance. AWS provides dedicated and exclusive (single tenant) access to AWS CloudHSM instances, isolated from other AWS customers.

Top benefits include:

- **Secure Key Storage** – As part of the service, you have dedicated access to HSM capabilities in the cloud. AWS CloudHSM protects your cryptographic keys with tamper-resistant HSM appliances that are designed to comply with international (Common Criteria EAL4+) and U.S. Government (National Institute of Standards and Technology [NIST] Federal Information Processing Standards [FIPS] 140-2) regulatory standards for cryptographic modules. You retain full control of your keys and cryptographic operations on the HSM, while Amazon manages and maintains the hardware without having access to your keys.
- **Contractual and Regulatory Compliance** – By protecting your keys in hardware and preventing them from being accessed by third parties, AWS CloudHSM can help you comply with the most stringent regulatory and contractual requirements for key protection.
- **On-Demand** – The AWS CloudHSM API, CLI Tools, and SDK let you start and stop dedicated AWS CloudHSM instances whenever you want.
- **Secure Connectivity** – AWS CloudHSM instances are in your VPC, so it is easy to use them with your Amazon EC2 applications. You use standard Amazon VPC

security mechanisms to control access to AWS CloudHSM instances.

- **Improved Application Performance** – By placing AWS CloudHSM instances in your Amazon VPC near your Amazon EC2 instances, you can reduce network latency and increase the performance of your AWS applications that use HSMs.
- **Integrated** – You can use AWS CloudHSM with Amazon Redshift, Amazon RDS Oracle, or third-party applications such as SafeNet ProtectV volume encryption for EBS, Apache (SSL termination), or Microsoft SQL Server (transparent data encryption). You can also use AWS CloudHSM when writing your own applications and continue to use the standard cryptographic libraries you're familiar with, including PKCS#11, Java JCA/JCE, and Microsoft CAPI and CNG.
- **Auditable** – If you need to track resource changes, or audit activities for security and compliance purposes, you can review all of the AWS CloudHSM API calls made from your account through AWS CloudTrail. Additionally, you can audit operations on the HSM appliance using syslog or send syslog log messages to your own collector.

87.2 Backup/Restore and Disaster Recovery

AWS CloudHSM is available in multiple regions and Availability Zones to help you build highly available applications that require strong key protection. The AWS CloudHSM CLI Tools can help you configure high availability groups that span multiple Availability Zones, so you can build resilient applications. In the unlikely event of a hardware failure, you can launch a new AWS CloudHSM instance and replicate the keys to the new HSM with a few commands. You can also use AWS CloudHSM with your compatible on-premises HSMs to securely store keys in your datacentre. This increases key durability and gives you the flexibility to securely migrate keys in and out of AWS.

87.3 Pricing Overview

Please visit <https://aws.amazon.com/pricing/> for further details on pricing AWS services.

87.4 Service Constraints

Please see <https://aws.amazon.com/cloudhsm/> for more information.

87.5 Technical Requirements

Please refer to the links below and <https://aws.amazon.com/documentation/cloudhsm/> and the following links for comprehensive technical documentation regarding AWS CloudHSM.

- **User Guide** – Provides key concepts of AWS CloudHSM and documents advanced product features and the AWS CLI. [HTML](#) | [PDF](#) | [Kindle](#)

API Reference – Describes the API operations for AWS CloudHSM. [HTML](#) | [PDF](#)

88.0 Service Definition – AWS CloudTrail

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

88.1 Service Overview

AWS CloudTrail is a web service that records AWS API calls for your account and delivers log files to you. The recorded information includes the identity of the API caller, the time of the API call, the source IP address of the API caller, the request parameters, and the response elements returned by the AWS service.

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

Top features and benefits include:

- **Increased Visibility** – AWS CloudTrail provides increased visibility into your user activity by recording AWS API calls. You can answer questions such as, what actions did a given user take over a given time period? For a given resource, which user has taken actions on it over a given time period? What is the source IP address of a given activity? Which activities failed due to inadequate permissions?
- **Durable and Inexpensive Log File Storage** – AWS CloudTrail uses Amazon S3 for log file storage and delivery, so log files are stored durably and inexpensively. You can use Amazon S3 lifecycle configuration rules to further reduce storage costs. For example, you can define rules to automatically delete old log files or archive them to Amazon Glacier for additional savings.
- **Easy Administration** – AWS CloudTrail is a fully managed service; you simply turn on AWS CloudTrail for your account using the AWS Management Console, the AWS CLI, or the AWS CloudTrail SDK and start receiving AWS CloudTrail log files in the Amazon S3 bucket that you specify.
- **Notifications for Log File Delivery** – AWS CloudTrail can be configured to publish a notification for each log file delivered, thus enabling you to automatically take action upon log file delivery. CloudTrail uses the Amazon SNS for notifications.
- **Log File Aggregation** – AWS CloudTrail can be configured to aggregate log files across multiple accounts and regions so that log files are delivered to a single bucket. For detailed instructions, refer to the Aggregating CloudTrail Log Files to a Single Amazon S3 Bucket section of the user guide.
- **Reliable and Timely Delivery** – AWS CloudTrail continuously transports events from AWS Cloud services using a highly available and fault tolerant processing

pipeline. CloudTrail typically delivers events within 15 minutes of the API call.

- **Troubleshoot operational or security issues** – You can troubleshoot operational issues or perform security analysis by looking up API activity that was captured for your AWS account. Using the AWS CloudTrail console, AWS CLI, or AWS SDKs, you can quickly and easily answer questions related to API activity for the last 7 days and take immediate action.
- **Receive SNS Notifications of API activity** – AWS CloudTrail can be configured to be deliver API activity to an Amazon CloudWatch Logs log group you specify. You can then create CloudWatch Alarms to receive Amazon SNS notifications when specific API activity occurs.

88.2 Backup/Restore and Disaster Recovery

AWS CloudTrail uses Amazon S3 for log file storage and delivery, so log files are stored durably and inexpensively. Refer to Amazon S3 for more detailed information.

88.3 Pricing Overview

This service is not discussed in the AWS UK G-Cloud 12 Pricing Document as it is available free of charge. AWS CloudTrail allows you to setup a trail that delivers a single copy of management events in each region free of charge. Once a CloudTrail trail is setup, Amazon S3 charges apply based on your usage. You will be charged for any data events or additional copies of management events recorded in that region:

- **Management events:** provide insights into the management (“control plane”) operations performed on resources in your AWS account. For example, CloudTrail delivers management events for API calls such as launching Amazon EC2 instances or creating Amazon S3 buckets. Management events are enabled by default when you configure a trail and record supported activity at the account level. The first copy of management events within each region is delivered free of charge. Additional copies of management events are charged according to the pricing spreadsheet.
- **Data events:** provide insights into the resource (“data plane”) operations performed on or within the resource itself. Data events are often high volume activities and include operations such as Amazon S3 object level APIs and Lambda function invoke API. For example, CloudTrail delivers data events for AWS Lambda Invoke API calls and Amazon S3 object level APIs such as Get, Put, Delete and List actions. Data events are recorded only for the Lambda functions and S3 buckets you specify and are charged according to the pricing spreadsheet.

88.4 Service Constraints

Please see <https://aws.amazon.com/cloudtrail/> for more information.

88.5 Technical Requirements

Please refer to the links below and <https://aws.amazon.com/documentation/cloudtrail/> and the following links for comprehensive technical documentation regarding AWS CloudTrail.

- **User Guide** – Provides detailed descriptions of product concepts and includes instructions for using the various features with both the console and the CLI. [HTML](#) | [PDF](#)
- **API Reference** – Describes all the API operations for AWS CloudTrail in detail. Also provides sample requests, responses, and errors for the supported web services protocols. [HTML](#) | [PDF](#)

89.0 Service Definition – AWS CodeBuild

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

89.1 Service Overview

AWS CodeBuild is a fully managed build service that compiles source code, runs tests, and produces software packages that are ready to deploy. With AWS CodeBuild, you don't need to provision, manage, and scale your own build servers. AWS CodeBuild scales continuously and processes multiple builds concurrently, so your builds are not left waiting in a queue. You can get started quickly by using pre-packaged build environments, or you can create custom build environments that use your own build tools. With AWS CodeBuild, you are charged by the minute for the compute resources you use.

Top benefits include:

- **Fully Managed Build Service** – AWS CodeBuild eliminates the need to set up, patch, update, and manage your own build servers and software. There is no software to install or manage.
- **Continuous Scaling** – AWS CodeBuild scales automatically to meet your build volume. It immediately processes each build you submit and can run separate builds concurrently, which means your builds are not left waiting in a queue.
- **Pay as You Go** – With AWS CodeBuild, you are charged based on the number of minutes it takes to complete your build. This means you no longer have to worry about paying for idle build server capacity.
- **Extensible** – You can bring your own build tools and programming runtimes to use with AWS CodeBuild by creating customised build environments in addition to the pre-packaged build tools and runtimes supported by AWS CodeBuild.
- **Enables Continuous Integration and Delivery** – AWS CodeBuild belongs to a family of [AWS Code Services](#), which you can use to create complete, automated

software release workflows for [continuous integration/continuous delivery](#) (CI/CD). You can also integrate AWS CodeBuild into your existing CI/CD workflow. For example, you can use AWS CodeBuild as a worker node for your existing Jenkins server setup for distributed builds.

- **Secure** – with AWS CodeBuild, your build artefacts are encrypted with customer specific keys that are managed by the [AWS KMS](#). CodeBuild is integrated with [AWS IAM](#), so you can assign user-specific permissions to your build projects.

89.2 Backup/Restore and Disaster Recovery

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

89.3 Pricing Overview

Please visit <https://aws.amazon.com/pricing/> for further details on pricing AWS services.

89.4 Service Constraints

Please see <https://aws.amazon.com/codebuild/> for more information.

89.5 Technical Requirements

AWS CodeBuild is a fully managed build service that compiles your source code, runs unit tests, and produces artefacts that are ready to deploy.

- **User Guide** – Describes how to get started building your source code using AWS CodeBuild. [HTML](#) | [PDF](#) | [Kindle](#)
- **API Reference** – Describes the API operations for AWS CodeBuild. Also provides details of related request and response syntax and errors. [HTML](#) | [PDF](#)
- **AWS CLI Reference for AWS CodeBuild** – Describes the AWS CLI commands that you can use to automate building your source code. [HTML](#)

90.0 Service Definition – AWS CodeCommit

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

90.1 Service Overview

AWS CodeCommit is a fully managed source control service that makes it easy for companies to host secure and highly scalable private Git repositories. AWS

CodeCommit eliminates the need to operate your own source control system or worry about scaling its infrastructure. You can use AWS CodeCommit to securely store anything from source code to binaries, and it works seamlessly with your existing Git tools.

Top benefits include:

- **Fully Managed** – AWS CodeCommit eliminates the need to host, maintain, back up, and scale your own source control servers. The service automatically scales to meet the growing needs of your project.
- **Secure** – AWS CodeCommit automatically encrypts your files in transit and at rest. AWS CodeCommit is integrated with AWS IAM, allowing you to assign user specific permissions to your repositories.
- **High Availability** – AWS CodeCommit has a highly scalable, redundant, and durable architecture. The service is designed to keep your repositories highly available and accessible.
- **Store Anything** – AWS CodeCommit allows you to store any type of file, and there are no repository size limits. This allows you to store and version application assets, like images and libraries, along with your code.
- **Faster Development Lifecycle** – AWS CodeCommit keeps your repositories close to your build, staging, and production environments in the AWS Cloud. You can transfer incremental changes instead of the entire application. This allows you to increase the speed and frequency of your development lifecycle.
- **Use Your Existing Tools** – AWS CodeCommit supports all Git commands and works with your existing Git tools. You can keep using your preferred development environment plugins, continuous integration/continuous delivery systems, and graphical clients with AWS CodeCommit.

90.2 Backup/Restore and Disaster Recovery

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

90.3 Pricing Overview

Please visit <https://aws.amazon.com/pricing/> for further details on pricing AWS services.

90.4 Service Constraints

Please see <https://aws.amazon.com/codecommit/> for more information.

90.5 Technical Requirements

AWS CodeCommit is a version control service that enables you to privately store and manage Git repositories in the AWS Cloud.

User Guide – Describes how to get started and store your assets using AWS CodeCommit. [HTML](#) | [PDF](#) | [Kindle](#)

API Reference – Describes all the API operations for AWS CodeCommit in detail. Also

provides details of request and response syntax and errors for the supported web services protocols. [HTML](#) | [PDF](#)

AWS CLI Reference for AWS CodeCommit – Describes the AWS CLI commands that you can use to automate management of your assets. [HTML](#)

91.0 Service Definition – AWS CodeDeploy

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

91.1 Service Overview

AWS CodeDeploy is a service that automates code deployments to any instance, including Amazon EC2 instances and instances running on-premises. AWS CodeDeploy makes it easier for you to rapidly release new features, helps you avoid downtime during application deployment, and handles the complexity of updating your applications. You can use AWS CodeDeploy to automate software deployments, eliminating the need for error-prone manual operations, and the service scales with your infrastructure so you can easily deploy to one instance or thousands.

Top benefits include:

- **Automated Deployments** – AWS CodeDeploy fully automates your code deployments, allowing you to deploy reliably and rapidly. You can consistently deploy your application across your development, test, and production environments on any instance, whether running in the cloud or on-premises. The service scales with your infrastructure so you can deploy to one instance or thousands.
- **Minimise Downtime** – AWS CodeDeploy helps maximise your application availability during the software deployment process. It performs rolling updates across your instances and tracks application health according to configurable rules. Code deployments can easily be stopped and rolled back if there are errors.
- **Centralised Control** – AWS CodeDeploy allows you to easily launch and track the status of your application deployments through the AWS Management Console or the AWS CLI. AWS CodeDeploy gives you a detailed report, allowing you to view when each application revision was deployed and to which instances. You can also create push notifications to receive live updates about your deployments.
- **Easy to Adopt** – AWS CodeDeploy is platform and language agnostic and works with any application. You can easily reuse your existing setup code. AWS CodeDeploy can also integrate with your existing software release process or continuous delivery toolchain (e.g., Jenkins).

91.2 Backup/Restore and Disaster Recovery

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

91.1 Pricing Overview

Please visit <https://aws.amazon.com/pricing/> for further details on pricing AWS services.

91.2 Service Constraints

Please see <https://aws.amazon.com/codedeploy/> for more information.

91.3 Technical Requirements

AWS CodeDeploy is a deployment service that enables developers to automate the deployment of applications to instances and to update the applications as required.

- **User Guide** – Describes how to get started and deploy applications using AWS CodeDeploy. [HTML](#) | [PDF](#) | [Kindle](#)
- **API Reference** – Describes all the API operations for AWS CodeDeploy. Also provides sample requests, responses, and errors for the supported web services protocols. [HTML](#) | [PDF](#)
- **AWS CLI Reference for AWS CodeDeploy** – Describes the AWS CLI commands that you can use to automate deployments. [HTML](#)

92.0 Service Definition – AWS CodePipeline

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

92.1 Service Overview

AWS CodePipeline is a [continuous delivery](#) service for fast and reliable application updates. AWS CodePipeline builds, tests, and deploys your code every time there is a code change, based on the release process models you define. This enables you to rapidly and reliably deliver features and updates. You can easily build out an end-to-end solution by using our prebuilt plugins for popular third-party services like GitHub or integrating your own custom plugins into any stage of your release process. With AWS CodePipeline, you only pay for what you use. There are no up-front fees or long-term commitments.

Top benefits include:

- **Rapid Delivery** – AWS CodePipeline automates your software release process, allowing you to rapidly release new features to users. With AWS CodePipeline, you can quickly iterate on feedback and get new features to customers faster.

- **Improved Quality** – Automating your build, test, and release process allows you to easily test each code change and catch bugs while they are small and simple to fix. You can assure the quality of your code by running each change through your standardised release process.
- **Configurable Workflow** – AWS CodePipeline allows you to model the different stages of your software release process through a graphical user interface. You can specify the tests to run and the steps to deploy your application and its dependencies.
- **Get Started Fast** – With AWS CodePipeline, you can immediately begin to model your software release process. There are no servers to provision or set up. AWS CodePipeline is a fully managed continuous delivery service that connects to your existing tools and systems.
- **Easy to Integrate** – AWS CodePipeline can easily be extended to adapt to your specific needs. You can use our prebuilt plugins or your own custom plugins in any step of your release process. For example, you can pull your source code from GitHub, use your on-premises Jenkins build server, run load tests using a third-party service, or pass on deployment information to your custom operations dashboard.

92.2 Backup/Restore and Disaster Recovery

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

92.3 Pricing Overview

Please visit <https://aws.amazon.com/pricing/> for further details on pricing AWS services.

92.4 Service Constraints

Please see <https://aws.amazon.com/codepipeline/> for more information.

92.5 Technical Requirements

AWS CodePipeline is a continuous delivery service that enables you to model, visualise, and automate the steps required to release your software.

- **User Guide** – Walks through how to set up AWS CodePipeline and integrate it with other services. [HTML](#) | [PDF](#) | [Kindle](#)
- **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. [HTML](#) | [PDF](#)
- **AWS CLI Reference for AWS CodePipeline** – Describes the AWS CLI commands that you can use to manage your assets. [HTML](#)

93.0 Service Definition – AWS CodeStar

The following subsections provide service definition information.

93.1 Service Overview

AWS CodeStar enables you to quickly develop, build, and deploy applications on AWS. AWS CodeStar provides a unified user interface, enabling you to easily manage your software development activities in one place. With AWS CodeStar, you can set up your entire [continuous delivery](#) toolchain in minutes, allowing you to start releasing code faster. AWS CodeStar makes it easy for your whole team to work together securely, allowing you to easily manage access and add owners, contributors, and viewers to your projects. Each AWS CodeStar project comes with a project management dashboard, including an integrated issue tracking capability powered by Atlassian JIRA Software. With the AWS CodeStar project dashboard, you can easily track progress across your entire software development process, from your backlog of work items to teams' recent code deployments.

There is no additional charge for using AWS CodeStar. You only pay for the AWS resources that you provision for developing and running your application (for example, Amazon EC2 instances).

Top benefits include:

- **Start developing on AWS in minutes** - AWS CodeStar makes it easy for you to set up your entire development and continuous delivery toolchain for coding, building, testing, and deploying your application code. To start a project, you can choose from a variety of AWS CodeStar templates for [Amazon EC2](#), [AWS Lambda](#), and [AWS Elastic Beanstalk](#). You have the option to choose [AWS CodeCommit](#) or GitHub to use as your project's source control. You also have the option to edit your source code using one of several options including [AWS Cloud9](#), Microsoft Visual Studio, or Eclipse. After you make your selections the underlying AWS services are provisioned in minutes, allowing you to quickly start coding and deploying your applications.
- **Manage software delivery in one place** - AWS CodeStar provides an easy way to coordinate your day-to-day development activities through a unified user interface, reducing the need to switch between various service consoles. AWS CodeStar's project dashboard lets you monitor application activity, and track progress across all stages of your software development process, including code commits, builds, tests, and deployments, from a central place. AWS CodeStar integrates [Atlassian JIRA Software](#), a third-party issue tracking and project management tool, allowing you to easily manage JIRA issues directly in the AWS CodeStar dashboard.
- **Work across your team securely** - AWS CodeStar enables you to collaborate on projects across your team in a secure manner. You can easily manage access for project owners, contributors, and viewers without needing to manually

configure your own policy for each service. AWS CodeStar simplifies the process of setting up project access for teams by providing built-in role-based policies that follow AWS [Identity and Access Management](#) best practices.

- **Choose from a variety of project templates** - With AWS CodeStar project templates, you can easily develop a variety of applications such as websites, web applications, web services, and Alexa skills. AWS CodeStar project templates include the code for getting started on supported programming languages including Java, JavaScript, PHP, Ruby, C#, and Python.

93.2 Backup/Restore and Disaster Recovery

This requirement is not applicable for AWS CodeStar. For additional information beyond what is described herein, please refer to <https://docs.aws.amazon.com/codestar/latest/userguide/welcome.html>.

93.3 Pricing Overview

Please visit <https://aws.amazon.com/pricing/> for further details on pricing AWS services.

93.4 Service Constraints

Please see <https://aws.amazon.com/codestar/faqs/> for more information.

93.5 Technical Requirements

AWS CodeStar lets you quickly develop, build, and deploy applications on AWS.

- **User Guide** - Provides conceptual overviews of AWS CodeStar and explains how to use it to develop software applications on AWS. [HTML](#) | [PDF](#)
- **API Reference** - Describes all the API operations for AWS CodeStar. Provides examples of request and response syntax. [HTML](#) | [PDF](#)

94.0 Service Definition – AWS Config

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

94.1 Service Overview

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

Top benefits include:

- **Configuration Visibility** – You can view continuously updated details of all configuration attributes associated with AWS resources. You are notified via Amazon SNS of every configuration change and you can process these notifications programmatically.
- **Fully Managed** – With AWS Config, there are no software agents to install, and no databases to manage. AWS Config automatically manages the underlying infrastructure required to record, store and report on the configuration details of your AWS resources.
- **Easy to get started** – You can enable AWS Config with a few clicks in the AWS Management Console. AWS Config will discover your AWS resources and start recording configuration changes. You can access information about the configuration of any resource using the AWS Management Console, CLI, or SDKs.
- **Low cost** – With AWS Config, there are no upfront costs. You pay-as-you-go based on the number of resources and configuration changes recorded for your AWS account.

94.2 Backup/Restore and Disaster Recovery

AWS Config uses an Amazon S3 bucket you specify to store the information recorded. Refer to Amazon S3 for more detailed information.

94.3 Pricing Overview

Please visit <https://aws.amazon.com/pricing/> for further details on pricing AWS services.

94.4 Service Constraints

Please see <https://aws.amazon.com/config/> for more information.

94.5 Technical Requirements

AWS Config provides a detailed view of the resources associated with your AWS account, including how they are configured, how they are related to one another, and how the configurations, and their relationships have changed over time.

- **Developer Guide** – Provides a conceptual overview of AWS Config and includes detailed development instructions for using the various features. [HTML](#) | [PDF](#)
- **CLI Reference** – Documents the AWS Config CLI. [HTML](#)
- **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. [HTML](#) | [PDF](#)

95.0 Service Definition – AWS Data Pipeline

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

95.1 Service Overview

AWS Data Pipeline is a web service that helps you reliably process and move data between different AWS compute and storage services, as well as on-premises data sources, at specified intervals. With AWS Data Pipeline, you can regularly access your data where it's stored, transform and process it at scale, and efficiently transfer the results to AWS Cloud services such as Amazon S3, Amazon RDS, Amazon DynamoDB, and Amazon EMR.

AWS Data Pipeline helps you easily create complex data processing workloads that are fault tolerant, repeatable, and highly available. You don't have to worry about ensuring resource availability, managing inter-task dependencies, retrying transient failures or timeouts in individual tasks, or creating a failure notification system. AWS Data Pipeline also allows you to move and process data that was previously locked up in on-premises data silos.

Top benefits include:

- **Reliable** – AWS Data Pipeline is built on a distributed, highly available infrastructure designed for fault-tolerant execution of your activities. If failures occur in your activity logic or data sources, AWS Data Pipeline automatically retries the activity. If the failure persists, AWS Data Pipeline sends you failure notifications via [Amazon SNS](#). You can configure your notifications for successful runs, delays in planned activities, or failures.
- **Simple** – Creating a pipeline is quick and easy via our drag-and-drop console. Common preconditions are built into the service, so you don't need to write any extra logic to use them. For example, you can check for the existence of an Amazon S3 file by simply providing the name of the Amazon S3 bucket and the path of the file that you want to check for, and AWS Data Pipeline does the rest. In addition to its easy visual pipeline creator, AWS Data Pipeline provides a library of pipeline templates. These templates make it simple to create pipelines for a number of more complex use cases, such as regularly processing your log files, archiving data to Amazon S3, or running periodic SQL queries.
- **Flexible** – AWS Data Pipeline allows you to take advantage of a variety of features such as scheduling, dependency tracking, and error handling. You can use activities and preconditions that AWS provides and/or write your own custom ones. This means that you can configure AWS Data Pipeline to take actions like run Amazon EMR jobs, execute SQL queries directly against databases, or execute custom applications running on Amazon EC2 or in your own data centre. This allows you to create powerful custom pipelines to analyse and process your data without having to deal with the complexities of reliably scheduling and

executing your application logic.

- **Scalable** – AWS Data Pipeline makes it equally easy to dispatch work to one machine or many, in serial or parallel. With AWS Data Pipeline's flexible design, processing a million files is as easy as processing a single file.
- **Low Cost** – AWS Data Pipeline is inexpensive to use and is billed at a low monthly rate. You can try it for free under the AWS Free Usage Tier.
- **Transparent** – You have full control over the computational resources that execute your business logic, making it easy to enhance or debug your logic. Additionally, full execution logs are automatically delivered to Amazon S3, giving you a persistent, detailed record of what has happened in your pipeline.

95.2 Backup/Restore and Disaster Recovery

This requirement is not applicable for AWS Data Pipeline. For additional information beyond what is described herein, please refer to <http://aws.amazon.com/documentation/>.

95.3 Pricing Overview

Please visit <https://aws.amazon.com/pricing/> for further details on pricing AWS services.

95.4 Service Constraints

Please see <https://aws.amazon.com/datapipeline/> for more information.

95.5 Technical Requirements

AWS Data Pipeline is a web service that you can use to automate the movement and transformation of data. With AWS Data Pipeline, you can define data-driven workflows, so that tasks can be dependent on the successful completion of previous tasks.

- **Developer Guide** – Provides a conceptual overview of AWS Data Pipeline and includes detailed development instructions for using the various features. [HTML](#) | [PDF](#) | [Kindle](#)
- **API Reference** – Describes all the API operations for AWS Data Pipeline in detail. Also provides sample requests, responses, and errors for the supported web services protocols. [HTML](#) | [PDF](#)

96.0 Service Definition – AWS Database Migration Service

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

96.1 Service Overview

AWS Database Migration Service helps you migrate databases to AWS easily and securely. The source database remains fully operational during the migration, minimizing downtime to applications that rely on the database. AWS Database Migration Service can migrate your data to and from most widely used commercial and open-source databases. The service supports homogenous migrations such as Oracle to Oracle, as well as heterogeneous migrations between different database platforms, such as Oracle to Amazon Aurora or Microsoft SQL Server to MySQL. It also allows you to stream data to Amazon Redshift from any of the supported sources including Amazon Aurora, PostgreSQL, MySQL, MariaDB, Oracle, and SQL Server, enabling consolidation and easy analysis of data in the petabyte-scale data warehouse.

Top benefits include:

- **Simple to Use** – AWS Database Migration Service is simple to use. There is no need to install any drivers or applications, and it does not require changes to the source database in most cases. You can begin a database migration with just a few clicks in the AWS Management Console. Once the migration has started, AWS manages all the complexities of the migration process including automatically replicating data changes that occur in the source database during the migration process.
- **Zero Downtime** – AWS Database Migration Service helps you migrate your databases to AWS with virtually no downtime. All data changes to the source database that occur during the migration are continuously replicated to the target, allowing the source database to be fully operational during the migration process. After the database migration is complete, the target database will remain synchronised with the source for as long as you choose, allowing you to switch over the database at a convenient time.
- **Supports Most Widely Used Databases** – AWS Database Migration Service can migrate your data to and from most of the widely used commercial and open source databases. It supports homogeneous migrations such as Oracle to Oracle, as well as heterogeneous migrations between different database platforms, such as Oracle to Amazon Aurora. Migrations can be from on premises databases to Amazon RDS or Amazon EC2, databases running on Amazon EC2 to Amazon RDS, or vice versa, as well as from one Amazon RDS database to another Amazon RDS database.
- **Low Cost** – AWS Database Migration Service is a low-cost service. You only pay for the compute resources used during the migration process and any additional log storage. Migrating a terabyte-sized database can be done for as little as a few dollars per month (see Supplier Pricing Document for pricing details). This applies to both homogeneous and heterogeneous migrations of any supported databases. This is in stark contrast to conventional database migration methods, which can be very expensive.
- **Fast and Easy to Set Up** – You can set up a migration task within minutes in the

AWS Management Console. A migration task is where you define the parameters AWS Database Migration Service uses to execute the migration. This includes setting up connections to the source and target databases, as well as choosing the replication instance used to run the migration process. Once set up, the same task can be used for test runs before performing the actual migration.

- **Reliable** – AWS Database Migration Service is highly resilient and self-healing. It continually monitors source and target databases, network connectivity, and the replication instance. In case of interruption, it automatically restarts the process and continues the migration from where it was halted. Detailed diagnostic information is available for you to take the necessary corrective action for errors that cannot be automatically resolved.

96.2 Backup/Restore and Disaster Recovery

This requirement is not applicable for AWS Database Migration Service. For additional information beyond what is described herein, please refer to <http://aws.amazon.com/documentation/>.

96.3 Pricing Overview

Please visit <https://aws.amazon.com/pricing/> for further details on pricing AWS services.

96.4 Service Constraints

Please see <https://aws.amazon.com/dms/> for more information.

96.5 Technical Requirements

AWS Database Migration Service is a web service you can use to migrate data from your database that is on-premises, on an Amazon RDS DB instance, or in a database on an Amazon EC2 instance to a database on an AWS Cloud service. These services can include a database on Amazon RDS or a database on an Amazon EC2 instance. You can also migrate a database from an AWS Cloud service to an on-premises database. You can migrate data between heterogeneous or homogenous database engines.

- **User Guide** – Describes all AWS DS concepts and provides instructions on using the various features with both the console and the AWS CLI. [HTML](#) | [PDF](#)
- **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. [HTML](#) | [PDF](#)

97.0 Service Definition – AWS DataSync

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

97.1 Service Overview

AWS DataSync is a data transfer service that makes it easy for you to automate moving data between on-premises storage and Amazon S3 or Amazon EFS.

The top features include:

- **Parallel, scalable architecture** – Connections between the local DataSync agent and the in-cloud service components are multi-threaded, parallelising large files to improve performance over the wide area network (WAN). Service can scale performance to fully utilise a 10 Gbps network link.
- **Automatic infrastructure management** – After you deploy the DataSync agent, connect it to your on-premises storage, and select your Amazon S3 bucket or Amazon EFS file system as the destination or source for data transfers, the DataSync service manages the rest of the infrastructure, including automatically scaling in-cloud resources.
- **File system integration and metadata preservation** – Where Amazon EFS is the destination for your data, DataSync preserves your existing directory structures and file metadata. Similarly, when copying data to Amazon S3, DataSync automatically converts each file to be a single Amazon S3 object in a 1:1 relationship and preserves file metadata as Amazon S3 object metadata.
- **Data encryption and validation** – All of your data is encrypted in transit with TLS. DataSync integrates with AWS KMS so you can encrypt data at rest. For each transfer, DataSync also performs integrity checks both in transit and at rest.
- **Cost effective** – You only pay for data copied by the service, at a flat, per-gigabyte rate. No software licenses, contracts, maintenance fees, or development cycles are required.
- **Integration with AWS infrastructure and management services** – DataSync works natively with AWS infrastructure and management services to make data movement simpler and to provide a consistent management experience.

Top benefits include:

- **Simplify and automate transfers** – AWS DataSync makes it easy for you to move data over the network between on-premises storage and AWS. DataSync automates both the management of data transfer processes and the infrastructure required for high-performance, secure data transfer.
- **Automatic encryption of data** - Transfer data rapidly over the network into AWS, up to 10 times faster than is common with open source tooling. DataSync uses a purpose-built network protocol and a parallel, multi-threaded architecture to accelerate your transfers.
- **Reduce operational costs** – Move data cost effectively with DataSync's flat, per-gigabyte pricing.

97.2 Backup/Restore and Disaster Recovery

This requirement is not applicable to AWS DataSync. For additional information beyond what is described herein, please refer to <http://aws.amazon.com/documentation/>.

97.3 Pricing Overview

Please visit <https://aws.amazon.com/pricing/> for further details on pricing AWS services.

97.4 Service Constraints

Please see <https://docs.aws.amazon.com/datasync/latest/userguide/datasync-limits.html> for more information.

97.5 Technical Requirements

Please refer to <https://docs.aws.amazon.com/datasync/latest/userguide/getting-started.html> for comprehensive technical documentation regarding DataSync.

98.0 Service Definition - AWS DeepRacer

Developers of all skill levels can get hands on with machine learning through a cloud based 3D racing simulator, fully autonomous 1/18th scale race car driven by reinforcement learning, and global racing league.

98.1 Service Overview

Top Features include

- Build models in Amazon SageMaker and train, test, and iterate quickly and easily on the track in the AWS DeepRacer 3D racing simulator.
- Experience the thrill of the race in the real-world when you deploy your reinforcement learning model onto AWS DeepRacer.
- Compete in the world's first global, autonomous racing league, to race for prizes and glory and a chance to advance to the Championship Cup.

Top benefits include

- 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 either a single-car time-trial format or dual-car head-to-head racing 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.

- Master the basics with time-trial racing- Develop your first RL model quickly and easily in the AWS DeepRacer console, using getting started tutorials. Experiment and grow your skills by tuning and retraining your model, to successfully navigate around the track and achieve the fastest lap time!
- Expand your skills with head-to-head racing- Experiment with additional sensors and new training algorithms, to create an RL model capable of skilfully avoiding obstacles on the road, or detecting and predicting driving behaviour of other cars in dual-car head-to-head races.

98.2 Backup/Restore and Disaster Recovery

For additional information beyond what is described herein, please refer to [_http://aws.amazon.com/documentation/_](http://aws.amazon.com/documentation/) (<http://aws.amazon.com/documentation/>).

98.3 Pricing Overview

Please see the AWS UK G Cloud 12 Pricing Document affiliated with this service in the Digital Marketplace.

98.4 Service Constraints

Please see the following links for information:

<https://aws.amazon.com/deepracer/faqs/>

98.5 Technical Requirements

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

- <https://docs.aws.amazon.com/deepracer/latest/developerguide/what-is-deepracer.html>
- <https://docs.aws.amazon.com/deepracer/latest/developerguide/deepracer-get-started.html>
- <https://docs.aws.amazon.com/deepracer/latest/developerguide/deepracer-setup.html>
- <https://aws.amazon.com/deepracer/faqs/>

99.0 Service Definition – AWS Device Farm

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

99.1 Service Overview

AWS Device Farm is an app testing service that lets you test and interact with your Android, iOS, and web apps on many devices at once or reproduce issues on a device in real time. View video, screenshots, logs, and performance data to pinpoint and fix issues and increase quality before shipping your app. AWS Device Farm lets you test your application on a shared fleet of 2500+ devices or on your own private device lab in the cloud.

The top features include:

- **Automated testing** – Test your app in parallel against a massive collection of physical devices in the AWS Cloud. Use one of our built-in frameworks to test your applications without having to write or maintain test scripts or use one of our supported automation testing frameworks.
- **Remote access** – You can gesture, swipe, and interact with devices in real time, directly from your web browser.
- **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 how 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.
- **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, and application data and by installing prerequisite apps to simulate real-world customer conditions.

Top benefits include:

- **Choose the tests that work for you** – Run our built-in test suite (no scripting required) or customise 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.
- **Set up 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.

99.2 Backup/Restore and Disaster Recovery

This requirement is not applicable to AWS Device Farm. For additional information beyond what is described herein, please refer to <http://aws.amazon.com/documentation>.

99.3 Pricing Overview

Please visit <https://aws.amazon.com/pricing/> for further details on pricing AWS services.

99.4 Service Constraints

Please see <https://docs.aws.amazon.com/devicefarm/latest/developerguide/limits.html> for more information.

99.5 Technical Requirements

Please refer to <https://docs.aws.amazon.com/devicefarm/latest/developerguide/welcome.html> and <https://docs.aws.amazon.com/devicefarm/latest/developerguide/getting-started.html> for comprehensive technical documentation regarding AWS Device Farm.

100.0 Service Definition – AWS Direct Connect

The following subsections provide service definition information.

100.1 Service Overview

AWS Direct Connect makes it easy to establish a dedicated network connection from your premises to AWS. Using AWS Direct Connect, you can establish private connectivity between AWS and your data centre, office, or colocation environment, which in many cases can reduce your network costs, increase bandwidth throughput, and provide a more consistent network experience than Internet-based connections.

AWS Direct Connect lets you establish a dedicated network connection between your network and one of the AWS Direct Connect locations. Using industry-standard 802.1q Virtual LANS (VLANs), this dedicated connection can be partitioned into multiple virtual interfaces. This allows you to use the same connection to access public resources such as objects stored in Amazon S3 using public IP address space and private resources such as Amazon EC2 instances running within an Amazon VPC using private IP space, all while maintaining network separation between the public and private environments.

Virtual interfaces can be reconfigured at any time to meet your changing needs.

Top benefits include:

- **Reduces Your Bandwidth Costs** – If you have bandwidth-heavy workloads that you wish to run in AWS, AWS Direct Connect reduces your network costs into and out of AWS in two ways. First, by transferring data to and from AWS directly, you can reduce your bandwidth commitment to your Internet service provider. Second, all data transferred over your dedicated connection is charged at the reduced AWS Direct Connect data transfer rate rather than Internet data transfer

rates.

- **Consistent Network Performance** – Network latency over the Internet can vary because the Internet is constantly changing how data gets from point A to B. With AWS Direct Connect, you choose the data that uses the dedicated connection and how that data is routed, which can provide a more consistent network experience over Internet-based connections.
- **Compatible with all AWS Cloud Services** – AWS Direct Connect is a network service and works with all AWS Cloud services that are accessible over the Internet, such as Amazon S3, Amazon EC2, and Amazon VPC.
- **Private Connectivity to Your Amazon VPC** – You can use AWS Direct Connect to establish a private virtual interface from your on-premises network directly to your Amazon VPC, providing you with a private, high-bandwidth network connection between your network and your VPC.
- **Elastic** – AWS Direct Connect makes it easy to scale your connection to meet your needs.
- **Simple** – You can sign up for AWS Direct Connect quickly and easily using the [AWS Management Console](#).

100.2 Backup/Restore and Disaster Recovery

This requirement is not applicable for AWS Direct Connect. For additional information beyond what is described herein, please refer to <http://aws.amazon.com/documentation>.

100.3 Pricing Overview

Please visit <https://aws.amazon.com/pricing/> for further details on pricing AWS services.

100.4 Service Constraints

Please see <http://aws.amazon.com/documentation/direct-connect/> for more information.

100.5 Technical Requirements

Please refer to <http://aws.amazon.com/documentation/directconnect/> and the following links for comprehensive technical documentation regarding AWS Direct Connect.

- **User Guide** – Provides a conceptual overview of AWS Direct Connect and includes instructions on using the various features with the CLI. Available in [HTML](#), [PDF](#), and [Kindle](#) formats.
- **API Reference** – Describes all the API operations for AWS Direct Connect in detail. Also provides sample requests, responses, and errors for the supported web services protocols. Available in [HTML](#) and [PDF](#) formats.

101.0 Service Definition – AWS Directory Service

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

101.1 Service Overview

AWS Directory Service makes it easy to setup and run Microsoft Active Directory (AD) in the AWS Cloud or connect your AWS resources with an existing on-premises Microsoft Active Directory. Once your directory is created, you can use it to manage users and groups, provide single sign-on to applications and services, create and apply group policy, domain join Amazon EC2 instances, as well as simplify the deployment and management of cloud-based Linux and Microsoft Windows workloads.

AWS Directory Service provides you with three directory types to choose from, including AWS Directory Service for Microsoft Active Directory (Enterprise Edition), also referred to as Microsoft AD, as well as Simple AD and AD Connector.

Top benefits include:

- **Simplify Deployment of Linux and Windows Workloads on AWS** – You can use AWS Directory Service to provision a managed Microsoft Active Directory, enabling you to manage users and group memberships, domain-joining Amazon EC2 Linux and Windows computers, perform Kerberos Single Sign-On (SSO), apply Group Policies and create trust relationships between domains. This makes it easier to deploy and manage Amazon EC2 instances and deploy directory-aware Windows-based workloads, including SharePoint, custom .NET, and SQL Server-based applications.
- **Easy to Get Started; Pay as You Go** – Getting started is easy. You can use the AWS Management Console, or the API to provision the directory type that meets your needs. Once your directory is running, you pay only for the directory hours you use, whether you require a directory for a handful of users or tens of thousands of users.
- **Seamless End User Access to Applications** – AWS Directory Service enables your end users to use their existing corporate credentials when accessing AWS applications, such as Amazon WorkSpaces, Amazon WorkDocs, and Amazon WorkMail, as well as directory-aware Microsoft applications, including SharePoint, custom .NET, and SQL Server-based applications.
- **Managed Service** – AWS Directory Service helps to reduce management tasks. 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.
- **Seamless Administration of AWS Resources** – AWS Directory Service

enables your IT administrators to use their existing corporate credentials to access the AWS Management Console via AWS IAM roles to manage all your AWS resources (e.g., Amazon EC2 instances or Amazon S3 buckets). There is no need to stand up and manage federation infrastructure of your own.

101.2 Backup/Restore and Disaster Recovery

This requirement is not applicable for AWS Directory Service. For additional information beyond what is described herein, please refer to <http://aws.amazon.com/documentation/>.

101.3 Pricing Overview

Please visit <https://aws.amazon.com/pricing/> for further details on pricing AWS services.

101.4 Service Constraints

Please see <http://aws.amazon.com/directoryservice/> for more information.

101.5 Technical Requirements

AWS Directory Service is a managed service that makes it easy to connect AWS Cloud services to your existing on-premises Microsoft Active Directory (AD Connector), or to set up and operate a new directory in the AWS Cloud (Simple AD and AWS Directory Service for Microsoft Active Directory). Your directory users and groups can access the AWS Management Console and AWS applications, such as Amazon WorkSpaces and Amazon WorkDocs, using their existing credentials.

- **Administration Guide** – Describes how to create and manage an AWS Directory Service directory. [HTML](#) | [PDF](#)
- **API Reference** – Describes the API operations for AWS Directory Service. [HTML](#) | [PDF](#)

102.0 Service Definition – AWS Elastic Beanstalk

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

102.1 Service Overview

AWS Elastic Beanstalk is an easy-to-use service for deploying and scaling web applications and services developed with Java, [.NET](#), PHP, Node.js, Python, Ruby, Go, and [Docker](#) on familiar servers such as Apache, Nginx, Passenger, and [IIS](#).

You can simply upload your code and AWS Elastic Beanstalk automatically handles the deployment, from capacity provisioning, load balancing, and auto scaling to application health monitoring. At the same time, you retain full control over the AWS resources powering your application and can access the underlying resources at any time.

Top benefits include:

- **Fast and Simple to Begin** – AWS Elastic Beanstalk is the fastest and simplest way to deploy your application on AWS. You simply use the AWS Management Console, a Git repository, or an Integrated Development Environment (IDE) such as Eclipse or Visual Studio to upload your application, and AWS Elastic Beanstalk automatically handles the deployment details of capacity provisioning, load balancing, auto scaling, and application health monitoring. Within minutes, your application will be ready to use without any infrastructure or resource configuration work on your part.
- **Developer Productivity** – AWS Elastic Beanstalk provisions and operates the infrastructure and manages the application stack (platform) for you, so you don't have to spend the time or develop the expertise. It will also keep the underlying platform running your application up to date with the latest patches and updates. Instead, you can focus on writing code rather than spending time managing and configuring servers, databases, load balancers, firewalls, and networks.
- **Impossible to Outgrow** – AWS Elastic Beanstalk automatically scales your application up and down based on your application's specific need using easily adjustable Auto Scaling settings. For example, you can use CPU utilization metrics to trigger Auto Scaling actions. With AWS Elastic Beanstalk, your application can handle peaks in workload or traffic while minimizing your costs.
- **Complete Resource Control** – You have the freedom to select the AWS resources, such as Amazon EC2 instance type, that are optimal for your application. Additionally, AWS Elastic Beanstalk lets you “open the hood” and 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 AWS Elastic Beanstalk's management capabilities.

102.2 Backup/Restore and Disaster Recovery

This requirement is not applicable for AWS Elastic Beanstalk. For additional information beyond what is described herein, please refer to <http://aws.amazon.com/documentation/>.

102.3 Pricing Overview

Please visit <https://aws.amazon.com/pricing/> for further details on pricing AWS services.

102.4 Service Constraints

Please see <https://aws.amazon.com/elasticbeanstalk/> for more information.

102.5 Technical Requirements

With AWS Elastic Beanstalk, you can quickly deploy and manage applications in the AWS Cloud without worrying about the infrastructure that runs those applications. AWS

Elastic Beanstalk reduces management complexity without restricting choice or control. You simply upload your application, and AWS Elastic Beanstalk automatically handles the details of capacity provisioning, load balancing, scaling, and application health monitoring.

- **Developer Guide** – Provides conceptual overviews of AWS Elastic Beanstalk as well as detailed instructions on the variety of tools for deploying applications and managing applications and environments. [HTML](#) | [PDF](#) | [Kindle](#)
- **API Reference** – Describes all the API operations for AWS Elastic Beanstalk in detail. Also provides sample requests, responses, and errors for the supported web services protocols. [HTML](#) | [PDF](#)
- **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. [HTML](#)

103.0 Service Definition - AWS Elemental MediaConnect

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

103.1 Service Overview

AWS Elemental MediaConnect is a high-quality transport service for live video. Today, broadcasters and content owners rely on satellite networks or fibre connections to send their high-value content into the cloud or to transmit it to partners for distribution. Both satellite and fibre approaches are expensive, require long lead times to set up, and lack the flexibility to adapt to changing requirements. To be more nimble, some customers have tried to use solutions that transmit live video on top of IP infrastructure, but have struggled with reliability and security.

Now you can get the reliability and security of satellite and fibre combined with the flexibility, agility, and economics of IP-based networks using AWS Elemental MediaConnect. MediaConnect enables you to build mission-critical live video workflows in a fraction of the time and cost of satellite or fibre services. You can use MediaConnect to ingest live video from a remote event site (like a stadium), share video with a partner (like a cable TV distributor), or replicate a video stream for processing (like an over-the-top service). MediaConnect combines reliable video transport, highly secure stream sharing, and real-time network traffic and video monitoring that allow you to focus on your content, not your transport infrastructure.

Top features include:

- **Comprehensive range of video industry standard protocols** - AWS Elemental MediaConnect supports a range of protocols for video delivery including the Zixi protocol, Real-Time Transport Protocol (RTP), and RTP with forward error correction (FEC).
- **Built-in security** - AWS Elemental MediaConnect lets you protect your content

using industry-standard, end-to-end AES encryption, and you can enable whitelisting to limit access only to trusted sources.

- **High-quality video sharing** - Create entitlements in MediaConnect, which 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 centre to share live video, you can now reach a larger number of customers quickly and easily using the AWS Cloud.
- **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.
- **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.

Top benefits include:

- **Reliably transport video** - MediaConnect adds a video-specific quality-of-service layer over standard IP transport, enabling uninterrupted, high-quality delivery.
- **Securely share live video streams** - MediaConnect lets you secure your live video using industry-standard encryption, and only share content with authorized customers.
- **Easily manage high-value live broadcasts** - MediaConnect provides real-time visibility into more than 15 critical metrics of your video stream quality, automatically adjusting settings to optimize performance.
- **Lower your costs by 30 % or more** - MediaConnect uses your IP infrastructure or AWS Direct Connect to enable cost-effective, high-quality live video to the AWS Cloud. Instead of long-term commitments to satellite transponders and fixed fibre networks, MediaConnect provides a pay-as-you-go transport alternative, saving you 30% or more compared to a typical satellite primary distribution use case with 70 destinations.

103.2 Backup/Restore and Disaster Recovery

This requirement is not applicable for AWS Elemental MediaConnect. For additional information beyond what is described herein, please refer to <https://aws.amazon.com/mediaconnect/resources/>.

103.3 Pricing Overview

Please visit <https://aws.amazon.com/pricing/> for further details on pricing AWS services.

103.4 Service Constraints

Please see <https://docs.aws.amazon.com/mediacore/latest/ug/limits.html> for more information.

103.5 Technical Requirements

Please refer to <https://docs.aws.amazon.com/mediacore/latest/ug/what-is.html> for comprehensive technical documentation regarding AWS Elemental MediaConnect.

104.0 Service Definition – AWS Elemental MediaConvert

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

104.1 Service Overview

AWS Elemental MediaConvert is a file-based video transcoding service with broadcast-grade features. It allows you to easily create video-on-demand (VOD) content for broadcast and multiscreen delivery at scale. The service combines advanced video and audio capabilities with a simple web services interface and pay-as-you-go pricing. With AWS Elemental MediaConvert, you can focus on delivering compelling media experiences without having to worry about the complexity of building and operating your own video processing infrastructure.

Top benefits include:

- **Broadcast-grade capabilities** - AWS Elemental MediaConvert lets you use a wide range of internet and professional media formats to produce high quality video outputs that look great on any device. With support for ultra-high definition resolutions, high dynamic range video, graphic overlays, advanced audio features, content protection, and closed captioning, AWS Elemental MediaConvert offers a full set of tools to deliver high-quality viewing experiences.
- **Reliable and easy to manage** - AWS Elemental MediaConvert does not require any set up, management, or maintenance of underlying infrastructure. Simply submit jobs with the video processing settings you want and get started without spending time or resources managing transcoding infrastructure. Not only does AWS Elemental MediaConvert provision the required resources to process your jobs, but it also monitors them automatically, so you don't need to worry about reliability.
- **Simple, predictable pricing** - AWS Elemental MediaConvert lets customers create high-quality, end-to-end video processing workflows in the cloud without upfront investment or capital expenditures for video processing infrastructure. You simply pay based on the duration of video that is processed and the features you use.

104.2 Backup/Restore and Disaster Recovery

This requirement is not applicable for AWS Elemental MediaConvert. For additional information beyond what is described herein, please refer to <https://aws.amazon.com/mediaconvert/>.

104.3 Pricing Overview

Please visit <https://aws.amazon.com/pricing/> for further details on pricing AWS services.

104.4 Service Constraints

Please see <https://aws.amazon.com/mediaconvert/resources/> for more information.

104.5 Technical Requirements

User Guide - Describes the components and features that AWS Elemental MediaConvert provides and how to use them. [HTML](#) | [PDF](#)

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. [HTML](#) | [PDF](#)

105.0 Service Definition – AWS Elemental MediaLive

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

105.1 Service Overview

AWS Elemental MediaLive is a broadcast-grade live video processing service. It lets you create high-quality video streams for delivery to broadcast televisions and internet connected multiscreen devices, like connected TVs, tablets, smart phones, and set-top boxes. The service works by encoding your live video streams in real-time, taking a larger-sized live video source and compressing it into smaller versions for distribution to your viewers. With AWS Elemental MediaLive, you can easily set up streams for both live events and 24x7 channels with advanced broadcasting features, high availability, and pay-as-you-go pricing. AWS Elemental MediaLive lets you focus on creating compelling live video experiences for your viewers without the complexity of building and operating broadcast-grade video processing infrastructure.

Top benefits include:

- **Simple deployment and management** - AWS Elemental MediaLive automates the provisioning and management of all the infrastructure used for video encoding, letting you deploy a simple live channel in minutes. The service transparently provisions resources and manages all the scaling, failover, monitoring, and reporting needed to power a live video stream. This lets you focus on your live content, not your encoding infrastructure.

- **Broadcast-grade capabilities** - AWS Elemental MediaLive makes it easy for anyone to produce broadcast-quality live streaming video. The service includes support for advanced capabilities such as ad marker support, audio features including audio normalization and Dolby audio, and multiple caption standards. You can use channel pre-sets to get started quickly, or adjust settings to precisely match your encoding requirements.
- **Highly available** - AWS Elemental MediaLive provides built-in reliability and resiliency. The service transparently manages resources across multiple Availability Zones, and automatically monitors their health, so that any potential issues are detected and resolved without disrupting live channels. With AWS Elemental MediaLive, you can exceed the reliability of infrastructure typically used for broadcast workloads with a straightforward pay-as-you-go model based on the hours of content processed.
- **Increased efficiency and reduced cost** - With AWS Elemental MediaLive, you only pay for the service as you use it, with no upfront investment in encoding infrastructure and no operational overhead devoted to managing physical resources. Pricing for AWS Elemental MediaLive uses a straightforward model based on the hours of content processed and delivered.

105.2 Backup/Restore and Disaster Recovery

This requirement is not applicable for AWS Elemental MediaLive. For additional information beyond what is described herein, please refer to <https://aws.amazon.com/medialive/resources/>.

105.3 Pricing Overview

Please visit <https://aws.amazon.com/pricing/> for further details on pricing AWS services.

105.4 Service Constraints

Please see <https://aws.amazon.com/medialive/resources/> for more information.

105.5 Technical Requirements

- **User Guide** – Describes the components and features that AWS Elemental MediaLive provides and how to use them. [HTML](#) | [PDF](#)
- **API Reference** – Describes all the AWS Elemental MediaLive API operations. Also provides sample requests, responses, and errors for the supported web services protocols. [HTML](#) | [PDF](#)

106.0 Service Definition – AWS Elemental MediaPackage

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

106.1 Service Overview

AWS Elemental MediaPackage reliably prepares and protects your video for delivery over the Internet. From a single video input, AWS Elemental MediaPackage creates video streams formatted to play on connected TVs, mobile phones, computers, tablets, and game consoles. It makes it easy to implement popular video features for viewers (start-over, pause, rewind, etc.), like those commonly found on DVRs. AWS Elemental MediaPackage can also protect your content using Digital Rights Management (DRM). AWS Elemental MediaPackage scales automatically in response to load, so your viewers will always get a great experience without you having to accurately predict in advance the capacity you'll need.

Top benefits include:

- Reach a wide range of connected devices - AWS Elemental
- MediaPackage makes it easy to package and distribute your content to a broad range of video playback devices, including web players, smart phones, game consoles, tablets, and connected TVs.
- Advanced video experiences and content protection - AWS Elemental MediaPackage lets you configure a DVR-like experience for viewers of your live stream. It offers support for more than 20 digital rights management (DRM) providers, supports advanced audio features, and multi-language subtitle tracks.
- Built-in scalability and reliability - AWS Elemental MediaPackage scales automatically as the audience for your video grows. It automatically manages resources across multiple Availability Zones, and monitors their health, so that any potential issues are detected and resolved without disrupting your live video stream.
- Easy integration with AWS cloud services - AWS Elemental MediaPackage is built to work with Amazon CloudFront CDN for global distribution and with AWS Elemental MediaLive for live encoding to form a complete solution for live video processing and delivery. Integration with Amazon CloudWatch gives you real-time monitoring and notifications.

106.2 Backup/Restore and Disaster Recovery

This requirement is not applicable for AWS Elemental MediaPackage. For additional information beyond what is described herein, please refer to <http://aws.amazon.com/documentation/>.

106.3 Pricing Overview

Please visit <https://aws.amazon.com/pricing/> for further details on pricing AWS services.

106.4 Service Constraints

Please see <https://aws.amazon.com/mediapackage/resources/> for more information.

106.5 Technical Requirements

- **User Guide** - Describes the components and features that AWS Elemental MediaPackage provides and how to use them. [HTML](#) | [PDF](#)
- **API Reference** - Describes all the AWS Elemental MediaPackage API operations in detail. Also provides sample requests, responses, and errors for the supported web services protocols. [HTML](#) | [PDF](#)

107.0 Service Definition – AWS Elemental MediaStore

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

107.1 Service Overview

AWS Elemental MediaStore is an AWS storage service optimized for media. It gives you the performance, consistency, and low latency required to deliver live streaming video content. AWS Elemental MediaStore acts as the origin store in your video workflow. Its high performance capabilities meet the needs of the most demanding media delivery workloads, combined with long-term, cost-effective storage.

Top benefits include:

- **High performance, optimized for video** - AWS Elemental Media Store is optimized to deliver performance to meet the unique requirements of high scale, high-quality media workloads (delivering low-latency reads and writes concurrently). This means you can deliver consistent quality-of-service to your viewers, lowering the risk of buffering video and reducing end-to-end latency.
- **Scale with your audience** - AWS Elemental MediaStore scales automatically with the volume of requests you receive. Increases in load don't negatively impact the quality of the viewing experience. This automatic scalability eliminates the expense and complexity of pre-provisioning capacity. If your audience size goes down, there is no need to de-scale - the service automatically sizes to varied capacity while you pay only for what you use.
- **Familiar management tools for access control** - AWS Elemental MediaStore integrates with AWS features for access control, using AWS Identity and Access Management (IAM) policies and roles, with support for resource policies, allowing you to specify granular access controls.

107.2 Backup/Restore and Disaster Recovery

This requirement is not applicable for AWS Elemental MediaStore. For additional information beyond what is described herein, please refer to <https://aws.amazon.com/mediastore/resources/>.

107.3 Pricing Overview

Please visit <https://aws.amazon.com/pricing/> for further details on pricing AWS services.

107.4 Service Constraints

Please see <https://aws.amazon.com/mediastore/resources/> for more information.

107.5 Technical Requirements

- **User Guide** - Describes the components and features that AWS Elemental MediaStore provides and how to use them. [HTML](#) | [PDF](#)
- **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. [HTML](#) | [PDF](#)

108.0 Service Definition – AWS Elemental MediaTailor

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

108.1 Service Overview

AWS Elemental MediaTailor lets video providers insert individually targeted advertising into their video streams without sacrificing broadcast-level quality-of-service. With AWS Elemental MediaTailor, viewers of your live or on-demand video each receive a stream that combines your content with ads personalized to them. But unlike other personalized ad solutions, with AWS Elemental MediaTailor your entire stream – video and ads – is delivered with broadcast-grade video quality to improve the experience for your viewers. AWS Elemental MediaTailor delivers automated reporting based on both client and server-side ad delivery metrics, making it easy to accurately measure ad impressions and viewer behaviour. You can easily monetize unexpected high-demand viewing events with no up-front costs using AWS Elemental MediaTailor. It also improves ad delivery rates, helping you make more money from every video, and it works with a wider variety of content delivery networks, ad decision servers, and client devices.

Top benefits include:

- **Easily deliver targeted ads to any platform** - AWS Elemental MediaTailor makes it easy to personalize ads for users across different devices from a single ad monetization service. Targeting criteria can be set up to fetch personalized ad content from the ad decision server which makes the determination about which ads to send based on viewer, time, and content. Support for popular ad insertion protocols gives you the flexibility to serve individual ads or groups of ads based on your specific business needs.
- **Improve viewing experiences** - Inserting advertising into content streams as viewers play them allows you to match the quality and format of your advertising to your video content stream. Unlike other approaches to ad insertion that can cause ads to play with variable quality, AWS Elemental MediaTailor ensures the video quality of both your ads and of your core content are the same. This gives your viewers an experience that is superior to other streaming video ad solutions.

- **Increase the accuracy of ad view reporting** - AWS Elemental MediaTailor lets you increase the accuracy of ad tracking by capturing data directly from the viewing device, reducing the effects of ad blocking software and adhering to established advertising industry standards. Other approaches often fail to accurately capture viewing behaviour due to ad blocking software, lack of support for certain devices, or dependence on inadequate third-party tools for tracking. With AWS Elemental MediaTailor, ad impressions and viewing behaviour are accurately measured across web, iOS, Android, and other connected viewing devices, helping you more effectively measure the revenue impact of every ad delivered.

108.2 Backup/Restore and Disaster Recovery

This requirement is not applicable for AWS Elemental MediaTailor. For additional information beyond what is described herein, please refer to <https://aws.amazon.com/mediatailor/resources/>.

108.3 Pricing Overview

Please visit <https://aws.amazon.com/pricing/> for further details on pricing AWS services.

108.4 Service Constraints

Please see <https://aws.amazon.com/mediastore/resources/> for more information.

108.5 Technical Requirements

User Guide – Describes the components and features that AWS Elemental MediaTailor provides and how to use them. [HTML](#) | [PDF](#)

109.0 Service Definition – AWS Fargate

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

109.1 Service Overview

AWS Fargate is a technology for Amazon ECS and EKS that allows you to run containers without having to manage servers or clusters. With AWS Fargate, you no longer have to provision, configure, and scale clusters of virtual machines to run containers. This removes the need to choose server types, decide when to scale your clusters, or optimize cluster packing. AWS Fargate removes the need for you to interact with or think about servers or clusters. Fargate lets you focus on designing and building your applications instead of managing the infrastructure that runs them. AWS Fargate will also support Amazon EKS in 2018.

Amazon ECS and EKS have two modes: Fargate launch type and EC2 launch type. With Fargate launch type, all you have to do is package your application in containers, specify the CPU and memory requirements, define networking and IAM policies, and launch the application. EC2 launch type allows you to have server-level, more granular

control over the infrastructure that runs your container applications. With EC2 launch type, you can use Amazon ECS and EKS to manage a cluster of servers and schedule placement of containers on the servers. Amazon ECS and EKS keeps track of all the CPU, memory and other resources in your cluster, and also finds the best server for a container to run on based on your specified resource requirements. You are responsible for provisioning, patching, and scaling clusters of servers. You can decide which type of server to use, which applications and how many containers to run in a cluster to optimize utilization, and when you should add or remove servers from a cluster. EC2 launch type gives you more control of your server clusters and provides a broader range of customization options, which might be required to support some specific applications or possible compliance and government requirements.

Top benefits include:

- **No Clusters to Manage** - With AWS Fargate, you only have to think about the containers so you can just focus on building and operating your application. AWS Fargate eliminates the need to manage a cluster of Amazon EC2 instances. You no longer have to pick the instance types, manage cluster scheduling, or optimize cluster utilization. All of this goes away with Fargate.
- **Seamless Scaling** - AWS Fargate makes it easy to scale your applications. You no longer have to worry about provisioning enough compute resources for your container applications. After you define your application requirements (e.g., CPU, memory, etc.), AWS Fargate manages all the scaling and infrastructure needed to run your containers in a highly-available manner. You no longer have to decide when to scale your clusters or pack them for optimal utilization. With Fargate, you can launch tens or tens of thousands of containers in seconds and easily scale to run your most mission-critical applications.
- **Integrated with Amazon ECS and EKS** - AWS Fargate seamlessly integrates with Amazon ECS. You just define your application as you do for Amazon ECS. You package your application into task definitions, specify the CPU and memory needed, define the networking and IAM policies that each container needs, and upload everything to Amazon ECS. After everything is setup, AWS Fargate launches and manages your containers for you.

109.2 Backup/Restore and Disaster Recovery

This requirement is not applicable for AWS Fargate. For additional information beyond what is described herein, please refer to

https://docs.aws.amazon.com/AmazonECS/latest/developerguide/AWS_Fargate.html

109.3 Pricing Overview

With AWS Fargate, there are no upfront payments and you only pay for the resources that you use. You pay for the amount of vCPU and memory resources consumed by your [containerized](#) applications. For additional information beyond what is described herein, please refer to <https://aws.amazon.com/fargate/pricing/>.

109.4 Service Constraints

AWS Fargate features and default service limits are found at <https://aws.amazon.com/fargate/faqs/>.

109.5 Technical Requirements

AWS Fargate is a technology that you can use with Amazon ECS to run [containers](#) without having to manage servers or clusters of EC2 instances. With AWS Fargate, you no longer have to provision, configure, and scale clusters of virtual machines to run containers. This removes the need to choose server types, decide when to scale your clusters, or optimize cluster packing.

When you run your tasks and services with the Fargate launch type, you package your application in containers, specify the CPU and memory requirements, define networking and IAM policies, and launch the application. For more information, please refer to https://docs.aws.amazon.com/AmazonECS/latest/developerguide/AWS_Fargate.html.

110.0 Service Definition – AWS Firewall Manager

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

110.1 Service Overview

AWS Firewall Manager is a security management service that makes it easier to centrally configure and manage AWS WAF rules across your accounts and applications. Using AWS Firewall Manager, you can easily roll out AWS WAF rules for your Application Load Balancers and Amazon CloudFront distributions across accounts in [AWS Organisations](#).

The top features include:

- **Automatic application of AWS WAF rules** – You can automatically enforce AWS WAF rules on AWS resources that currently exist or are created in the future, thereby ensuring compliance with firewall rules across the organisation.
- **Multi-account resource groups** – Within AWS Firewall Manager, you are able to group resources by account, by resource type, and by tag. Now, your security team can write AWS WAF rules or enable DDoS protection easily for all resources within a particular group or across accounts in the organisation.
- **Cross-account protection policies** – AWS Firewall Manager is integrated with [AWS Organisations](#), so you can apply protections to resources across accounts.
- **Historical 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.
- **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.

Top benefits include:

- **Automatic subscription** – AWS Firewall Manager allows you to subscribe all member accounts in an AWS Organisations organisation to AWS Shield Advanced and automatically subscribes new in-scope accounts that join the organisation.
- **Ensure compliance** – AWS Firewall Manager can ensure existing and newly created resources comply with a mandatory set of security policies automatically. The service can discover new Application Load Balancer or Amazon CloudFront resources as they are created across accounts. For example, if you are required to meet US Department of Treasury's Office of Foreign Assets Control (OFAC) regulations, you can use AWS Firewall Manager to deploy a rule to block traffic from embargoed countries across your Application Load Balancer and Amazon CloudFront accounts.
- **Simplify rule management** – AWS Firewall Manager is integrated with [AWS Organisations](#) so you can enable AWS WAF across multiple AWS accounts and resources from a single place. You can group rules, build policies, and consistently apply those policies across your entire Application Load Balancer and Amazon CloudFront infrastructure. For example, you can delegate the creation of application-specific rules within an account while retaining the ability to enforce global security policies across accounts.
- **Easily deploy Managed Rules for AWS WAF** – AWS Firewall Manager also integrates with [Managed Rules for AWS WAF](#), which gives you an easy way to deploy preconfigured AWS WAF rules in front of your applications. This means you can choose a managed rule provided and updated by an AWS Marketplace seller and deploy it consistently across your Application Load Balancer and Amazon CloudFront infrastructure with just a few clicks in the console. For example, you can easily protect your entire organisation from zero-day vulnerabilities by subscribing to a Managed Rule for AWS WAF from the AWS Marketplace that provides CVE patch updates.
- **Compliance visibility** – With AWS Firewall Manager you can quickly view the compliance status for each policy by looking at how many accounts are included in the scope of the policy and how many of those are compliant. Further, for each policy configured on AWS Firewall Manager, you get a compliance dashboard. The central compliance dashboard allows you to view which accounts are non-compliant to a given policy, which specific resources are non-compliant, and also provides information about the reason why a particular resource is not compliant.
- **Rapid response** – With AWS Firewall Manager your security team can be notified of threats so they can respond and rapidly mitigate an attack. For example, when Amazon GuardDuty discovers a malicious IP address accessing your application, you can quickly deploy a firewall protection policy to apply a

block for the IP address across all your Application Load Balancers and Amazon CloudFront distributions across your entire organisation.

- **Alerting and monitoring** – For each AWS Firewall Manager policy created, you can aggregate CloudWatch metrics for each rule in the rule group, indicating how many requests were allowed or blocked across the entire organisation. This gives you a central place to set up alerts for threats across your organisation. You can create new Amazon SNS notification channels to receive real-time notifications when new non-compliant resources are discovered.

110.2 Backup/Restore and Disaster Recovery

This requirement is not applicable to AWS Firewall Manager. For additional information beyond what is described herein, please refer to <http://aws.amazon.com/documentation/>.

110.3 Pricing Overview

Please visit <https://aws.amazon.com/pricing/> for further details on pricing AWS services.

110.4 Service Constraints

Please see <https://docs.aws.amazon.com/waf/latest/developerguide/fms-limits.html> for more information.

110.5 Technical Requirements

Please see <https://aws.amazon.com/documentation/firewall-manager/> for comprehensive technical documentation regarding AWS Firewall Manager.

111.0 Service Definition – AWS Global Accelerator

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

111.1 Service Overview

AWS Global Accelerator is a networking service that improves the availability and performance of the applications that you offer to your global users. Today, if you deliver applications to your global users over the public internet, your users might face inconsistent availability and performance as they traverse through multiple public networks to reach your application. These public networks can be congested, and each hop can introduce availability and performance risk. AWS Global Accelerator uses the highly available and congestion-free AWS global network to direct internet traffic from your users to your applications on AWS, making your users' experience more consistent.

The top features include:

- **Static anycast IP addresses** – AWS Global Accelerator provides you with static IP addresses that serve as fixed entry points to your applications hosted in any

number of AWS Regions.

- **Fault tolerance using network zones** – AWS Global Accelerator has a fault-isolating design that increases the availability of your applications.
- **Global performance-based routing** – AWS Global Accelerator supports management of network layer traffic at a global level, offering a single front end for both TCP and UDP workloads.
- **DDoS resiliency at the edge** – By default, AWS Global Accelerator is protected by [AWS Shield](#) Standard, which minimises application downtime and latency from denial of service attacks by using always-on network flow monitoring and automated in-line mitigation.
- **Client affinity** – AWS Global Accelerator enables you to build applications where keeping state is essential.

The top benefits include

- **Improve global application availability** – AWS Global Accelerator continually monitors the health of your application endpoints, such as your Network Load Balancers, Application Load Balancers, or Elastic IPs, instantly reacting to changes in their health or configuration. AWS Global Accelerator will then redirect user traffic to healthy endpoints that deliver the best performance and availability to your users.
- **Accelerate your global applications** – AWS Global Accelerator optimises the network path, taking advantage of the vast, congestion-free AWS global network. Regardless of where your users are located, AWS Global Accelerator intelligently routes traffic to the endpoint that provides the best application performance.
- **Easily manage endpoints** – AWS Global Accelerator's static IP addresses make it easy to move endpoints between Availability Zones or AWS Regions without needing to update your DNS configuration or change client-facing applications.
- **Support latency-sensitive applications** – Many applications, such as gaming, media, mobile applications, and financial applications, need very low latency for a great user experience. To improve performance for your users, AWS Global Accelerator directs traffic to edge locations that are nearest to users, routing that traffic to the closest endpoint. AWS Global Accelerator quickly reacts to changes in endpoint performance, improving your users' application performance.
- **Scale your network to support more users** – As your application usage grows, you need to scale and adapt your network. AWS Global Accelerator allows you to scale your network up or down, easily adding or removing endpoints without needing to update the IP addresses in your client applications. AWS Global Accelerator lets you associate regional resources to static IP addresses and whitelist these addresses in your client applications, firewalls, and DNS records.
- **Disaster recovery** – You need to rely on your network to always be available. If

there is an application failure in the primary AWS Region or endpoints, AWS Global Accelerator instantly re-routes traffic to a secondary AWS Region or endpoint.

- **Automatic health checks** – AWS Global Accelerator continuously monitors the health of your application endpoints and routes around failed applications instantly. It supports TCP, HTTP, and HTTP(s) health checks.
- **Fine-grained traffic control** – AWS Global Accelerator gives you the option to dial up or down the traffic to a specific AWS Region using values between zero and one hundred for your endpoint groups. This is especially valuable for testing performance and updates. In the case of a failure, the traffic is reassigned to the other endpoints to maintain high availability. By default, traffic dials are set at 100% across all endpoint groups, letting AWS Global Accelerator select the best endpoint for your applications.
- **Fault tolerance** – AWS Global Accelerator has a fault-isolating design that increases the availability of your applications. Upon provisioning an AWS Global Accelerator resource, you are allocated two IPv4 static addresses that are serviced by independent network zones. Similar to Availability Zones, these network zones are isolated units with their own set of physical infrastructures and service IP addresses from a unique IP subnet. If one IP address from a network zone becomes unavailable due to IP address blocking by certain client networks or network disruptions, then client applications can retry on the healthy static IP address from the other isolated network zone.
- **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 any number of AWS Regions. These IP addresses are anycast from AWS edge locations, meaning that these IP addresses are announced from multiple AWS edge locations at the same time, enabling traffic to ingress onto the AWS global network as close to your users as possible. You can associate these addresses to regional AWS resources or endpoints, such as Network Load Balancers, Application Load Balancers, and Elastic IP addresses. You don't need to make any client-facing changes or update DNS records as you modify or replace endpoints. AWS Global Accelerator's IP addresses are static and serve as the front-end interface of your applications.
- **DDoS protection at the edge** – By default, AWS Global Accelerator is protected by [AWS Shield](#) Standard, which minimises 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 edge for your applications.

111.2 Backup/Restore and Disaster Recovery

This requirement is not applicable to AWS Global Accelerator. For additional information beyond what is described herein, please refer to <http://aws.amazon.com/documentation/>.

111.3 Pricing Overview

Please visit <https://aws.amazon.com/pricing/> for further details on pricing AWS services.

111.4 Service Constraints

Please see <https://docs.aws.amazon.com/global-accelerator/latest/dg/limits-global-accelerator.html> for more information.

111.5 Technical Requirements

Please see https://docs.aws.amazon.com/global-accelerator/?id=docs_gateway for comprehensive technical documentation regarding AWS Global Accelerator.

112.0 Service Definition – AWS Glue

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

112.1 Service Overview

AWS Glue is a fully managed extract, transform, and load (ETL) service that makes it easy for customers to prepare and load their data for analytics. You can create and run an ETL job with a few clicks in the AWS Management Console. You simply point AWS Glue to your data stored in AWS, and AWS Glue discovers your data and stores the associated metadata (e.g., table definition and schema) in the AWS Glue Data Catalog. Once catalogued, your data is immediately searchable, queryable, and available for ETL.

The top features include:

- **Integrated data catalogue** – The AWS Glue Data Catalog is your persistent metadata store for all your data assets, regardless of where they are located.
- **Automatic schema discovery** – AWS Glue crawlers connect to your source or target data store, progress through a prioritised 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 catalogue and used in the authoring process of your ETL jobs.
- **Code generation** – AWS Glue automatically generates the code to ETL your data. Simply point AWS Glue to your data source and target, and AWS Glue creates ETL scripts to transform, flatten, and enrich your data. The code is generated in Scala or Python and written for Apache Spark.

- **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 favourite 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.
- **Flexible job scheduler** – 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.

112.2 Top Benefits

AWS Glue is integrated across a wide range of AWS Cloud services, meaning less hassle for you when on boarding. AWS Glue natively supports data stored in Amazon Aurora and all other Amazon RDS engines, Amazon Redshift, and Amazon S3, as well as common database engines and databases in your Amazon VPC running on Amazon EC2.

AWS Glue is serverless. There is no infrastructure to provision or manage. AWS Glue handles provisioning, configuration, and scaling of the resources required to run your ETL jobs on a fully managed, scale-out Apache Spark environment. You pay only for the resources used while your jobs are running.

AWS Glue automates much of the effort in building, maintaining, and running ETL jobs. AWS Glue crawls your data sources, identifies data formats, and suggests schemas and transformations. AWS Glue automatically generates the code to execute your data transformations and loading processes.

112.3 Backup/Restore and Disaster Recovery

AWS Glue is a managed service, meaning that Amazon manages any backups required to support the service.

112.4 Pricing Overview

Please visit <https://aws.amazon.com/pricing/> for further details on pricing AWS services.

112.5 Service Constraints

Please see <https://docs.aws.amazon.com/glue/latest/dg/troubleshooting-service-limits.html> for more information.

112.6 Technical Requirements

Please refer to <https://docs.aws.amazon.com/glue/latest/dg/what-is-glue.html> for comprehensive technical documentation regarding AWS Glue.

113.0 Service Definition – AWS Greengrass

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

113.1 Service Overview

AWS Greengrass is software that lets you run local compute, messaging, data caching, sync, and ML inference capabilities for connected devices in a secure way. With AWS Greengrass, connected devices can run [AWS Lambda](#) functions, keep device data in sync, and communicate with other devices securely – even when not connected to the Internet. Using AWS Lambda, Greengrass ensures your IoT devices can respond quickly to local events, use Lambda functions running on Greengrass Core to interact with local resources, operate with intermittent connections, stay updated with over the air updates, and minimize the cost of transmitting IoT data to the cloud.

[ML Inference](#) is a feature of AWS Greengrass that makes it easy to perform machine learning inference locally on Greengrass Core devices using models that are built and trained in the cloud.

AWS Greengrass seamlessly extends AWS to devices so they can act locally on the data they generate, while still using the cloud for management, analytics, and durable storage. With Greengrass, you can use familiar languages and programming models to create and test your device software in the cloud, and then deploy it to your devices. AWS Greengrass can be programmed to filter device data and only transmit necessary information back to the cloud. AWS Greengrass authenticates and encrypts device data at all points of connection using the security and access management capabilities of [AWS IoT Core](#). This way, data is never exchanged between devices when they communicate with each other and the cloud, without proven identity.

Top benefits include:

- **Respond to Local Events in Near Real-time** - AWS Greengrass devices can act locally on the data they generate, while still using the cloud for management, analytics, and durable storage.
- **Operate Offline** - AWS Greengrass lets connected devices operate even with intermittent connectivity to the cloud. Once the device reconnects, Greengrass synchronizes the data on the device with [AWS IoT](#), providing seamless functionality regardless of connectivity.
- **Simplified Device Programming with AWS Lambda** - AWS Greengrass uses the same AWS Lambda programming models you use in the cloud so you can create and test your device software in the cloud first, and then deploy it seamlessly to your devices. Greengrass lets you execute Lambda functions locally, reducing the complexity of developing embedded software.
- **Reduce the Cost of Running IoT Applications** - With AWS Greengrass you can program the device to filter device data locally and only transmit the data you need for your applications to cloud. This reduces the amount of raw data

transmitted to the cloud and lowers cost, and increases the quality of the data you send to the cloud so you can achieve rich insight at a lower cost.

113.2 Backup/Restore and Disaster Recovery

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

113.3 Pricing Overview

Please visit <https://aws.amazon.com/pricing/> for further details on pricing AWS services.

113.4 Service Constraints

Please see <https://aws.amazon.com/greengrass/> for more information.

113.5 Technical Requirements

AWS Greengrass seamlessly extends AWS onto physical devices so they can act locally on the data they generate, while still using the cloud for management, analytics, and durable storage. AWS Greengrass ensures your devices can respond quickly to local events and operate with intermittent connectivity. AWS Greengrass minimizes the cost of transmitting data to the cloud by allowing you to author AWS Lambda functions that execute locally.

- **Developer Guide** - Walks through how to set up AWS Greengrass and integrate it with other services. [HTML](#) | [PDF](#)
- **AWS Greengrass sections of the AWS CLI Reference** - Describes the AWS CLI commands that you can use to administer and manipulate AWS Greengrass. [Greengrass](#)
- **API Reference** - Describes all the API operations for AWS Greengrass in detail. Also provides sample requests, responses, and errors for the supported web services protocols. [HTML](#) | [PDF](#)

114.0 Service Definition – AWS Identity and Access Management (IAM)

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

114.1 Service Overview

AWS IAM enables you to securely control access to AWS Cloud services and resources for your users. Using AWS IAM, you can create and manage AWS users and groups and use permissions to allow and deny their access to AWS resources.

Top features include:

- **Manage AWS IAM Users and their Access** – You can create users in AWS IAM, assign them individual security credentials (e.g., access keys, passwords, multi-factor authentication devices) or request temporary security credentials to provide users access to AWS Cloud services and resources. You can manage permissions in order to control which operations a user can perform.
- **Manage AWS IAM Roles and their Permissions** – You can create roles in AWS IAM and manage permissions to control which operations can be performed by the entity or AWS Cloud service that assumes the role. You can also define which entity is allowed to assume the role.
- **Manage Federated Users and their Permissions** – You can enable identity federation to allow existing identities in your enterprise to access the AWS Management Console, to call AWS APIs, and to access resources without the need to create an AWS IAM user for each identity.
- **Resource-Based Policies** - Resource-based policies are attached to a resource. For example, you can attach resource-based policies to Amazon S3 buckets, Amazon SQS queues, and AWS Key Management Service encryption keys.
- **Permissions Boundaries for IAM Entities** - AWS supports permissions boundaries for IAM entities (users or roles). A permissions boundary is an advanced feature for using a managed policy to set the maximum permissions that an identity-based policy can grant to an IAM entity. An entity's permissions boundary allows it to perform only the actions that are allowed by both its identity-based policies and its permissions boundaries.
- **IAM Access Analyzer** - AWS IAM Access Analyzer helps you identify the resources in your account, such as Amazon S3 buckets or IAM roles, that are shared with an external entity. This lets you identify unintended access to your resources and data, which is a security risk. Access Analyzer identifies resources that are shared with external principals by using logic-based reasoning to analyse the resource-based policies in your AWS environment.

114.2 Backup/Restore and Disaster Recovery

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

114.3 Pricing Overview

This service is not discussed in the AWS UK G-Cloud 12 Pricing Document as it is available free of charge. AWS IAM is a feature of your AWS account offered at no additional charge.

114.4 Service Constraints

Please see <http://aws.amazon.com/documentation/iam/> for more information.

114.5 Technical Requirements

Please refer to <http://aws.amazon.com/documentation/iam> and the following links for comprehensive technical documentation regarding AWS IAM.

- **Using AWS IAM** – Introduces you to AWS IAM, helps you set up an account, and walks you through a simple example to help you use AWS IAM for the first time. Also provides tips and links to advanced product features and resources. Available in [HTML](#), [PDF](#), and [Kindle](#) formats.
- **AWS IAM Section of AWS CLI Reference** – Describes the AWS CLI commands that you can use to administer AWS IAM. Provides syntax, options, and usage examples for each command. Available in [HTML](#) format.
- **AWS IAM API Reference** – Describes all the API operations for AWS IAM in detail. Also provides sample requests, responses, and errors for the supported web services protocols. Available in [HTML](#) and [PDF](#) formats.
- **Security Token Service (STS) Section of AWS CLI Reference** – Describes the AWS CLI commands that you can use to generate temporary security credentials. Provides syntax, options, and usage examples for each command. Available in [HTML](#) format.
- **STS API Reference** – Describes all the API operations for AWS STS in detail. Also provides sample requests, responses, and errors for the supported web services protocols. Available in [HTML](#) and [PDF](#) formats.

115.0 Service Definition – AWS IoT

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

This service listing includes the following AWS services:

- AWS IoT 1-Click
- AWS IoT Analytics
- AWS IoT Core
- AWS IoT Device Defender
- AWS IoT Device Management
- AWS IoT GreenGrass
- AWS IoT Events
- AWS IoT SiteWise (preview)

- AWS IoT Things Graph

115.1 Service Overview

AWS IoT is a managed cloud platform that lets connected devices easily and securely interact with cloud applications and other devices. AWS IoT 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, your applications can keep track of and communicate with all your devices, all the time, even when they aren't connected.

AWS IoT makes it easy to use AWS Cloud services like AWS Lambda, Amazon Kinesis, Amazon S3, Amazon Machine Learning, Amazon DynamoDB, Amazon CloudWatch, AWS CloudTrail, and Amazon Elasticsearch Service with built-in Kibana integration to build IoT applications that gather, process, analyse and act on data generated by connected devices, without having to manage any infrastructure.

Top benefits include:

- **Connect and Manage Your Devices** – AWS IoT allows you to easily connect devices to the cloud and to other devices. AWS IoT supports HTTP, WebSockets, and MQTT, a lightweight communication protocol specifically designed to tolerate intermittent connections, minimise the code footprint on devices, and reduce network bandwidth requirements. AWS IoT also supports other industry-standard and custom protocols, and devices can communicate with each other even if they are using different protocols.
- **Secure Device Connections and Data** – AWS IoT provides authentication and end-to-end encryption throughout all points of connection, so that data is never exchanged between devices and AWS IoT without proven identity. In addition, you can secure access to your devices and applications by applying policies with granular permissions.
- **Process and Act Upon Device Data** – With AWS IoT, you can filter, transform, and act upon device data on the fly, based on business rules you define. You can update your rules to implement new device and application features at any time. AWS IoT makes it easy to use AWS Cloud services like AWS Lambda, Amazon Kinesis, Amazon S3, Amazon Machine Learning, Amazon DynamoDB, Amazon CloudWatch, and Amazon Elasticsearch Service for even more powerful IoT applications.
- **Read and Set Device State at Any Time** – AWS IoT stores the latest state of a device so that it can be read or set at any time, making the device appear to your applications as if it were online all the time. This means that your application can read a device's state even when it is disconnected and also allows you to set a device state and have it implemented when the device reconnects.

115.2 Backup/Restore and Disaster Recovery

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

115.3 Pricing Overview

Please visit <https://aws.amazon.com/pricing/> for further details on pricing AWS services.

115.4 Service Constraints

Please see <https://aws.amazon.com/iot/> for more information.

115.5 Technical Requirements

AWS IoT enables secure, bi-directional communication between Internet-connected things (such as sensors, actuators, embedded devices, or smart appliances) and the AWS Cloud over MQTT and HTTP.

- **Developer Guide** – Walks through how to set up AWS IoT and integrate it with other services. [HTML](#) | [PDF](#)
- **API Reference** – Describes all the API operations for AWS IoT in detail. Also provides sample requests, responses, and errors for the supported web services protocols. [HTML](#) | [PDF](#)
- **AWS IoT sections of the AWS CLI Reference** – Describes the AWS CLI commands that you can use to administer and manipulate AWS IoT. [IoT](#) | [IoT-Data](#)
- **Thing Shadow REST API** – Describes all of the REST API operations for working with Thing Shadows. [HTML](#)

115.6 Service Overview – AWS IoT Analytics

AWS IoT Analytics is a fully-managed service that makes it easy to run sophisticated analytics on massive volumes of IoT data without having to worry about all the cost and complexity typically required to build your own IoT analytics platform. It is the easiest way to run analytics on IoT data and get insights to make better and more accurate decisions for IoT applications and machine learning use cases.

IoT data is highly unstructured which makes it difficult to analyse with traditional analytics and business intelligence tools that are designed to process structured data. IoT data comes from devices that often record fairly noisy processes (such as temperature, motion, or sound), and as a result the data from these devices can frequently have significant gaps, corrupted messages, and false readings that must be cleaned up before analysis can occur. Also, IoT data is often only meaningful in the context of other data from external sources. For example, to determine when to water their crops, vineyard irrigation systems often enrich humidity sensor data with data on rainfall at the vineyard, allowing them to be more efficient with water usage while maximizing their harvest yield.

AWS IoT Analytics automates each of the difficult steps that are required to analyse data from IoT devices. IoT Analytics filters, transforms, and enriches IoT data before storing it in a time-series data store for analysis. You can setup the service to collect only the data you need from your devices, apply mathematical transforms to process the data, and enrich the data with device-specific metadata such as device type and location before storing the processed data. Then, you can analyse your data by running ad hoc or scheduled queries using the built-in SQL query engine, or perform more complex analytics and machine learning inference. IoT Analytics makes it easy to get started with machine learning by including pre-built models for common IoT use cases so you can quickly answer questions like which devices are about to fail or which customers are at risk of abandoning their wearable devices.

AWS IoT Analytics is fully managed and scales automatically to support up to petabytes of IoT data. With IoT Analytics, you can analyse data from millions of devices and build fast, responsive IoT applications without managing hardware or infrastructure.

For more information, visit the [AWS IoT Analytics documentation](#) page.

Top benefits include:

- **Easily Run Queries on IoT Data** - With AWS IoT Analytics, you can run simple, ad-hoc queries using the built-in IoT Analytics SQL query engine. The service allows you to use standard SQL queries to extract data from the data store to answer questions like the average distance travelled for a fleet of connected vehicles or how many doors are locked after 7pm in a smart building. These queries can be re-used even if connected devices, fleet size, and analytic requirements change.
- **Run Time-Series Analytics** - AWS IoT Analytics also supports time-series analyses so you can analyse the performance of devices over time and understand how and where they are being used, continuously monitor device data to predict maintenance issues, and monitor sensors to predict and react to environmental conditions.
- **Data Storage Optimized for IoT** - AWS IoT Analytics stores the processed device data in a time-series data store that is optimized to deliver fast response times on IoT queries that typically include time as a criteria. The raw data is also automatically stored for later processing or to reprocess it for another use case.
- **Prepares Your IoT Data for Analysis** - AWS IoT Analytics includes data preparation techniques that make it easy to prepare and process your data for analysis. IoT Analytics is integrated with AWS IoT Core so it is easy to ingest device data directly from connected devices. It cleans false readings, fills gaps in the data, and performs mathematical transformations of message data. As the data is ingested, IoT Analytics can process it using conditional statements, filter data to collect just the data you want to analyse, and enrich it with information from the AWS IoT registry. You can also use AWS Lambda functions to enrich your device data from external sources like the Weather Service, HERE Maps, Salesforce, or Amazon DynamoDB. For example, you could combine weather

data and mapping information to create better information about a device's environment.

- **Tools for Machine Learning** - AWS IoT Analytics makes it easy to apply machine learning to your IoT data with hosted Jupyter notebooks. You can directly connect your IoT data to the notebook and build, train, and execute models right from the IoT Analytics console without having to manage any of the underlying infrastructure. Using AWS IoT Analytics, you can apply machine learning algorithms to your device data to produce a health score for each device in your fleet. For example, an auto manufacturer can detect which of their customers have worn brake pads and alert them to seek maintenance for their vehicles.
- **Automated Scaling with Pay-As-You Go Pricing** - AWS IoT Analytics is a fully managed and pay-as-you go service that scales automatically to support up to petabytes of IoT data. With IoT Analytics, you can analyse your entire fleet of connected devices without managing hardware or infrastructure. As your needs change, compute power and the data store automatically scale up or down so you always have the right capacity for your IoT applications and you only pay for the resources that you use.

115.6.1 Backup/Restore and Disaster Recovery

This requirement is not applicable for AWS IoT Analytics. For additional information beyond what is described herein, please refer to <http://aws.amazon.com/documentation/>.

115.6.2 Pricing Overview

Please visit <https://aws.amazon.com/pricing/> for further details on pricing AWS services.

115.6.3 Service Constraints

Please see <https://aws.amazon.com/documentation/iotanalytics/> for more information.

115.6.4 Technical Requirements

AWS IoT Analytics provides advanced data analysis for AWS IoT. You can collect large amounts of device data, process messages, and store them. You can then query the data and run sophisticated analytics to make accurate decisions in your IoT applications and machine learning use cases.

- **User Guide** - Walks through how to set up the service and integrate with other AWS services. [HTML](#) | [PDF](#)
- **User Guide** - Describes all the API operations for AWS IoT Analytics in detail. Also provides sample requests, responses, and errors for the supported web services protocols. [HTML](#) | [PDF](#)

115.7 Service Overview – AWS IoT Button

The AWS IoT Button is a programmable button based on the Amazon Dash Button hardware. This simple Wi-Fi device is easy to configure and designed for developers to get started with [AWS IoT Core](#), [AWS Lambda](#), [Amazon DynamoDB](#), [Amazon SNS](#), and many other Amazon Web Services without writing device-specific code.

You can code the button's logic in the cloud to configure button clicks to count or track items, call or alert someone, start or stop something, order services, or even provide feedback. For example, you can click the button to unlock or start a car, open your garage door, call a cab, call your spouse or a customer service representative, track the use of common household chores, medications or products, or remotely control your home appliances.

The button can be used as a remote control for Netflix, a switch for your Philips Hue light bulb, a check-in/check-out device for Airbnb guests, or a way to order your favourite pizza for delivery. You can integrate it with third-party APIs like Twitter, Facebook, Twilio, Slack or even your own company's applications. Connect it to things we haven't even thought of yet. We can't wait to see what you will build with the AWS IoT Button!

Top benefits include:

- Easy to configure
- Designed for developers
- Get started with AWS IoT Core, AWS Lambda, Amazon DynamoDB

For more information, please refer to <https://aws.amazon.com/iotbutton/>.

115.8 Service Overview – AWS IoT Core

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.

AWS IoT Core makes it easy to use AWS services like [AWS Lambda](#), [Amazon Kinesis](#), [Amazon S3](#), [Amazon Machine Learning](#), [Amazon DynamoDB](#), [Amazon CloudWatch](#), [AWS CloudTrail](#), and [Amazon Elasticsearch Service](#) with built-in [Kibana](#) integration, to build IoT applications that gather, process, analyse and act on data generated by connected devices, without having to manage any infrastructure.

Top benefits include:

- **Connect and manage your devices** - AWS IoT Core allows you to easily connect devices to the cloud and to other devices. AWS IoT Core supports HTTP, WebSockets, and MQTT, a lightweight communication protocol specifically designed to tolerate intermittent connections, minimize the code

footprint on devices, and reduce network bandwidth requirements. AWS IoT Core also supports other industry-standard and custom protocols, and devices can communicate with each other even if they are using different protocols.

- **Process and act upon device data** - With AWS IoT Core, you can filter, transform, and act upon device data on the fly, based on business rules you define. You can update your rules to implement new device and application features at any time. AWS IoT Core makes it easy to use AWS services like AWS Lambda, Amazon Kinesis, Amazon S3, Amazon Machine Learning, Amazon DynamoDB, Amazon CloudWatch, and Amazon Elasticsearch Service for even more powerful IoT applications.
- **Read and set device state at any time** - AWS IoT Core stores the latest state of a device so that it can be read or set at any time, making the device appear to your applications as if it were online all the time. This means that your application can read a device's state even when it is disconnected, and also allows you to set a device state and have it implemented when the device reconnects.

115.8.1 Backup/Restore and Disaster Recovery

This requirement is not applicable for AWS IoT Core. For additional information beyond what is described herein, please refer to <http://aws.amazon.com/documentation/>.

115.8.2 Pricing Overview

Please visit <https://aws.amazon.com/pricing/> for further details on pricing AWS services.

115.8.3 Service Constraints

Please see <https://aws.amazon.com/iot-core/> for more information.

115.8.4 Technical Requirements

AWS IoT Core is a platform that enables you to connect devices to AWS Services and other devices, secure data and interactions, process and act upon device data, and enable applications to interact with devices even when they are offline.

- **Developer Guide** - [PDF Version or HTML Version](#)
- **API References** - [PDF Version or HTML Version](#)

115.9 Service Overview – AWS IoT Device Management

AWS IoT Device Management makes it easy to securely onboard, organise, monitor, and remotely manage IoT devices at scale. IoT Device Management lets you register your devices individually or in bulk, and manage permissions so that devices remain secure. Then, you use the IoT Device Management console to organise your devices into groups, monitor and troubleshoot device functionality, and send remote updates to your devices. AWS IoT Device Management allows you to scale your device fleets and reduce the cost and effort of managing large IoT device deployments.

Top benefits include:

- **Fast Device On boarding** - With AWS IoT Device Management, you can securely add device attributes like device name, type and manufacturing year, certificates and access policies to the AWS IoT Registry in bulk, assign them to devices, and put large fleets of devices into service quickly.
- **Simple Device Organisation** - AWS IoT Device Management lets you organise your devices into groups and manage access policies for these groups. This makes it easy to track, operate, and manage your devices according to business and security requirements, such as deploying a firmware update for all devices in a building or defining how devices communicate with each other. You can create a hierarchy for your groups such as grouping multiple sensors within a single vehicle and grouping multiple vehicles in a fleet. Then, your devices will inherit access policies based on the group hierarchy.
- **Locate Devices Quickly** - AWS IoT Device Management lets you quickly search and find any device across your entire device fleet in near real-time. You can easily find devices based on a combination of attributes like device ID, device state and type, and quickly find specific devices so that you can take action or troubleshoot your devices.
- **Easy Remote Management** - AWS IoT Device Management makes it easy for you to maintain the health of your device fleet. With IoT Device Management, you can remotely update the software running on your devices after they have been deployed in the field – allowing you to ensure that devices are always running on the latest software. You can also remotely execute actions on the devices such as reboots, factory resets, software updates, and security patches..

115.9.1 Backup/Restore and Disaster Recovery

This requirement is not applicable for AWS IoT Device Management. For additional information beyond what is described herein, please refer to <http://aws.amazon.com/documentation/>.

115.9.2 Pricing Overview

Please visit <https://aws.amazon.com/pricing/> for further details on pricing AWS services.

115.9.3 Service Constraints

Please see <https://aws.amazon.com/documentation/iot-device-management/> for more information.

115.9.4 Technical Requirements

AWS IoT Device Management is a cloud-based device management service that makes it easy for customers to securely manage IoT devices throughout their lifecycle. Customers can use AWS IoT Device Management to on board device information and configuration, organise their device inventory, monitor their fleet of devices, and remotely manage devices deployed across many locations. This remote management includes over-the-air (OTA) updates to device software.

- **User Guide** - Describes device provisioning, thing groups, and jobs that are sent to and executed on one or more devices connected to AWS IoT Core. [HTML](#) | [PDF](#)
- **API Reference** - Describes all the API operations for AWS IoT Device Management in detail. Provides examples of request and response syntax. [HTML](#) | [PDF](#)

116.0 Service Definition – AWS Key Management Service (AWS KMS)

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

116.1 Service Overview

AWS KMS is a managed service that makes it easy for you to create and control the encryption keys used to encrypt your data, and uses HSMs to protect the security of your keys. AWS Key Management Service is integrated with other AWS Cloud services including Amazon EBS, Amazon S3, Amazon RDS, and Amazon Redshift. AWS KMS is also integrated with AWS CloudTrail to provide you with logs of all key usage to help meet your regulatory and compliance needs.

Top benefits include:

- **Centralised Key Management** – AWS Key Management Service provides you with centralised control of your encryption keys. AWS KMS presents a single view into all of the key usage in your organisation. You can easily create keys, implement key rotation, create usage policies, and enable logging from the AWS Management Console, or by using the API.
- **Integrated with AWS Cloud Services** – AWS KMS is integrated with Amazon S3, Amazon EBS, Amazon Redshift, Amazon RDS, and Amazon EMR to make it easy to encrypt the data you store with these services using keys that you manage.
- **Encryption for all your applications** – AWS KMS makes it easy to manage encryption keys used to encrypt data stored by your applications regardless of where you store it. AWS KMS provides an SDK for programmatic integration of encryption and key management into your applications.
- **Built-in Auditing** – AWS KMS works with AWS CloudTrail to provide you with logs of API calls made to or by AWS KMS. These logs help you meet compliance and regulatory requirements by providing details of when keys were accessed and who accessed them.
- **Fully Managed** – AWS KMS is a fully managed service, so you can focus on the encryption needs of your applications while AWS handles availability, physical security, and hardware maintenance of the underlying infrastructure.

- **Low-cost** – There is no charge for the storage of default keys in your account. You pay only for additional master keys that you create and your key usage.
- **Secure** – AWS KMS provides you a secure location to store and use encryption keys, using hardened systems where your unencrypted keys are only used in memory. AWS KMS keys are never transmitted outside of the AWS Regions 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.
- **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.

116.2 Backup/Restore and Disaster Recovery

AWS KMS is a managed service. As your usage of AWS KMS encryption keys grows, you do not have to buy additional key management hardware or software or manage any infrastructure. AWS KMS automatically scales to meet your encryption key needs.

Key storage is highly durable. AWS KMS stores multiple copies of encrypted versions of your keys in systems that are designed for 99.999999999% durability to help assure you that your keys will be available when you need to access them.

AWS KMS is deployed in multiple Availability Zones within an AWS Region to provide high availability for your encryption keys.

116.3 Pricing Overview

Please visit <https://aws.amazon.com/pricing/> for further details on pricing AWS services.

116.4 Service Constraints

Please see <https://aws.amazon.com/kms/> for more information.

116.5 Technical Requirements

AWS KMS is an encryption and key management service scaled for the cloud. KMS keys and functionality are used by other AWS Cloud services, and you can use them to protect data in your own applications that use AWS.

- **Developer Guide** – Provides conceptual overviews of AWS KMS and explains how to use it to protect data in your own applications that use AWS. [HTML](#) | [PDF](#)
- **API Reference** – Describes all the API operations for AWS KMS in detail. Also provides sample requests, responses, and errors for the supported web services protocols. [HTML](#) | [PDF](#)

117.0 Service Definition – AWS Lake Formation

AWS Lake Formation is a service that makes it easy to set up a secure data lake in days. A data lake is a centralized, curated, and secured repository that stores all your data, both in its original form and prepared for analysis. A data lake enables you to break down data silos and combine different types of analytics to gain insights and guide better business decisions.

117.1 Service Overview

Top features include:

- Import data from MySQL, PostgreSQL, SQL Server, MariaDB and Oracle databases running in Amazon RDS, Amazon EC2, from on-premise using a JDBC connector, or from Amazon S3.
- Automatically crawl and read data sources to extract metadata and schema definitions to create a searchable catalog of data that your users can search to discover available data sets. Add your own custom labels to provide additional context, and secure data based on these labels.
- Perform data transformations, such as rewriting of mismatched date formats, for consistency. Transform using common standards such as Parquet and ORC using AWS Glue, or build a transformation in Apache Spark.
- De-duplicate data using machine learning transformations, such as FindMatches, without knowing anything about machine learning.
- Optimise partitioning of data in Amazon S3 to improve performance and reduce cost. Raw data that is read in from partitions that are too small requires more reads, whereas additional data is read in when partitions are too large. Lake Formation organises data by size, time period or relevant keys, enabling both fast scans and parallel, distributed reads.
-

Top benefits include:

- Build data lakes quickly by moving, storing, centrally cataloguing and cleaning data faster. Point Lake Formation at data sources, and have it crawl, move and organise data appropriately.
- Simplify security management by defining security, governance and audit policies in one, central place, without the need to configure policies in IAM, S3, KMS, Redshift, Athena and EMR. Lake Formation will ensure these policies are consistently implemented.
- Provide self-service access to data by providing data catalogues with information visible within only accessible to those with access to the data.

117.2 Backup/Restore and Disaster Recovery

AWS Lake Formation relies on a number of AWS products in order to function, including Amazon S3, Amazon Glue and Amazon Athena. Back-up, restoration and disaster

recovery should be taken in the context of Amazon S3. For further information, please contact an AWS solutions architect.

117.3 Pricing Overview

Please see the AWS UK G Cloud 12 Pricing Document affiliated with this service in the Digital Marketplace.

117.4 Service Constraints

Please see <https://docs.aws.amazon.com/lake-formation/latest/dg/limitations.html> for more information regarding the constraints in place when using AWS Lake Formation.

117.5 Technical Requirements

Please refer to <https://docs.aws.amazon.com/lake-formation/> and the following links for comprehensive technical documentation regarding Amazon Lake formation:

- <https://docs.aws.amazon.com/lake-formation/latest/dg/what-is-lake-formation.html>
- <https://docs.aws.amazon.com/lake-formation/latest/dg/how-it-works>
- <https://docs.aws.amazon.com/lake-formation/latest/dg/getting-started-setup>
- <https://docs.aws.amazon.com/lake-formation/latest/dg/getting-started>

118.0 Service Definition – AWS Lambda

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

118.1 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 AWS Lambda, you can run code for virtually any type of application or backend service, all with zero administration. Just upload your code and AWS 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 Cloud services or call it directly from any web or mobile application.

Top benefits include:

- **No Servers to Manage** – AWS Lambda automatically runs your code without requiring you to provision or manage servers. Just write the code and upload it to AWS Lambda.
- **Continuous Scaling** – AWS Lambda automatically scales your application by running code in response to each trigger. Your code runs in parallel and processes each trigger individually, scaling precisely with the size of the workload.

- **Sub second Metering** – With AWS Lambda, you are charged for every 100ms your code executes and the number of times your code is triggered. You don't pay anything when your code isn't running.

118.2 Backup/Restore and Disaster Recovery

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

118.3 Pricing Overview

Please visit <https://aws.amazon.com/pricing/> for further details on pricing AWS services.

118.4 Service Constraints

Please see <https://aws.amazon.com/lambda/> for more information.

118.5 Technical Requirements

AWS Lambda is a zero-administration compute platform for back-end web developers that runs your code for you in the AWS Cloud and provides you with a fine-grained pricing structure. AWS Lambda runs your back-end code on its own AWS compute fleet of Amazon EC2 instances across multiple Availability Zones in a region, which provides the high availability, security, performance, and scalability of the AWS infrastructure.

- **Developer Guide** – Provides a conceptual overview of AWS Lambda, includes detailed instructions for using the various features, and provides a complete API reference for developers. [HTML](#) | [PDF](#)

119.0 Service Definition – AWS License Manager

AWS License Manager makes it easier to manage your software licenses from software vendors such as Microsoft, SAP, Oracle, and IBM across AWS and on-premises environments. AWS License Manager lets administrators create customized licensing rules that emulate the terms of their licensing agreements, and then enforces these rules when an instance of EC2 gets launched. Administrators can use these rules to help prevent licensing violations, such as using more licenses than an agreement stipulates. The rules in AWS License Manager enable you to help prevent a licensing breach by stopping the instance from launching or by notifying administrators about the infringement. Administrators gain control and visibility of all their licenses with the AWS License Manager dashboard and reduce the risk of non-compliance, misreporting, and additional costs due to licensing overages.

119.1 Service Overview

Top Features include

- Enables a centralized team in your organization to interpret software licensing agreements and create rules to govern license usage across AWS and on-premises based on your licensing rules.

- Set hard or soft limits to control license usage and prevent the launch of a new, non-compliant instance.
- Seamlessly track license usage throughout the lifecycle of your AWS resources.
- Integrated with Amazon EC2, AWS Systems Manager, AWS Organizations, and AWS Service Catalog, and AWS Marketplace.
- 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.

Top benefits include

- **Gain control over license usage**-The way organizations manage licenses can vary from simple spreadsheets to highly customized solutions. Often, these approaches can be problematic, as they require manual and ad-hoc reporting that can be inaccurate and quickly outdated. With AWS License Manager, license administrators can create custom licensing rules to help manage license usage centrally. AWS License Manager gives organizations 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. 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. This helps you reduce the risk of non-compliance.
- **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. Rules can be attached whenever an instance gets created using golden AMIs, Amazon EC2 launch templates, AWS CloudFormation templates (<https://aws.amazon.com/cloudformation/aws-cloudformation-templates/>), or AWS Service Catalog (<https://aws.amazon.com/servicecatalog/>). 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, that BYOL license is released and is available for re-use.
- **Limit non-compliance proactively**- Set hard or soft limits to control license usage and prevent the launch of a new, non-compliant instance. These limits get evaluated 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 (<https://aws.amazon.com/sns/>). Notifications are sent as emails, text messages or alerts to inform them 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 discover software running on existing EC2 instances using AWS Systems Manager (<https://aws.amazon.com/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. Additionally, administrators can discover software usage on instances using AWS Organizations (<https://aws.amazon.com/organizations/>) by going through a one-time multi-account set-up and creating policies that centrally control AWS service use across multiple AWS accounts.
- Centralize license management and reporting- Get a centralized view of license usage across AWS and on-premises 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 your eligible licenses that require Dedicated Hosts, AWS License Manager allows you to specify Dedicated Host management preferences for host allocation and host capacity utilization. AWS then takes care of these tasks on your behalf, so that you can seamlessly launch instances without performing upfront host allocation or managing capacity utilization.
- Built-in AWS integration- Seamlessly track license usage throughout the lifecycle of your AWS resources. AWS License Manager is integrated with Amazon EC2 (<https://aws.amazon.com/ec2/>), AWS Systems Manager, AWS Organizations, and AWS Service Catalog, and AWS Marketplace (<https://aws.amazon.com/marketplace>). License administrators can add rules in AWS Service Catalog, which allows them to create and manage catalogues of IT services that are approved for use on all their AWS accounts. Through seamless integration with AWS Systems Manager and AWS Organizations, administrators can manage licenses across all the AWS accounts in an organization and on-premises environments. AWS Marketplace buyers can also use AWS License Manager to track bring your own license (BYOL) software obtained from the Marketplace and keep a consolidated view of all their licenses.

119.2 Backup/Restore and Disaster Recovery

For additional information beyond what is described herein, please refer to <http://aws.amazon.com/documentation>.

- <https://docs.aws.amazon.com/license-manager/latest/userguide/disaster-recovery-resiliency.html>
- (<https://docs.aws.amazon.com/license-manager/latest/userguide/disaster-recovery-resiliency.html>)

119.3 Pricing Overview

Please see the AWS UK G Cloud 12 Pricing Document affiliated with this service in the Digital Marketplace.

119.4 Service Constraints

Please see the following links for information:

- <https://aws.amazon.com/license-manager/faqs>
- <https://aws.amazon.com/license-manager/faqs>
- <https://docs.aws.amazon.com/license-manager/latest/userguide/license-manager.html>
- <https://docs.aws.amazon.com/license-manager/latest/userguide/license-manager.html>

119.5 Technical Requirements

Please refer to <https://aws.amazon.com/license-manager> and the following links for comprehensive technical documentation regarding AWS License Manager.

- <https://docs.aws.amazon.com/license-manager/latest/userguide/license-manager.html>
- <https://docs.aws.amazon.com/license-manager/latest/APIReference/Welcome.html>

120.0 Service Definition – AWS Migration Hub

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

120.1 Service Overview

AWS Migration Hub provides a single location to track the progress of application migrations across multiple AWS and partner solutions. Migration Hub allows you to choose the AWS and partner migration tools that best fit your needs while providing visibility into the status of migrations across your portfolio of applications. Migration Hub also provides key metrics and progress for individual applications, regardless of which tools are being used to migrate them. For example, you might use AWS DMS, AWS SMS, and partner migration tools such as ATADATA ATAmotion, CloudEndure Live Migration, or RiverMeadow Server Migration SaaS to migrate an application comprised of a database, virtualised web servers, and a bare metal server. Using Migration Hub,

you can view the migration progress of all the resources in the application. This allows you to quickly get progress updates across all of your migrations, easily identify and troubleshoot any issues, and reduce the overall time and effort spent on your migration projects.

Top Features include:

- **Import your on-premises server details** – AWS Migration Hub import allows you to import information about on-premises servers and applications into the Migration Hub so you can track the status of application migrations.
- **Simple and intuitive migration dashboard** – The AWS Migration Hub dashboard shows the latest status and metrics for your entire migration portfolio. This allows you to quickly understand the progress of your migrations as well as identify and troubleshoot any issues that arise.
- **Application migration tracking** – AWS Migration Hub provides all application details in a central location. This allows you to track the status of all the moving parts across all migrations, making it easier to view overall migration progress and reducing the time spent determining current status and next steps.
- **Migration tool integration** – AWS Migration Hub provides the flexibility to use the migration tools that work best for your organisation. Whether you use AWS migration tools like AWS SMS and AWS DMS or partner tools like ATADATA, CloudEndure, and RiverMeadow, Migration Hub makes it easy for you to track migrations from all of those tools in a central location.
- **Multi-Region migrations** – AWS Migration Hub lets you track the status of your migrations into any AWS Region supported by your migration tools. Regardless of which Regions you migrate into, the migration status will appear in Migration Hub when using an integrated tool.

Top Benefits include:

- **Centralised tracking** – Migrations involve many components. AWS Migration Hub helps address this by providing a central location to track the status of all these components, making it easier to view overall migration progress and reducing the time spent determining current status and next steps.
- **Migration flexibility** – AWS Migration Hub provides the flexibility to use the migration tools that work best for your organisation.
- **Improved visibility** – AWS Migration Hub helps plan your migrations by allowing you to group related servers and resources that should be migrated together.
- **No additional charges** – There is no additional charge for AWS Migration Hub. You only pay for the cost of the migration tools you use and any resources being consumed on AWS.
- **Understand your environment** – AWS Migration Hub helps you understand your IT environment by letting you explore information collected by AWS discovery tools and stored in the AWS Application Discovery Service's

repository.

Use the migration tools you choose – Migration tools can publish your status to AWS Migration Hub by writing to the [AWS Migration Hub API](#).

Authentication and access control – AWS Migration hub is integrated with IAM for authentication and access control.

120.2 Backup/Restore and Disaster Recovery

This requirement is not applicable to AWS Migration Hub. For additional information beyond what is described herein, please refer to <http://aws.amazon.com/documentation/>.

120.3 Pricing Overview

Please visit <https://aws.amazon.com/pricing/> for further details on pricing AWS services.

120.4 Service Constraints

Please see <https://docs.aws.amazon.com/migrationhub/latest/ug/limits.html> for more information.

120.5 Technical Requirements

Please refer to <https://docs.aws.amazon.com/migrationhub/latest/ug/whatishub.html> for comprehensive technical documentation regarding AWS Migration Hub.

121.0 Service Definition – AWS Mobile Hub

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

121.1 Service Overview

AWS Mobile Hub lets customers easily add and configure features for mobile apps, including user authentication, data storage, back-end logic, push notifications, content delivery, and analytics. AWS Mobile Hub gives customers easy access to testing on real devices, as well as analytics dashboards to track app usage, all from a single, integrated console. AWS Mobile Hub includes integration with Amazon Lex, so customers can build mobile apps that use speech and text to a mobile app. Amazon Lex lets customers use the technology that powers Amazon Alexa to create engaging speech- and text-based conversational bots in their own app. AWS Mobile Hub's SaaS Connectors make it easy to securely access data in third-party enterprise SaaS applications from a mobile apps.

Top benefits include:

- **Start for free** - Get started building your app for free. Every AWS service has a free tier. Pay only for the services you use beyond the free tier.
- **Add cloud services fast** - Create amazing cloud-enabled apps in minutes. Use

serverless technology to run code without provisioning or managing servers.

- **Deliver quality apps** - Automate your DevOps pipeline with build, test, and deploy services for your iOS, Android, and web applications.
- **Engage your audience** - Analyse app usage, and engage in meaningful conversations with your users using email, two-way SMS, and mobile push.

121.2 Backup/Restore and Disaster Recovery

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

121.3 Pricing Overview

When you use AWS Mobile, you pay only for the underlying AWS services you use. There are no additional charges for using the AWS Mobile Hub or the Mobile SDKs. AWS services are priced separately based on usage. Each service has a free tier available. Please visit <https://aws.amazon.com/pricing/> for further details on pricing AWS services.

121.4 Service Constraints

Please see <http://aws.amazon.com/waf/> for more information.

121.5 Technical Requirements

AWS Mobile Hub provides an integrated console that helps you build, test, and monitor your mobile apps. Use the console to choose the features you want to include in your app. Mobile Hub then provisions and configures the necessary AWS services on your behalf and creates a working sample app for you.

Developer Guide – Walks you through how to get started and use the AWS Mobile Hub console. [HTML](#) | [PDF](#)

122.0 Service Definition – AWS OpsWorks

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

This service listing includes the following AWS services:

- AWS OpsWorks
- AWS OpsWorks for Puppet Enterprise
- AWS OpsWorks for Chef Automate

- AWS OpsWorks Stacks

122.1 Service Overview

AWS OpsWorks is a configuration management service that helps you configure and operate applications of all shapes and sizes using Chef. You can define the application's architecture and the specification of each component, including package installation, software configuration, and resources such as storage. Start from templates for common technologies like application servers and databases or build your own to perform any task that can be scripted. AWS OpsWorks includes automation to scale your application based on time or load and dynamic configuration to orchestrate changes as your environment scales.

Top benefits include:

- **Supports Any Application** – AWS OpsWorks supports a wide variety of architectures, from simple web applications to highly complex custom applications and any software that has a scripted installation. Since AWS OpsWorks supports Chef recipes and Bash/PowerShell scripts, you can leverage community-built configurations such as MongoDB and Elasticsearch. You start by modelling and visualizing your application with layers that define resource and software configuration. You control every aspect of your application's configuration to match your needs, processes, and tools.
- **Configuration as Code** – AWS OpsWorks lets you define configurations for your entire environment in a format that you can maintain and version just like your application source code. You can reproduce the software configuration on new instances and apply changes to all running instances, ensuring consistent configuration at any time. You can deploy your application from source repositories such as Git and Subversion to one or hundreds of instances with the click of a button.
- **Automation to Run at Scale** – AWS OpsWorks provides dynamic configuration and orchestration that allows you to efficiently manage your applications over their lifetime, including support for automatic instance scaling and auto healing. Each new instance that comes online is built to specification automatically and can change its configuration in response to system events. This lets you, for example, change the configuration of a web application firewall when a new web server comes online without any manual steps.
- **Resource Organisation** – AWS OpsWorks lets you model and visualise your application using concepts such as stacks, layers, and apps. The AWS OpsWorks dashboard shows the status of your stacks across all AWS Regions. AWS OpsWorks tags your resources with the stack and layer names to make discovery easier and to support logical groupings for monitoring, cost allocation, and permissions. You can grant AWS IAM users access to specific stacks, making management of multi-user environments easier.
- **Supports Any Server** – AWS OpsWorks helps you automate operational tasks

like software configurations, package installations, database setups, and code deployment on any Linux or Windows server, including existing Amazon EC2 instances or servers running in your own data centre. You can use a single configuration management service to deploy and operate applications across your hybrid architecture.

122.2 Backup/Restore and Disaster Recovery

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

122.3 Pricing Overview

Please visit <https://aws.amazon.com/pricing/> for further details on pricing AWS services.

122.4 Service Constraints

Please see <https://aws.amazon.com/opsworks/> for more information.

122.5 Technical Requirements

AWS OpsWorks provides a simple and flexible way to create and manage stacks and applications. With AWS OpsWorks, you can provision AWS resources, manage their configuration, deploy applications to those resources, and monitor their health.

- **User Guide** – Provides detailed descriptions of all the AWS OpsWorks concepts and provides instructions on using the various features with both the console and the CLI. [HTML](#) | [PDF](#) | [Kindle](#)
- **API Reference** – Describes all the API operations for AWS OpsWorks in detail. In addition, it provides sample requests, responses, and errors for the supported web services protocols. [HTML](#) | [PDF](#)

123.0 Service Definition – AWS Organizations

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

123.1 Service Overview

AWS Organizations offers policy-based management for multiple AWS accounts. With Organizations, you can create groups of accounts and then apply policies to those groups. Organizations enables you to centrally manage policies across multiple accounts, without requiring custom scripts and manual processes.

Using AWS Organizations, you can create Service Control Policies (SCPs) that centrally control AWS service use across multiple AWS accounts. You can also use Organizations to help automate the creation of new accounts through APIs.

Organizations helps simplify the billing for multiple accounts by enabling you to setup a single payment method for all the accounts in your organisation through consolidated billing. AWS Organizations is available to all AWS customers at no additional charge. Top benefits include:

- **Centrally manage policies across multiple AWS accounts** - AWS Organizations helps you manage policies for multiple AWS accounts. With Organizations, you can create groups of accounts, and then attach policies to a group to ensure the correct policies are applied across the accounts. Organizations enables you to centrally manage policies across multiple accounts, without requiring custom scripts and manual processes.
- **Control access to AWS services** - With AWS Organizations, you can create Service Control Policies (SCPs) that centrally control AWS service use across multiple AWS accounts. SCPs put bounds around the permissions that AWS Identity and Access Management (IAM) policies can grant to entities in an account, such as IAM users and roles. For example, IAM policies for an account in your organisation cannot grant access to AWS Direct Connect if access is not also allowed by the SCP for the account. Entities can only use the services allowed by both the SCP and the IAM policy for the account.
- **Automate AWS account creation and management** - You can use the AWS Organizations APIs to automate the creation and management of new AWS accounts. The AWS Organizations APIs enable you to create new accounts programmatically, and to add the new accounts to a group. The policies attached to the group are automatically applied to the new account. For example, you can automate the creation of sandbox accounts for developers and grant entities in those accounts access only to the necessary AWS services.
- **Consolidate billing across multiple AWS accounts** - AWS Organizations enables you to set up a single payment method for all the AWS accounts in your organisation through consolidated billing. With consolidated billing, you can see a combined view of charges incurred by all your accounts, as well as take advantage of pricing benefits from aggregated usage, such as volume discounts for Amazon EC2 and Amazon S3.
- **Configure AWS services across multiple accounts** - AWS Organizations helps you configure AWS services and share resources across accounts in your organization. For example, Organizations integrates with AWS Single Sign-on to enable you to easily provision access for all of your developers to accounts in your organization from a single place. You can make central changes to access permissions and have them automatically updated on accounts in your organization.

123.2 Backup/Restore and Disaster Recovery

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

123.3 Pricing Overview

Please visit <https://aws.amazon.com/pricing/> for further details on pricing AWS services.

123.4 Service Constraints

Please see <https://aws.amazon.com/organizations/> for more information.

123.5 Technical Requirements

AWS Organizations is a service that enables Amazon Web Services (AWS) customers to consolidate and centrally manage multiple AWS accounts. With AWS Organizations, you can create accounts and invite existing accounts to join your organisation. You can organise those accounts into groups and attach policy-based controls. If you already have a Consolidated Billing family of accounts, those accounts automatically become part of your organisation.

- **User Guide** - Introduces you to AWS Organizations, helps you set up an organisation by inviting other accounts to join, and shows you how to organise your accounts into groups to control access to your AWS resources by applying service control policies. [HTML](#) | [PDF](#) | [Kindle](#)
- **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. [HTML](#) | [PDF](#)
- **AWS Organisations 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. [HTML](#)

124.0 Service Definition - AWS Resource Access Manager

AWS Resource Access Manager (RAM) is a service that enables you to easily and securely share AWS resources with any AWS account or within your AWS Organization.

124.1 Service Overview

Many organizations use multiple accounts to create administrative or billing isolation, and to limit the impact of errors. RAM eliminates the need to create duplicate resources in multiple accounts, reducing the operational overhead of managing those resources in every single account you own. You can create resources centrally in a multi-account environment, and use RAM to share those resources across accounts in three simple steps: create a Resource Share, specify resources, and specify accounts. RAM is available to you at no additional charge.

Top Features include

- Integrates with the following services: Amazon Aurora, AWS CodeBuild, Amazon EC2, Amazon EC2 Image Builder, AWS License Manager, AWS Resource Groups, Amazon Route 53

- Share resources with any AWS Account
- Integrates with AWS Organizations
-

Top benefits include

- Reduces operational overhead - Create resources centrally and use AWS RAM to share those resources with other accounts. This eliminates the need to provision duplicate resources in every account, which reduces operational overhead.
- Provides security and consistency—Govern consumption of shared resources using existing policies and permissions, to achieve security and control. AWS RAM offers a consistent experience for sharing different types of AWS resources.
- Provides visibility and auditability—View usage details for shared resources through integration with Amazon CloudWatch and AWS CloudTrail. AWS RAM provides comprehensive visibility into shared resources and accounts.

124.2 Backup/Restore and Disaster Recovery

This requirement is not applicable for AWS Resource Access Manager. For additional information beyond what is described herein, please refer to <https://docs.aws.amazon.com/ram/>

124.3 Pricing Overview

Please see the AWS UK G Cloud 12 Pricing Document affiliated with this service in the Digital Marketplace.

124.4 Service Constraints

Please see <https://docs.aws.amazon.com/ram/latest/userguide/what-is.html#what-is-limits> for more information.

125.0 Service Definition – AWS RoboMaker

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

125.1 Service Overview

AWS RoboMaker is a service that makes it easy to develop, test, and deploy intelligent robotics applications at scale. AWS RoboMaker extends the most widely used open source robotics software framework, Robot Operating System (ROS), with connectivity to cloud services. This includes AWS machine learning services, monitoring services, and analytics services that enable a robot to stream data, navigate, communicate, comprehend, and learn. AWS RoboMaker provides a robotics development environment for application development, a robotics simulation service to accelerate application testing, and a robotics fleet management service for remote application deployment, update, and management.

Robots are machines that sense, compute, and take action. Robots need instructions to accomplish tasks, and these instructions come in the form of applications that developers code to determine how the robot will behave. Receiving and processing sensor data, controlling actuators for movement, and performing a specific task are all functions that are typically automated by these intelligent robotics applications. Intelligent robots are being increasingly used in warehouses to distribute inventory, in homes to carry out tedious housework, and in retail stores to provide customer service. Robotics applications use machine learning in order to perform more complex tasks like recognising an object or face, having a conversation with a person, following a spoken command, or navigating autonomously. Until now, developing, testing, and deploying intelligent robotics applications was difficult and time consuming. Building intelligent robotics functionality using machine learning is complex and requires specialised skills. Setting up a development environment can take each developer days, and building a realistic simulation system to test an application can take months due to the underlying infrastructure needed. Once an application has been developed and tested, a developer needs to build a deployment system to deploy the application into the robot and later update the application while the robot is in use.

AWS RoboMaker provides the tools to make building intelligent robotics applications more accessible, a fully managed simulation service for quick and easy testing, and a deployment service for lifecycle management. AWS RoboMaker removes the heavy lifting from each step of robotics development.

The top features include:

- **Cloud extensions for ROS** – ROS is the most widely used open source robotics software framework, providing software libraries that help you build robotics applications. AWS RoboMaker provides cloud extensions for ROS so that you can offload the more resource-intensive computing processes that are typically required for intelligent robotics applications to the cloud and free up local compute resources.
- **Computer vision with Amazon Kinesis and Amazon Rekognition** – You can use Amazon Kinesis and Amazon Rekognition to build a computer vision application that offloads compute resources to the cloud.
- **Voice command with Amazon Lex and Amazon Polly** – Amazon Lex provides high-quality speech recognition and NLU, plus intent chaining, so you can simplify complex conversations directed towards the robot by breaking them into smaller components.
- **Monitoring and logging with Amazon CloudWatch** – Amazon CloudWatch gives you actionable insights that help you optimise application performance, manage resource utilisation, and understand system-wide operational health of your fleet of robots.
- **Development environment** – AWS RoboMaker provides a robotics development environment for building and editing robotics applications. The

AWS RoboMaker development environment is based on AWS Cloud9, so you can launch a dedicated workspace to edit, run, and debug robotics application code.

Top benefits include:

- **Get started quickly** – AWS RoboMaker includes sample robotics applications to help you get started quickly.
- **Build intelligent robots** – Because AWS RoboMaker is pre-integrated with popular AWS analytics, machine learning, and monitoring services, it's easy to add functions like video streaming, face and object recognition, voice command and response, or metrics and logs collection to your robotics application.
- **Lifecycle management** – Manage the lifecycle of a robotics application from building and deploying the application to monitoring and updating an entire fleet of robots.
- **Preconfigured ROS tools** – ROS is preinstalled and configured in the development environment so that you can start editing right away.
- **Pre-integrated sample applications** – AWS RoboMaker provides a number of sample applications pre-integrated and ready for download in the development environment.
- **Fully featured editor** – The AWS RoboMaker development environment 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 to code faster and avoid errors.
- **Simulation** – You can use simulation for testing and fine-tuning robotics applications before deploying to physical hardware. AWS RoboMaker provides a fully managed robotics simulation service that supports large-scale and parallel simulations and automatically scales the underlying infrastructure based on the complexity of the simulation.
- **Automatic scaling** – AWS RoboMaker scales the underlying infrastructure automatically based on the complexity of your robotics application and simulation application. AWS RoboMaker takes care of infrastructure-related tasks, such as capacity planning, compute resource provision, software update, and OS patching, so you don't have to. You only pay for the resources your simulation job consumes.
- **Fleet management** – Once an application has been developed or modified, you'd build an over-the-air (OTA) system to securely deploy the application into the robot and later update the application while the robot is in use.
- **AWS IoT Greengrass integration** – AWS RoboMaker fleet management is integrated with AWS IoT Greengrass so that you can take advantage of additional AWS IoT Greengrass features such as local Lambda functions, local messaging, and machine learning inference.

125.2 Backup/Restore and Disaster Recovery

This requirement is not applicable to AWS RoboMaker. The service does not store your robot application or simulation application. You will upload your applications to your Amazon S3 bucket and refer to the Amazon S3 object path during robot application and simulation application creation.

125.3 Pricing Overview

Please visit <https://aws.amazon.com/pricing/> for further details on pricing AWS services.

125.4 Service Constraints

Please see <https://docs.aws.amazon.com/robomaker/latest/dg/what-is-robomaker.html> for more information.

125.5 Technical Requirements

Please refer to <https://docs.aws.amazon.com/robomaker/latest/dg/what-is-robomaker.html> for comprehensive technical documentation regarding AWS RoboMaker.

126.0 Service Definition – AWS Secrets Manager

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

126.1 Service Overview

AWS Secrets Manager helps you to protect secrets needed by your applications and services in order to operate effectively. The service lets you easily rotate, manage, and retrieve database credentials, API keys, and other secrets by making calls to the Secrets Manager API, eliminating the need to hard-code sensitive information in plain text. Secrets Manager can automatically rotate credentials used by Amazon RDS, Amazon Redshift, and Amazon DocumentDB. The service is extensible to other types of secret, and access to secrets can be controlled using a fine-grained permissions system. Secrets Manager lets you audit secret usage and rotation for secrets in the AWS Cloud, third-party services, and on premises.

The top features include:

- Secrets Manager can rotate Amazon RDS, Amazon Redshift, and Amazon DocumentDB secrets without affecting application uptime. This feature is extensible to other AWS Cloud services and on-premises systems using Lambda.
- Control access to secrets using IAM policies. For example, you could only allow access to certain secrets from your corporate network.
- Encrypt secrets using the same keys you use with AWS KMS.
- Audit secret storage and rotation in AWS CloudTrail and use integrations with

Amazon CloudWatch to alert on infrequently used secrets and secret rotation.

- Manage secrets for workloads subject to controls such as HIPAA, PCI DSS, ISO/IEC 27001, ISO/IEC 27017, ISO/IEC 27018, and ISO 9001.

Top benefits include:

- Meet your security and compliance requirements by eliminating any need to store secrets in plain text.
- Make it easy to rotate credentials without requiring a code deployment.
- Use integrations to Amazon Redshift, Amazon RDS, and Amazon DocumentDB or extend to other systems using Lambda functions.
- Manage access to secrets using IAM policies.
- Encrypt your secrets with keys from AWS KMS.
- Audit secret operations, such as rotations, with AWS CloudTrail and alert based on action using Amazon Cloudwatch.

126.2 Backup/Restore and Disaster Recovery

AWS Secrets Manager is a managed service. Customers do not need to define their own backup, restore, or disaster recovery strategy.

126.3 Pricing Overview

Please see the AWS UK G Cloud 12 Pricing Document affiliated with this service in the Digital Marketplace.

126.4 Service Constraints

Technical constraints related to the usage of AWS Secrets Manager are listed at https://docs.aws.amazon.com/secretsmanager/latest/userguide/reference_limits.html.

126.5 Technical Requirements

Please refer to <https://docs.aws.amazon.com/secretsmanager/latest/userguide/intro.html> for comprehensive technical documentation regarding AWS Secrets Manager.

127.0 Service Definition – AWS Security Hub

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

127.1 Service Overview

AWS Security Hub gives you a comprehensive view of your high-priority security alerts and compliance status across AWS accounts. As an organisation, you'll have a variety of powerful security tools at your disposal, from firewalls and endpoint protection to

vulnerability and compliance scanners, but operating these often leads to your team switching back and forth between tools to deal with alerts.

AWS Security Hub aggregates, organises, and prioritises findings from multiple products, such as Amazon GuardDuty and Amazon Inspector, as well as from AWS partner solutions. It visually summarises the findings and lets you continuously monitor your infrastructure using automated compliance checks based on AWS best practices and industry standards.

The top features include:

- AWS Security Hub collects and aggregates findings from AWS security services enabled in your accounts, such as intrusion detection from Amazon GuardDuty, vulnerability scans from Amazon Inspector, and findings from integrated AWS partner solutions.
- Findings are grouped together, making it easy to highlight emerging trends. For instance, insights relating to Amazon EC2 instances missing security patches group together, as do insights relating to Amazon S3 buckets with public read or write permissions.
- Create your own insights or customise AWS Security Hub's preconfigured insights. For instance, create an insight to identify EC2 instances that don't meet your organisation's security standard and are tagged as "production."
- Associate AWS Security Hub with multiple accounts and aggregate findings across those accounts. Identify which accounts need urgent remediation using in-built analytics.
- Automate compliance checks using standard industry benchmarks, such as CIS AWS Foundations. This provides clear, step-by-step remediation guidelines on an account-by-account basis.
- Integrate services from partners such as AlertLogic, Armor, Barracuda, CheckPoint, CrowdStrike, CyberArk, Demisto, Dome9, F5 Networks, Fortinet, GuardCore, IBM, McAfee, Palo Alto Networks, Qualys, Rapid7, Redlock, Sophos, Splunk, Sumo Logic, Symantec, Tenable, Trend Micro, Turbot, and Twistlock by making use of a standardised findings format. This eliminates the need to standardise before aggregating.

Top benefits include:

- Easily spot trends and identify potential issues using integrated dashboards to bring together your findings.
- Save time with aggregated findings by ingesting data using a common format and prioritising findings across providers.
- Run continuous, automated account-level configuration and compliance checks based on industry standards, such as the CIS AWS Foundations benchmark.
- Integrate findings with Amazon CloudWatch Events to send findings to ticketing,

chat, email, or remediation systems automatically.

127.2 Backup/Restore and Disaster Recovery

Findings are stored for 90 days within AWS Security Hub. You can export findings at any time.

127.3 Pricing Overview

Please see the AWS UK G Cloud 12 Pricing Document affiliated with this service in the Digital Marketplace.

127.4 Service Constraints

Please see

https://docs.aws.amazon.com/securityhub/latest/userguide/securityhub_limits.html for more information.

127.5 Technical Requirements

Please refer to <https://docs.aws.amazon.com/securityhub/latest/userguide/what-is-securityhub.html> for comprehensive technical documentation regarding AWS Security Hub.

128.0 Service Definition – AWS Server Migration Service (SMS)

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

128.1 Service Overview

AWS Server Migration Service (SMS) is an agentless service which makes it easier and faster for you to migrate thousands of on-premises workloads to AWS. AWS SMS allows you to automate, schedule, and track incremental replications of live server volumes, making it easier for you to coordinate large-scale server migrations.

Top benefits include:

- **Easy to Get Started** - Start and manage server migration with a few clicks via the AWS Management Console. AWS Server Migration Service will automatically replicate live server volumes to AWS and create Amazon Machine Images (AMI) as needed.
- **Control** - Create and manage a customized replication schedule designed for large-scale migrations, and track the progress of each migration.
- **Agility** - Perform migrations faster while minimizing network bandwidth, by migrating only incremental changes made to on-premises servers.
- **Cost-Effective** - AWS Server Migration Service is free to use; pay only for the storage resources used during the migration process.

- **Minimize Downtime** - Incremental server replication allows you to reduce server downtime significantly.

128.2 Backup/Restore and Disaster Recovery

This requirement is not applicable for AWS Server Migration Service (SMS). For additional information beyond what is described herein, please refer to <https://aws.amazon.com/documentation/>

128.3 Pricing Overview

You can use the AWS Server Migration Service to migrate your on-premises workloads to AWS at no charge.

128.3.1 Additional Charges

The AWS Server Migration Service creates a new EBS snapshot with every replication. You will incur additional charges for [EBS snapshots](#). To prevent additional charges, delete snapshot copies you no longer need.

The AWS Server Migration Service replicates server volumes from your on premises environment to S3 temporarily and purges them from S3 right after creating EBS snapshots, incurring a transient charge for [S3](#).

Please see <https://aws.amazon.com/server-migration-service/pricing/> for more information.

128.4 Service Constraints

Please see <https://aws.amazon.com/server-migration-service/faqs/> for more information on service constraints for AWS Server Migration Service (SMS).

128.5 Technical Requirements

AWS Server Migration Service (SMS) combines data collection tools with automated server replication to speed the migration of on-premises servers to AWS.

- **User Guide** - Describes key concepts of AWS SMS and provides instructions for using the features of AWS SMS. [HTML](#) | [PDF](#)
- **SMS section of AWS CLI Reference** - Documents the AWS SMS commands available in the AWS CLI. [HTML](#)

129.0 Service Definition – AWS Serverless Application Repository

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

129.1 Service Overview

The AWS Serverless Application Repository is a collection of serverless applications and serverless application components provided by AWS and other AWS partners and developers. With a growing selection of applications to choose from, the AWS Serverless Application Repository makes it easy to get started with the AWS serverless platform. The AWS Serverless Application Repository includes applications for Alexa Skills, chatbots, data processing, IoT, real time stream processing, web and mobile back-ends, social media trend analysis, image resizing, and more from publishers on AWS.

Top benefits include:

- **Deploy quickly** – quickly deploy code samples, components, and complete applications for common use cases such as web and mobile back-ends, event and data processing, logging, monitoring, IoT, and more.
- **Easily define AWS resources used** - each application is packaged with an [AWS Serverless Application Model \(SAM\)](#) template that defines the AWS resources used.
- **Publishing Applications to the Repository** - Publish applications to the AWS Serverless Application Repository to share solutions with developers or to help your customers quickly understand the value of products and services you sell and support.
- **Easy to share your published applications within your team**, across your organisation, or with the community at large. Publicly shared applications must include a link to the application's source code so others can view what the application does and how it works.

129.2 Backup/Restore and Disaster Recovery

This requirement is not applicable for AWS Serverless Application Repository. For additional information beyond what is described herein, please refer to <https://aws.amazon.com/documentation/>

129.3 Pricing Overview

Please visit <https://aws.amazon.com/pricing/> for further details on pricing AWS services.

129.4 Service Constraints

AWS Serverless Application Repository features and default service limits are found at <https://aws.amazon.com/serverless/serverlessrepo/faqs/>

129.5 Technical Requirements

With AWS Serverless Application Repository, you can quickly find and deploy serverless applications in the AWS Cloud. You can browse applications by category, or

search for them by name, publisher, or event source. To use an application, you simply select it, configure required fields, and deploy it with a few clicks.

As a serverless application developer, you can also easily publish applications, sharing them across teams and organisations and with other AWS users. To publish serverless applications, you can use the AWS Management Console, AWS CLI, or AWS SDKs to upload the code, along with a simple manifest file.

- **Developer Guide** - Provides a conceptual overview of AWS Serverless Application Repository, includes detailed instructions for using features, and provides a complete API reference for developers. [HTML](#) | [PDF](#)

130.0 Service Definition – AWS Service Catalog

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

130.1 Service Overview

AWS Service Catalog allows organisations to create and manage catalogues of IT services that are approved for use on AWS. These IT services can include everything from virtual machine images, servers, software, and databases to complete multi-tier application architectures. AWS Service Catalog allows you to centrally manage commonly deployed IT services and helps you achieve consistent governance and meet your compliance requirements, while enabling users to quickly deploy only the approved IT services they need.

Top benefits include:

- **Ensure Compliance with Corporate Standards – AWS Service Catalog** provides a single location where organisations can centrally manage catalogues of IT services. With AWS Service Catalog, you can control which IT services and versions are available, the configuration of the available services, and permission access by individual, group, department, or cost centre.
- **Help Employees Quickly Find and Deploy Approved IT Services – With AWS Service Catalog**, you define your own catalogue of AWS Cloud services and AWS Marketplace software and make them available for your organisation. Then, end users can quickly discover and deploy IT services using a self-service portal.
- **Centrally Manage IT Service Lifecycle – AWS Service Catalog** enables you to add new versions of IT services as necessary, and end users are notified so they can easily keep abreast of the latest updates. With AWS Service Catalog you can control the use of IT services by specifying constraints, such as the AWS Region in which a product can be launched or allowed IP ranges.

130.2 Backup/Restore and Disaster Recovery

This requirement is not applicable for AWS Service Catalog. For additional information beyond what is described herein, please refer to <http://aws.amazon.com/documentation/>.

130.3 Pricing Overview

Please visit <https://aws.amazon.com/pricing/> for further details on pricing AWS services.

130.4 Service Constraints

Please see <https://aws.amazon.com/servicecatalog/> for more information.

130.5 Technical Requirements

AWS Service Catalog allows IT administrators to create, manage, and distribute portfolios of approved products to end users, who can then access the products they need in a personalised portal. Typical products include servers, databases, websites, or applications that are deployed using AWS resources (for example, an Amazon EC2 instance or an Amazon RDS database). You can control which users have access to specific products to enforce compliance with organisational business standards, manage product lifecycles, and help users find and launch products with confidence.

- **Administrator Guide** – Provides a conceptual overview of AWS Service Catalog and includes detailed instructions for using the service as an administrator. [HTML](#) | [PDF](#)
- **User Guide** – Provides a conceptual overview of AWS Service Catalog and includes detailed instructions for using the service as an end user. [HTML](#) | [PDF](#)

131.0 Service Definition – AWS Shield

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

This service listing includes the following AWS services:

- AWS Shield Advanced
- AWS Shield Standard

131.1 Service Overview

AWS Shield is a managed DDoS protection service that safeguards web applications running on AWS. AWS Shield provides always-on detection and automatic inline mitigations that minimise application downtime and latency, so there is no need to engage AWS Support to benefit from DDoS protection. There are two tiers of AWS Shield – Standard and Advanced. All AWS customers benefit from the automatic

protections of AWS Shield Standard, at no additional charge. AWS Shield Standard defends against most common, frequently occurring network and transport layer DDoS attacks that target your website or applications.

For higher levels of protection against attacks targeting your web applications running on Elastic Load Balancing, Amazon CloudFront, and Amazon Route 53 resources, you can subscribe to AWS Shield Advanced. In addition to the common network and transport layer protections that come with Standard, AWS Shield Advanced provides additional detection and mitigation against large and sophisticated DDoS attacks, near real-time visibility into attacks, and integration with AWS WAF, a web application firewall. AWS Shield Advanced also gives you access to the AWS DRT and protection against DDoS related spikes in your Elastic Load Balancing, Amazon CloudFront, or Amazon Route 53 charges.

AWS Shield Advanced is available globally on all Amazon CloudFront and Amazon Route 53 edge locations. You can protect your web applications hosted anywhere in the world by deploying Amazon CloudFront in front of your application. Your origin servers can be Amazon S3, Amazon EC2, Elastic Load Balancing, or a custom server outside of AWS. You can also enable AWS Shield Advanced directly on Elastic Load Balancing in the following AWS Regions: Northern Virginia, Oregon, Ireland, and Tokyo.

Top benefits include:

- **Seamless Integration and Deployment** – With AWS Shield your AWS resources are automatically protected from common, most frequently occurring network and transport layer DDoS attacks. You can achieve a higher level of defence by simply enabling AWS Shield Advanced protection for Elastic Load Balancing, Amazon CloudFront or Amazon Route 53 resources you want to protect using the management console or APIs.
- **Customisable Protection** – With AWS Shield Advanced, you have the flexibility to write customised rules to mitigate sophisticated application layer attacks. These customisable rules can be deployed instantly, allowing you to quickly mitigate attacks. You can set up proactive rules to automatically block bad traffic, or respond to incidents as they occur. You can also engage the 24x7 AWS DRT, who can write rules on your behalf to mitigate application layer DDoS attacks.
- **Cost Efficient** – As an AWS customer, you automatically get network layer protection against some of the most common DDoS attacks with AWS Shield.

This protection does not require additional cost, resources, or time to initiate. With AWS Shield Advanced, you get “DDoS cost protection”, a feature that protects your AWS bill from Elastic Load Balancing, Amazon CloudFront, and Amazon Route 53 usage spikes as a result of a DDoS attack.

131.2 Backup/Restore and Disaster Recovery

This requirement is not applicable for AWS Shield Advanced. For additional information beyond what is described herein, please refer to <http://aws.amazon.com/documentation/>.

131.3 Pricing Overview

Please visit <https://aws.amazon.com/pricing/> for further details on pricing AWS services.

131.4 Service Constraints

Please see <https://aws.amazon.com/shield/> for more information.

131.5 Technical Requirements

For additional protection against DDoS attacks, AWS also offers AWS Shield Advanced.

AWS Shield Advanced provides expanded DDoS attack protection for your Amazon CloudFront distributions, Amazon Route 53 hosted zones, and Elastic Load Balancing load balancers. AWS Shield Advanced incurs additional charges.

- **Developer Guide** – Describes how to get started with AWS WAF and AWS Shield Advanced, explains key concepts, and provides step-by-step instructions that show you how to use the features. [HTML](#) | [PDF](#) | [Kindle](#)
- **API Reference** – Describes all the API operations for AWS Shield Advanced in detail. [HTML](#) | [PDF](#)

132.0 Service Definition – AWS Single Sign-On

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

132.1 Service Overview

AWS Single Sign-On (SSO) is a cloud SSO service that makes it easy to centrally manage SSO access to multiple AWS accounts and business applications. You can easily manage SSO access and user permissions to all accounts within an AWS Organisation centrally and connect to other business applications such as Salesforce, Box, Office 365, or your own service using SAML 2.0. User credentials can be configured in AWS SSO, or existing corporate credentials may be used through integration with existing directory services.

Top features include:

- Centrally manage access and user permissions for all AWS accounts managed through AWS Organizations, without additional setup in each account. Assign users permissions based on common job functions and customise these permissions to meet your specific security requirements.
- Create and manage users within AWS SSO, organising users into groups, or connect to Microsoft Active Directory through AWS Directory Service, so users can sign in with their existing Active Directory credentials.
- Enable a highly available SSO service for your organisation in minutes, with no infrastructure to maintain or deploy. Meet your security requirements by having all actions logged in AWS CloudTrail to get visibility into SSO actions in one

place.

- Access services such as Salesforce, Box, and Office 365 using SAML integrations, or create your own SAML integration to existing corporate applications.

Top benefits include:

- Centrally manage access and user permissions for all AWS accounts managed through AWS Organizations, without additional setup in each account. Assign users permissions based on common job functions and customise these permissions to meet your specific security requirements.
- Create and manage users within AWS SSO, organising users into groups, or connect to Microsoft Active Directory through AWS Directory Service, so users can sign in with their existing Active Directory credentials.
- Enable a highly available SSO service for your organisation in minutes, with no infrastructure to maintain or deploy. Meet your security requirements by having all actions logged in AWS CloudTrail to get visibility into SSO actions in one place.
- Access services such as Salesforce, Box, and Office 365 using SAML integrations, or create your own SAML integration to existing corporate applications.

132.2 Backup/Restore and Disaster Recovery

AWS SSO creates a fault-tolerant, resilient, and highly available SSO service on your behalf. AWS SSO is responsible for backing up and restoring.

132.3 Service Constraints

Please see <https://docs.aws.amazon.com/singlesignon/latest/userguide/limits.html> for more information.

132.4 Pricing Overview

Please see the AWS UK G Cloud 12 Pricing Document affiliated with this service in the Digital Marketplace.

132.5 Technical Requirements

Please refer to <https://docs.aws.amazon.com/singlesignon/latest/userguide/what-is.html> for comprehensive technical documentation regarding AWS Single Sign-On.

133.0 Service Definition – AWS Snowball Edge

The following subsections provide service definition information.

133.1 Service Overview

AWS Snowball Edge is a 100TB data transfer device with on-board storage and compute capabilities. You can use Snowball Edge to move large amounts of data into and out of AWS, as a temporary storage tier for large local datasets, or to support local workloads in remote or offline locations.

Snowball Edge connects to your existing applications and infrastructure using standard storage interfaces, streamlining the data transfer process and minimizing setup and integration. Snowball Edge can cluster together to form a local storage tier and process your data on-premises, helping ensure your applications continue to run even when they are not able to access the cloud.

Top benefits include:

- **Process Data Locally** - You can deploy AWS Lambda code on Snowball Edge to perform tasks such as analysing data streams or processing data locally. Data is collected, stored, and processed using the Snowball Edge independent from any other storage or compute resource.
- **Stand-Alone Storage** - You can use Snowball Edge as stand-alone storage. You can point existing applications to Snowball Edge devices as a file or object storage destination without changing application logic. You can also cluster Snowball Edge devices together into a single, larger, storage tier with increased durability. When a Snowball Edge is full, it can be removed from the cluster, returned to AWS, and replaced with a new Snowball Edge.
- **Secure** - Snowball Edge devices use tamper-resistant enclosures, 256-bit encryption, and industry-standard Trusted Platform Modules (TPM) designed to ensure both security and full chain-of-custody for your data. Encryption keys are managed with the [AWS Key Management Service \(KMS\)](#) and they are never stored on the device.
- **Scalable** - Snowball Edge devices can transport multiple terabytes of data and multiple devices can be used in parallel to transfer petabytes of data into or out of AWS. Snowball Edge is currently available in select regions and your location will be verified once you create a job in the AWS Management Console.

133.2 Backup/Restore and Disaster Recovery

This requirement is not applicable for AWS Snowball Edge. For additional information beyond what is described herein, please refer to <https://docs.aws.amazon.com/snowball/latest/developer-guide/whatisedge.html>.

133.3 Pricing Overview

Snowball Edge has two pricing elements. First, you pay a service fee per data transfer job, which includes 10 days of onsite Snowball Edge device usage. Shipping days, including the day the device is received and the day it is shipped back to AWS, are not counted toward the 10 days.

Second, if the device is kept for more than 10 days, you can incur an additional fee for each day beyond 10 days.

Data transferred into AWS does not incur any data transfer fees, and standard pricing applies to data once it is stored in the AWS Cloud. Please see <https://aws.amazon.com/snowball-edge/pricing/> for more information.

133.4 Service Constraints

Please see <https://aws.amazon.com/documentation/snowball/> for more information.

133.5 Technical Requirements

AWS Snowball Edge is a service for customers who want to transport terabytes or petabytes of data to and from AWS, or who want to access the storage and compute power of the AWS Cloud locally and cost effectively in places where connecting to the internet may not be an option.

AWS Snowball Edge Developer Guide - Provides a conceptual overview of how to use AWS Snowball with an Edge appliance, and includes guidance for local storage and compute, clustering, importing and exporting data into Amazon S3, and other features of an Edge appliance. [HTML](#) | [PDF](#)

134.0 Service Definition – AWS Snowmobile

The following subsections provide service definition information.

134.1 Service Overview

AWS Snowmobile is an Exabyte-scale data transfer service used to move extremely large amounts of data to AWS. You can transfer up to 100PB per Snowmobile, a 45foot long ruggedized shipping container, pulled by a semi-trailer truck. Snowmobile makes it easy to move massive volumes of data to the cloud, including video libraries, image repositories, or even a complete data centre migration. Transferring data with Snowmobile is secure, fast and cost effective.

After an initial assessment, a Snowmobile will be transported to your data centre and AWS personnel will configure it for you so it can be accessed as a network storage target. When your Snowmobile is on site, AWS personnel will work with your team to connect a removable, high-speed network switch from Snowmobile to your local network and you can begin your high-speed data transfer from any number of sources within your data centre to the Snowmobile. After your data is loaded, Snowmobile is driven back to AWS where your data is imported into Amazon S3 or Amazon Glacier.

Snowmobile uses multiple layers of security designed to protect your data including dedicated security personnel, GPS tracking, alarm monitoring, 24/7 video surveillance, and an optional escort security vehicle while in transit. All data is encrypted with 256-bit encryption keys managed through the [AWS Key Management Service](#) (KMS) and designed to ensure both security and full chain-of-custody of your data.

Top benefits include:

- **Fast Transfer Even at Massive Scale** - Even with high-speed internet connections, it can take decades to transfer extremely large amounts of data. With Snowmobile, you can move 100 petabytes of data in as little as a few weeks, plus transport time. That same transfer could take more than 20 years to accomplish over a direct connect line with a 1Gbps connection.
- **Strong Encryption** - Your data is encrypted with keys that you provide before it is written to Snowmobile. All data is encrypted with 256-bit encryption and you can manage your encryption keys with AWS Key Management Service (AWS KMS). Encryption keys used by encrypting servers on the Snowmobile are never written to disk. Should power be removed for any reason, the keys are securely erased.
- **Rugged, Durable, and Secure** - Snowmobile is tamper-resistant, water-resistant and temperature controlled. The data container is only operated by AWS personnel and physical access is limited via security access hardware controls. Snowmobile is protected by 24/7 video surveillance and alarm monitoring, GPS tracking and may optionally be escorted by a security vehicle while in transit.
- **Customized for Your Needs** - As physical sites can have different migration requirements, AWS will work with you to help ensure all needs are factored in before Snowmobile arrives on site.
- **Massively Scalable** - One Snowmobile can transport up to one hundred petabytes of data in a single trip, the equivalent of using about 1,250 [AWS Snowball devices](#).

134.2 Backup/Restore and Disaster Recovery

This requirement is not applicable for AWS Snowmobile. For additional information beyond what is described herein, please refer to <https://aws.amazon.com/snowmobile/faqs/>.

134.3 Pricing Overview

Please see the AWS UK G Cloud 12 Pricing Document affiliated with this service in the Digital Marketplace. Please note that this price is for provisional storage.

134.4 Service Constraints

Please see <https://aws.amazon.com/snowmobile/faqs/> for more information.

134.5 Technical Requirements

A Snowmobile comes with a removable connector cabinet that needs to be mounted on one of your data centre racks where it can be connected directly to your high-speed network backbone. The connector rack provides multiple 40Gb/s interfaces that can transfer up to 1 Tb/s in aggregate. Please refer to <https://aws.amazon.com/snowmobile/> for more information on how to get started.

135.0 Service Definition – AWS Step Functions

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

135.1 Service Overview

AWS Step Functions makes it easy to coordinate the components of distributed applications and microservices using visual workflows. Building applications from individual components that each perform a discrete function lets you scale and change applications quickly. AWS Step Functions is a reliable way to coordinate components and step through the functions of your application. AWS Step Functions provides a graphical console to arrange and visualise the components of your application as a series of steps. This makes it simple to build and run multi-step applications. AWS Step Functions automatically triggers and tracks each step, and retries when there are errors, so your application executes in order and as expected. AWS Step Functions logs the state of each step, so when things do go wrong, you can diagnose and debug problems quickly. You can change and add steps without even writing code, so you can easily evolve your application and innovate faster.

AWS Step Functions manages the operations and underlying infrastructure for you to help ensure your application is available at any scale.

Top benefits include:

- **Productivity: Build Applications Quickly** – AWS Step Functions includes a visual console and blueprints for commonly-used workflows that make it easy to coordinate the components of distributed applications into parallel and/or sequential steps. You can build applications in a matter of minutes, and then visualise and track the execution of each step to help ensure the application is operating as intended
- **Resilience: Scale and Recover Reliably** – AWS Step Functions automatically triggers each step so your application executes in order and as expected. It can handle millions of steps simultaneously to help ensure your application is available as demand increases. AWS Step Functions tracks the state of each step and handles errors with built-in retry and fall-back, whether the step takes seconds or months to complete.
- **Agility: Evolve Applications Easily** – AWS Step Functions makes it easy to change workflows and edit the sequence of steps without revising the entire application. You can re-use components and steps without even changing their code to experiment and innovate faster. Your workflow can support thousands of individual components and steps, so you can freely build increasingly complex applications.

135.2 Backup/Restore and Disaster Recovery

This requirement is not applicable for AWS Step Functions. For additional information beyond what is described herein, please refer to <http://aws.amazon.com/documentation/>.

135.1 Pricing Overview

Please see the AWS UK G Cloud 12 Pricing Document affiliated with this service in the Digital Marketplace.

135.2 Service Constraints

Please see <https://aws.amazon.com/step-functions/> for more information.

135.3 Technical Requirements

AWS Step Functions makes it easy to coordinate the components of distributed applications as a series of steps in a visual workflow. You can quickly build and run state machines to execute the steps of your application in a reliable and scalable fashion.

- **Developer Guide** – Describes key concepts of AWS Step Functions and provides instructions for using the features of AWS Step Functions. [HTML](#) | [PDF](#)
- **API Reference** – Documents the AWS Step Functions API. [HTML](#) | [PDF](#)

136.0 Service Definition – AWS Storage Gateway

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

136.1 Service Overview

AWS Storage Gateway is a hybrid storage service that enables your on-premises applications to seamlessly use AWS Cloud storage. You can use the service for backup and archiving, disaster recovery, cloud data processing, storage tiering, and migration. The service helps you reduce and simplify your data centre and branch or remote office storage infrastructure. Your applications connect to the service through a virtual machine or hardware gateway appliance using standard storage protocols, such as NFS, SMB, and iSCSI. The gateway connects to AWS storage services, such as Amazon S3, Amazon S3 Glacier, Amazon S3 Glacier Deep Archive, Amazon EBS, and AWS Backup, providing storage for files, volumes, snapshots, and virtual tapes in AWS. The service includes a highly optimised data transfer mechanism, with bandwidth management, automated network resilience, and efficient data transfer, along with a local cache for low-latency, on-premises access to your most active data.

Top features include:

- **Integrated** – Hybrid cloud storage means your data can be used on premises and stored durably in AWS Cloud storage services, including Amazon S3,

Amazon S3 Glacier, Amazon S3 Glacier Deep Archive, and Amazon EBS. Once data is moved to AWS, you can apply AWS compute, machine learning, and big data analytics services to it. Additionally, you can leverage the full AWS portfolio of security and management services, including AWS Backup, AWS KMS, IAM, Amazon SNS workflows, Amazon CloudWatch, and AWS CloudTrail.

- **Performance** – AWS Storage Gateway caches data in the local VM or hardware gateway appliance, providing low-latency disk and network performance for your most active data, with optimised data transfers occurring to AWS Cloud storage tiers in the background. Users and applications continue to operate using a local storage model while you take advantage of a cloud backend.
- **Optimised transfers** – Compression, encryption, and bandwidth management are built in. AWS Storage Gateway manages local cache offloads to the cloud based on your desired performance parameters, so you can fine-tune the balance of latency and scale for your workloads. Only data that changes is transferred, so you can optimise your network bandwidth.
- **Simple** – No disruptions are required. Download and install the virtual machine or deploy the dedicated hardware appliance, select an interface, and assign local cache capacity. The advanced networking and protocol support are all included, which means there are no clients to install and no network and/or firewall settings to tune. Additionally, the virtual appliance can run both on premises as well as in Amazon EC2 to serve your in-cloud applications.
- **Cloud scale** – Cloud storage is delivered and billed on demand, so you always have just the right amount. Workloads can expand and contract, backup and archive storage can expand without upfront media costs, and you can provision additional storage capacity without new hardware.

Top benefits include:

- **Durable and secure** – Data stored through AWS Storage Gateway benefits from the durability and security embedded in AWS Cloud storage services. Storage management tools like versioning, cross-region replication, and lifecycle management policies can lower the cost of long-term archiving, simplify audit and compliance requirements, and safeguard all of your data—not just the parts kept on premises. All data that AWS Storage Gateway transfers to AWS is encrypted in transit and at rest in AWS.
- **Standard storage protocols** – AWS Storage Gateway seamlessly connects to your local production or backup applications with NFS, SMB, iSCSI, or iSCSI-VTL, so you can adopt AWS Cloud storage without needing to modify your applications. Its protocol conversion and device emulation enables you to access block data on volumes managed by AWS Storage Gateway on top of Amazon S3, store files as native Amazon S3 objects, and keep virtual tape backups online in a virtual tape library backed by Amazon S3 or move the backups to a tape archive tier on Amazon S3 Glacier.
- **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.

- **AWS integrated** – As a native AWS service, AWS Storage Gateway integrates with other AWS Cloud services for storage, backup, and management. The service stores files as native Amazon S3 objects, archives virtual tapes in Amazon S3 Glacier, and stores Amazon EBS Snapshots generated by the Volume Gateway with Amazon EBS. AWS 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.

136.2 Backup/Restore and Disaster Recovery

This requirement is not applicable to AWS Storage Gateway. 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 12 Pricing Document affiliated with this service in the Digital Marketplace.

136.4 Service Constraints

Please see <https://docs.aws.amazon.com/storagegateway/latest/userguide/Performance.html> for more information.

136.5 Technical Requirements

Please refer to <https://docs.aws.amazon.com/storagegateway/latest/userguide/WhatIsStorageGateway.html> and <https://docs.aws.amazon.com/storagegateway/latest/userguide/GettingStarted.html> for comprehensive technical documentation regarding AWS Storage Gateway.

137.0 Service Definition – AWS Systems Manager

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

137.1 Service Overview

AWS Systems Manager gives you visibility and control of your infrastructure on AWS. AWS Systems Manager provides a unified user interface so you can view operational data from multiple AWS services and allows you to automate operational tasks across your AWS resources. With Systems Manager, you can group resources, like [Amazon EC2](#) instances, [Amazon S3](#) buckets, or [Amazon RDS](#) instances, by application, view operational data for monitoring and troubleshooting, and take action on your groups of

resources. Systems Manager simplifies resource and application management, shortens the time to detect and resolve operational problems, and makes it easy to operate and manage your infrastructure securely at scale.

Top benefits include:

- **Shorten the Time to Detect Problems** - AWS Systems Manager helps you quickly view operational data for groups of resources, so you can quickly identify any issues that might impact applications that use those resources. Resources can be grouped by applications, application layers, production versus development environments, or anything else you choose. Systems Manager presents the operational data for your resource groups in a single, easy to read dashboard so you don't have to navigate to other AWS consoles. For example, if you have an application that uses [Amazon EC2](#), [Amazon S3](#), and [Amazon RDS](#), you can use Systems Manager to create a resource group for the application and easily see the software installed on your [Amazon EC2](#) instances, any changes in your [Amazon S3](#) objects, or database instances that have stopped.
- **Easy to Use Automation** - AWS Systems Manager allows you to automate operational actions to help make your teams more efficient. You can automate maintenance and deployment tasks on Amazon EC2 and on-premises instances, or automatically apply patches, updates, and configuration changes across any resource group. Systems Manager provides predefined automation documents for common operational tasks, such as stopping and restarting an EC2 instance, that you can customize to your own specific use cases. Systems Manager also has built-in safety controls, allowing you to incrementally roll out new changes and automatically halt the roll-out when errors occur.
- **Improve Visibility and Control** - AWS Systems Manager helps you easily understand and control the current state of your resource groups. With Systems Manager, you can view detailed system configurations, operating system patch levels, software installations, application configurations, and other details about your environment through the Systems Manager dashboard. Systems Manager is integrated with [AWS Config](#) so you can easily view changes across your resources as they occur over time.
- **Manage Hybrid Environments** - With AWS Systems Manager, you can manage servers running on AWS and in your on-premises data centre through a single interface. Systems Manager securely communicates with a lightweight agent installed on your servers to execute management tasks. This helps you manage resources for Windows and Linux operating systems running on [Amazon EC2](#) or on-premises.
- **Maintain Security and Compliance** - AWS Systems Manager helps maintain security and compliance by scanning your instances against your patch, configuration, and custom policies. You can define patch baselines, maintain up-to-date anti-virus definitions, and enforce firewall policies. You can also remotely manage your servers at scale without manually logging in to each server. Systems Manager also provides a centralized store to manage your configuration

data, whether its plain text, such as database strings, or secrets, such as passwords. This allows you to separate your secrets and configuration data from code.

137.2 Backup/Restore and Disaster Recovery

This requirement is not applicable for AWS Systems Manager. For additional information beyond what is described herein, please refer to <http://aws.amazon.com/documentation/>.

137.1 Pricing Overview

Please see the AWS UK G Cloud 12 Pricing Document affiliated with this service in the Digital Marketplace.

137.2 Service Constraints

Please see <https://aws.amazon.com/documentation/systems-manager/> for more information.

137.3 Technical Requirements

Use AWS Systems Manager to organise, monitor, and automate management tasks on your AWS resources.

- **User Guide** – Provides detailed descriptions of AWS Systems Manager concepts, and includes instructions for using both the console and the command line interface to configure, manage, and automate tasks on groups of AWS instances and resources. [HTML](#) | [PDF](#) | [Kindle](#)
- **AWS Systems Manager in the AWS CLI Reference** – Describes the AWS CLI commands that you can use to automate systems management tasks. [HTML](#)
- **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. [HTML](#) | [PDF](#)

138.0 Service Definition – AWS Transfer for SFTP (AWS SFTP)

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

138.1 Service Overview

AWS Transfer for SFTP (AWS SFTP) is a fully managed service that enables the transfer of files directly into and out of Amazon S3 using the Secure File Transfer Protocol (SFTP), also known as Secure Shell (SSH) File Transfer Protocol. AWS helps you seamlessly migrate your file transfer workflows to AWS SFTP by integrating with

existing authentication systems and providing DNS routing with Amazon Route 53, so nothing changes for your customers and partners or their applications.

Top features include:

- **Fully managed highly available infrastructure** – AWS transparently operates and manages all of the compute, storage, and other infrastructure necessary to maintain high availability and performance for your SFTP endpoint.
- **Elastic resources** – AWS SFTP can meet the needs of your dynamic workloads with elastic compute infrastructure. Built-in automatic scaling means that you don't have to worry about provisioning additional resources as your data increases or spikes during certain periods.
- **Supports multiple user authentication methods** – AWS SFTP supports common user authentication systems, including Microsoft Active Directory and Lightweight Directory Access Protocol (LDAP). Alternatively, you can also choose to store and manage user credentials directly within the service.
- **Automated conversion of files into native Amazon S3 objects** – Store the files you exchange in an Amazon S3 data lake so you can extract business insights faster. The key piece that makes this exchange possible is AWS SFTP's conversion of files into objects, a process that preserves file metadata as object metadata. For example, with your files stored in Amazon S3, you can use Amazon Translate to make process documents more legible for international audiences. Or you can use Amazon Comprehend to extract relationships and insights from text files or even use Amazon Athena to query CSV files to analyse historical data.
- **Simple user experience** – An intuitive user interface and API makes it simple for you to configure your SFTP 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 behaviour to continue sharing data with you in the cloud.

Top benefits include:

- **No SFTP servers to manage** – You no longer have to purchase and run your own SFTP servers and storage to securely exchange data with partners and customers. AWS SFTP manages the infrastructure behind your SFTP endpoint for you, which includes automatic scaling capacity and maintaining high availability with a Multi-AZ architecture.
- **Seamless SFTP migrations** – AWS SFTP is fully compatible with the SFTP standard, connects to your Active Directory, LDAP, and other identity systems, and works with Amazon Route 53 DNS routing. For you, this means you can migrate SFTP-based workflows to AWS without changing your existing authentication systems, domain, and hostnames. Your external customers and partners can continue to exchange files with you, without changing their applications, processes, client software configurations, or behaviour.

- **Natively integrated with Amazon S3 for API-driven processes** – Because AWS SFTP stores your data in Amazon S3, you can use it easily with API-driven data processing and analytics workflows. AWS SFTP is natively integrated with AWS management services, simplifying your security, monitoring, and auditing operations.
- **Integration** – As AWS SFTP stores the files securely in an Amazon S3 bucket, you can use additional AWS Cloud services to extract business insights faster. For example, you can use Amazon Translate to make process documents more legible for international audiences, or you can use Amazon Comprehend to extract relationships and insights from text files or even use Amazon Athena to query CSV files to analyse historical data.
- **Familiar and comprehensive AWS management services** – With AWS SFTP, you can use IAM for security and identity management and Amazon CloudWatch for monitoring and event triggers to start post-upload processing. You can use AWS KMS or Amazon S3 server-side encryption to control encryption at rest with the service. Additionally, AWS CloudTrail helps you meet compliance requirements with granular auditing of user and API activity.
- **Availability** – AWS SFTP provides you with a fully managed, highly available SFTP service.
- **Cost** – With AWS SFTP you only pay for the resources you use: the time your SFTP server is provisioned and the amount of data it transfers. There are no upfront costs and no resources to manage yourself.

138.2 Backup/Restore and Disaster Recovery

This requirement is not applicable to AWS SFTP. For additional information beyond what is described herein, please refer to https://aws.amazon.com/backup-restore/services/#Amazon_Simple_Storage_Service .28Amazon S3.29.

138.3 Pricing Overview

Please see the AWS UK G Cloud 12 Pricing Document affiliated with this service in the Digital Marketplace.

138.4 Service Constraints

This requirement is not applicable to AWS SFTP. For additional information beyond what is described herein, please refer to <https://docs.aws.amazon.com/transfer/latest/userguide/what-is-aws-transfer-for-sftp.html>.

138.5 Technical Requirements

Please see <https://docs.aws.amazon.com/transfer/latest/userguide/what-is-aws-transfer-for-sftp.html> for comprehensive technical documentation regarding AWS SFTP.

139.0 Service Definition – AWS Transit Gateway

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

139.1 Service Overview

AWS Transit Gateway is a service that enables customers to connect their Amazon VPCs and their on-premises networks to a single gateway.

Top features include:

- Centralised routing policies across VPCs and on premises
- Scales to support thousands of VPCs across multiple accounts
- Increase connectivity throughput with multi-VPN connections.
- Flexible segmentation and routing rules
- Horizontally scalable
- Simplified management

Top benefits include:

- **Routing** – AWS Transit Gateways supports dynamic and static layer 3 routing between Amazon 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.
- **Easier connectivity** – Simplify how you interconnect all of your of VPCs across thousands of AWS accounts and into your on-premises networks. You can easily and quickly connect into a single centrally managed gateway, rapidly growing the size of your network.
- **Build a global network** – As your network grows to support more users in different parts of the world, you will need to scale AWS Cloud services within your network. With AWS Transit Gateway, you can easily share AWS Cloud services, such as DNS, Active Directory, and IPS/IDS, across all of your Amazon VPCs.
- **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.
- **Better visibility and control** – You can segment your network by creating multiple route tables in an AWS Transit Gateway and associate Amazon VPCs and VPNs to them. This will allow you to create isolated networks inside an AWS Transit Gateway, similar to virtual routing and forwarding (VRFs) in traditional networks. The AWS Transit Gateway will have a default route table. The use of

multiple route tables is optional.

- **On-demand bandwidth** – Expand your network quickly to get the bandwidth you need to transfer large amounts of data for your applications or to enable your migration to the cloud. Quickly add Amazon VPCs to your network without having to provision additional connections from your on-premises networks to AWS.
- **Security** – AWS Transit Gateway is integrated with 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.
- **Management** – You can use the 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.
- **Metrics and 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 see 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.
- **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 EFS in other Amazon VPCs that are also attached to the AWS Transit Gateway.

139.2 Backup/Restore and Disaster Recovery

This requirement is not applicable to AWS Transit Gateway. 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 12 Pricing Document affiliated with this service in the Digital Marketplace.

139.4 Service Constraints

Please see <https://docs.aws.amazon.com/vpc/latest/tgw/transit-gateway-limits.html> for more information.

139.5 Technical Requirements

Please see <https://docs.aws.amazon.com/vpc/latest/tgw/what-is-transit-gateway.html> for comprehensive technical documentation regarding AWS Transit Gateway.

140.0 Service Definition – AWS Trusted Advisor

AWS Trusted Advisor is an online tool that provides you real time guidance to help you provision your resources following AWS best practices. Trusted Advisor checks help optimize your AWS infrastructure, increase security and performance, reduce your overall costs, and monitor service limits. Whether establishing new workflows, developing applications, or as part of ongoing improvement, take advantage of the recommendations provided by Trusted Advisor on a regular basis to help keep your solutions provisioned optimally.

AWS Basic Support and AWS Developer Support customers get access to 6 security checks (S3 Bucket Permissions, Security Groups - Specific Ports Unrestricted, IAM Use, MFA on Root Account, EBS Public Snapshots, RDS Public Snapshots) and 50 service limit checks. AWS Business Support and AWS Enterprise Support customers get access to all 115 Trusted Advisor checks (14 cost optimization, 17 security, 24 fault tolerance, 10 performance, and 50 service limits) and recommendations.

140.1 Service Overview

Top features include:

- **Email notifications.** The Trusted Advisor notification feature helps you stay up-to-date with a summary of status across your checks. You will be notified by weekly email when you opt in for this service, and it is totally free.
- **Recent changes.** You can track recent changes or check status on the console dashboard. The most recent changes appear at the top of the list to bring them to your attention. Available with AWS Business Support and AWS Enterprise Support.
- **Exclude items.** The "exclude items" allows you to customize the Trusted Advisor report. You can exclude items from the check result if they are not relevant; the excluded items appear separately, and you can restore (include) them at any time.
- **Action links.** Items in a Trusted Advisor report have hyperlinks to the AWS Management Console, where you can take action on the Trusted Advisor recommendations.
- **Access management.** You can use AWS Identity and Access Management (IAM) to control access to specific checks or check categories.
- **AWS Support API.** You can refresh and retrieve Trusted Advisor results programmatically using AWS Support API. Available with AWS Business Support and AWS Enterprise Support.

- **Refresh.** You can refresh individual checks or refresh all the checks at once by clicking the Refresh All button in the top-right corner of the summary dashboard. The minimum refresh interval varies based on the check.
- **Amazon CloudWatch integration.** You can use Amazon CloudWatch Events to detect and react to changes in the status of Trusted Advisor checks. Available with AWS Business Support and AWS Enterprise Support.

Top benefits include:

- **Cost optimization.** AWS Trusted Advisor can save you money on AWS by eliminating unused and idle resources or by making commitments to reserved capacity.
- **Performance.** AWS Trusted Advisor can improve the performance of your service by checking your service limits, ensuring you take advantage of provisioned throughput, and monitoring for overutilized instances.
- **Security.** AWS Trusted Advisor can improve the security of your application by closing gaps, enabling various AWS security features, and examining your permissions.
- **Fault tolerance.** AWS Trusted Advisor can increase the availability and redundancy of your AWS application by take advantage of auto scaling, health checks, multi AZ, and backup capabilities.
- **Service limits.** AWS Trusted Advisor checks for service usage that is more than 80% of the service limit. Values are based on a snapshot, so your current usage might differ. Limit and usage data can take up to 24 hours to reflect any changes.

140.2 Backup/Restore and Disaster Recovery

This requirement is not applicable for AWS Trusted Advisor. For additional information beyond what is described herein, please refer to <http://aws.amazon.com/documentation/>.

The AWS global infrastructure is built around AWS Regions and Availability Zones. AWS Regions provide multiple physically separated and isolated 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 centre infrastructures.

For more information about AWS Regions and Availability Zones, see AWS Global Infrastructure (<http://aws.amazon.com/about-aws/global-infrastructure/>).

140.3 Pricing Overview

Please see the AWS UK G Cloud 12 Pricing Document affiliated with this service in the Digital Marketplace.

140.4 Service Constraints

Please see <https://docs.aws.amazon.com/whitepapers/latest/cost-optimization-reservation-models/aws-trusted-advisor.html> for more information regarding the constraints in place when using AWS Trusted Advisor

140.5 Technical Requirements

Please refer to <https://docs.aws.amazon.com/whitepapers/latest/cost-optimization-reservation-models/aws-trusted-advisor.html> for comprehensive technical documentation regarding AWS Trusted Advisor.

141.0 Service Definition – AWS Web Application Firewall (AWS WAF)

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

141.1 Service Overview

AWS WAF is a web application firewall that helps protect your web applications from common web exploits that could affect application availability, compromise security, or consume excessive resources. AWS WAF gives you control over which traffic to allow or block to your web applications by defining customizable web security rules. You can use AWS WAF to create custom rules that block common attack patterns, such as SQL injection or cross-site scripting, and rules that are designed for your specific application. New rules can be deployed within minutes, letting you respond quickly to changing traffic patterns. In addition, AWS WAF includes a full-featured API that you can use to automate the creation, deployment, and maintenance of web security rules.

Top benefits include:

- **Increased Protection Against Web Attacks** – AWS WAF protects web applications from attacks by filtering traffic based on rules that you create. For example, you can filter web requests based on IP addresses, HTTP headers, HTTP body, or Uniform Resource Identifier (URI) strings, which allows you to block common attack patterns, such as SQL injection or cross-site scripting.
- **Security Integrated with How You Develop Applications** – Every feature in AWS WAF can be configured using either the AWS WAF API or the AWS Management Console. This allows you to define application-specific rules that increase web security as you develop your application. This lets you put web

security at multiple points in the development chain, from the hands of the developer initially writing code, to the DevOps engineer deploying software, to the security experts conducting an audit.

- **Ease of Deployment and Maintenance** – AWS WAF is easy to deploy, protects any application deployed on Amazon CloudFront content delivery service, and there is no additional software to deploy. You can centrally define your rules and reuse them across all the web applications that you need to protect.
- **Improved Web Traffic Visibility** – You can set up AWS WAF to just monitor requests that match your filter criteria. AWS WAF gives near real-time visibility into your web traffic, which you can use to create new rules or alerts in Amazon CloudWatch.
- **Cost-Effective Web Application Protection** – With AWS WAF you pay only for what you use. AWS WAF provides a customizable, self-service offering, and pricing is based on how many rules you deploy and how many web requests your web application receives. There are no minimum fees and no up-front commitments.

141.2 Backup/Restore and Disaster Recovery

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

141.1 Pricing Overview

Please see the AWS UK G Cloud 12 Pricing Document affiliated with this service in the Digital Marketplace.

141.2 Service Constraints

Please see <http://aws.amazon.com/waf/> for more information.

141.3 Technical Requirements

AWS WAF is a web application firewall service that lets you monitor web requests for Amazon CloudFront distributions and restrict access to your content. Use AWS WAF to block or allow requests based on conditions that you specify, such as the IP addresses that requests originate from or values in the requests.

- **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. [HTML](#) | [PDF](#)
- **API Reference** – Describes all the API operations for AWS WAF in detail. [HTML](#) | [PDF](#)

142.0 Service Definition – AWS Well-Architected Tool

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

142.1 Service Overview

The AWS Well-Architected Tool helps you review the state of your workloads and compares them to the latest AWS architectural best practices. The tool is based on the AWS Well-Architected Framework, developed to help cloud architects build secure, high-performing, resilient, and efficient application infrastructure. This framework provides a consistent approach for customers and partners to evaluate architectures, has been used in tens of thousands of workload reviews conducted by the AWS solutions architecture team, and provides guidance to help implement designs that scale with application needs over time.

Top features include:

- Provides a plan on how to architect for the cloud using established best practices.
- Free tool used to define your workloads and review against operational excellence, security, reliability, performance efficiency, and cost optimisation.
- Delivers an action plan with step-by-step guidance on how to build better workloads for the cloud.

Top benefits include:

- Free architectural guidance and access to knowledge and best practices used by AWS architects.
- Single tool and consistent process to help you review and measure your cloud architectures. Use the results that the tool provides to identify next steps for improvement, drive architectural decisions, and bring architecture considerations into your corporate governance process.
- Enables you to monitor the overall status of your workloads and helps you to understand potential risks.
- Supports continuous improvement throughout the workload lifecycle. The tool makes it easy to save point-in-time milestones and track changes to your workload. Whenever you want, initiate a new review process to help ensure your architecture continues to improve over time.

142.2 Backup/Restore and Disaster Recovery

This requirement is not applicable to the AWS Well-Architected Tool. 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 12 Pricing Document affiliated with this service in the Digital Marketplace.

142.4 Service Constraints

This requirement is not applicable to the AWS Well-Architected Tool. For additional information beyond what is described herein, please refer to <http://aws.amazon.com/documentation>.

142.5 Technical Requirements

Please refer to <https://aws.amazon.com/well-architected-tool/> and <https://aws.amazon.com/well-architected-tool/faqs/> for comprehensive technical documentation regarding the AWS Well-Architected Tool.

143.0 Service Definition – AWS X-Ray

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

143.1 Service Overview

AWS X-Ray helps developers analyse 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 analyse both applications in development and in production, from simple three-tier applications to complex microservices applications consisting of thousands of services.

Top benefits include:

- **Review Request Behaviour** - AWS X-Ray traces user requests as they travel through your entire application. It aggregates the data generated by the individual services and resources that make up your application, providing you an end-to-end view of how your application is performing.
- **Discover Application Issues** - With AWS X-Ray, you can glean insights into how your application is performing and discover root causes. With X-Ray's tracing features, you can follow request paths to pinpoint where in your application and what is causing performance issues. X-Ray provides annotations so you can append metadata to traces, making it possible to tag and filter trace data so you can discover patterns and diagnose issues.
- **Improve Application Performance** - AWS X-Ray helps you identify performance bottlenecks. X-Ray's service maps let you see relationships between services and resources in your application in real time. You can quickly

see where high latencies are occurring in your application and then drill down into the specific services and paths that are performing poorly.

- **Ready to Use with AWS** - AWS X-Ray works with [Amazon EC2](#), [Amazon EC2 Container Service](#) (Amazon ECS), and [AWS Elastic Beanstalk](#). You can use XRay with applications written in Java, Node.js, and .NET that are deployed on these services. Support for [AWS Lambda](#) is coming soon.
- **Designed for a Variety of Applications** - AWS X-Ray works for both simple and complex applications, either in development or in production. You can analyse simple asynchronous event calls, three-tier web applications, or complex microservices applications consisting of thousands of services. With X-Ray, you can trace requests made to applications that span multiple AWS accounts, AWS Regions, and Availability Zones.

143.2 Backup/Restore and Disaster Recovery

This requirement is AWS X-Ray. For additional information beyond what is described herein, please refer to <http://aws.amazon.com/documentation/>.

143.1 Pricing Overview

Please see the AWS UK G Cloud 12 Pricing Document affiliated with this service in the Digital Marketplace.

143.2 Service Constraints

Please see <https://aws.amazon.com/xray/> for more information.

143.3 Technical Requirements

AWS X-Ray makes it easy for developers to analyse the behaviour of their distributed applications by providing request tracing, exception collection, and profiling capabilities.

- **Developer Guide** - Describes how to get started and instrument your applications using AWS X-Ray. [HTML](#) | [PDF](#)
- **API Reference** - Describes all the API operations for AWS X-Ray in detail. Also provides details of request and response syntax and errors for the supported web services protocols. [HTML](#) | [PDF](#)

144.0 Service Definition - Service Quotas

Service Quotas is an AWS service that enables you to view and manage your quotas from a central location. Quotas, also referred to as limits, are the maximum value for your resources, actions, and items in your AWS account. Each AWS service defines its quotas and establishes default values for those quotas. Depending on your business needs, it may be necessary to increase your service quota values. The Service Quotas

console makes it easy to quickly look up available service quotas, and request increases for those quotas.

144.1 Service Overview

Top Features include

- The Service Quotas console provides quick access to the AWS default quota values for your account, across all commercial Regions. When you select a service in the Service Quotas console, you'll see the quotas and whether the quota is adjustable.
- For any adjustable service quotas, you can use Service Quotas to request a quota increase.
- If your account has been active a while and a resource has been used, you can view a graph of your quota utilization.
- For supported services, you can manage your quotas by configuring CloudWatch alarms to monitor usage and alert you to approaching quotas.

Top benefits include

- View AWS Service Quotas - When you select a service in the Service Quotas console, you'll see the quotas and whether the quota is adjustable. Applied quotas are overrides, or increases for a particular quota, over the AWS default value.
- Request a Service Quota Increase - To request a quota increase, in the console simply select the service and the specific quota, and choose Request quota increase. You can also use the API or command line interface (CLI) tools to request service quota increases.
- Control Who Manages Service Quotas - You can attach AWS Identity and Access Management (IAM) permission policies to your users, groups, and roles that grant or deny permission to manage the service quotas in your AWS account. For example, you could create a "quota administrator" who can view, manage, or request increases for any service quotas in your account. You could also attach one policy to a group whose members need the ability to manage the quotas for a specific service.

144.2 Backup/Restore and Disaster Recovery

This requirement is not applicable for Service Quotas. For additional information beyond what is described herein, please refer to <https://docs.aws.amazon.com/index.html>.

The AWS global infrastructure is built around AWS Regions and Availability Zones. AWS Regions provide multiple physically separated and isolated 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 centre infrastructures.

For more information about AWS Regions and Availability Zones, see AWS Global Infrastructure (<http://aws.amazon.com/about-aws/global-infrastructure/>).

144.3 Pricing Overview

Please see the AWS UK G Cloud 12 Pricing Document affiliated with this service in the Digital Marketplace.

144.4 Service Constraints

Please see

https://docs.aws.amazon.com/servicequotas/latest/userguide/reference_limits.html for more information.

144.5 Technical Requirements

Please refer to <https://docs.aws.amazon.com/servicequotas/index.html> and the following links for comprehensive technical documentation regarding Service Quotas.

- <https://docs.aws.amazon.com/servicequotas/latest/userguide/intro.html>
- <https://docs.aws.amazon.com/servicequotas/latest/userguide/getting-started.html>
- <https://docs.aws.amazon.com/servicequotas/latest/userguide/security.html>
- <https://docs.aws.amazon.com/servicequotas/latest/userguide/terms-concepts.html>

145.0 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.

145.1 Availability

AWS Cloud services are hosted within our global data centre footprint, allowing customers to consume services without having to build or manage facilities or equipment. AWS Cloud services are offered in separate Regions in a number of separate geographic areas. A Region is a physical location in the world where we have multiple, isolated locations known as Availability Zones that are engineered to be isolated from failures in other zones (see **Figure 1** below). Availability Zones consist of one or more discrete data centres, each with redundant power, networking, and connectivity, and housed in separate facilities.

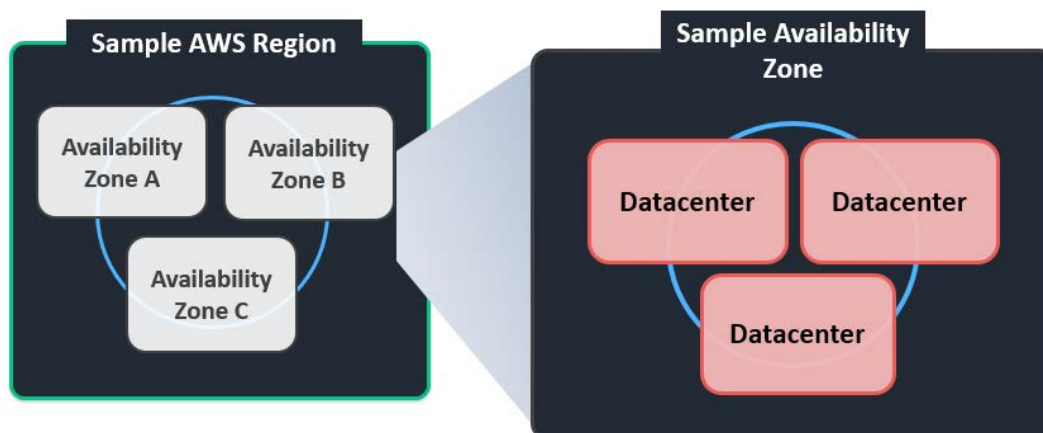


Figure 1 – Each Availability Zones can consist of multiple data centres, and at full scale can contain hundreds of thousands of servers. Every AWS Region contains 2+ zones and some Regions have as many as 6 zones

Availability Zones are engineered to be isolated from failures in other zones, and to provide inexpensive, low-latency network connectivity to other zones in the same Region. By hosting workloads in separate zones, you can protect your applications from the failure of a single location. Availability Zones offer 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.

including six Regions in the EU—United Kingdom (London), Ireland (Dublin), Italy (Milan), Germany (Frankfurt), France (Paris), and Sweden (Stockholm) —with an additional EU Region in Spain coming soon. Customers can choose to use one Region, all Regions, or any combination of Regions. Together, our EU Regions allow customers to architect highly fault tolerant applications while storing their data in the EU.

The AWS Cloud provides customers with the flexibility to run workloads and store data within multiple geographic Regions as well as across multiple Availability Zones within each Region. You decide which AWS Region(s) house your data, and it resides only in the Region(s) you specify, for as long as you choose. For example, a customer can choose to deploy their AWS Cloud services and data exclusively in the London Region, and customer content is not moved outside of London unless the customer decides to move it.

We are steadily expanding global infrastructure to help our customers achieve lower latency and higher throughput. As our customers grow their businesses, AWS will continue to provide infrastructure that meets their global requirements. The AWS products and services that are available in each Region are listed in [our Region Table](#).

Figure 2 displays our 24 global Regions and 76 Availability Zones (and one local Region in Osaka, Japan¹). Three more AWS Regions in Jakarta, Osaka, and Spain have been announced.



Figure 2 – AWS’s Global Infrastructure Consists of 24 Regions and 76 Availability Zones

To deliver content to end users with lower latency, Amazon CloudFront uses a global network of 187 Points of Presence (176 Edge Locations and 11 Regional Edge Caches) in 69 cities across 30 countries. Visit our website for a list of current [Amazon CloudFront edge locations](#).

145.1.1 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.

¹ For customers who specifically need to replicate their data or applications over greater geographic distances, there are AWS Local Regions. An AWS Local Region is a single datacentre designed to complement an existing AWS Region. Like all AWS Regions, AWS Local Regions are completely isolated from other AWS 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/>.

145.2 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, including the two discussed in the following bullets. See <http://aws.amazon.com/documentation/gettingstarted/> for more details.

- **Getting Started with AWS** – This guide provides an introduction to AWS, examples of what you can do with AWS, basic information that you need to know to get started, and links to resources and documentation that will help you learn more no matter what your use case is. The guide is available in [HTML](#), [PDF](#), and [Kindle](#) formats.
- **AWS Management Console** – This guide provides an overview of the AWS Management Console and instructions on starting various services, steps to customising the navigation bar, and tips on accessing region settings, security credentials, and billing updates. The guide is available in [HTML](#) format.

AWS allows customers to move data as needed off AWS storage using the public internet or AWS Cloud services such as AWS Direct Connect, AWS Import/Export, and more.

With AWS, you can provision compute power, storage, and other resources, gaining access to a suite of elastic IT infrastructure services as your business demands them. With minimal cost and effort, you can move your application to the AWS Cloud and reduce capital expenses, minimise support and administrative costs, and retain the performance, security, and reliability requirements your business demands. To see a step-by-step migration strategy, refer to the [Migrating Your Existing Applications to the AWS Cloud](#) whitepaper.

145.3 Service Management Details

All AWS Cloud services are driven by robust APIs that allow for a wide variety of monitoring and management tools to integrate easily with your AWS Cloud resources.

Common tools from vendors such as Microsoft, VMware, BMC Software, Okta, RightScale, Eucalyptus, CA, Xceedium, Symantec, Racemi, and Dell, to name just a few, already support AWS.

145.3.1.1 **AWS Management Console**

The AWS Management Console is a single destination for managing all of your AWS resources, from Amazon EC2 instances to Amazon DynamoDB tables. Use the console to perform any number of tasks, from deploying new applications to monitoring the health of your applications. The console enables you to manage all aspects of your AWS account, including accessing your monthly spending by service, managing

security credentials, or even setting up new IAM users. The console supports all AWS Regions and lets you provision resources across multiple regions.

145.3.1.2 ***AWS Command Line Interface (AWS CLI)***

The AWS CLI is a unified tool used to manage your AWS Cloud services. With just one tool to download and configure, you can control multiple AWS Cloud services from the command line and automate them through scripts.

145.3.1.3 ***Use Your Existing Management Tools***

It is likely that many of the tools that your organisation is using to manage your on premises environments can be integrated with AWS. Integrating your AWS environment can provide a simpler and quicker path for cloud adoption, because your operations team does not need to learn new tools or develop completely new processes.

145.4 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:

[Alexa for Business Service Level Agreement](#)

[Amazon AppStream 2.0 Service Level Agreement](#)

[Amazon API Gateway Service Level Agreement](#)

[Amazon Athena Service Level Agreement](#)

[Amazon Aurora Service Level Agreement](#)

[Amazon Chime Service Level Agreement](#)

[Amazon Cloud Directory Service Level Agreement](#)

[Amazon CloudFront Service Level Agreement](#)

[Amazon CloudSearch Service Level Agreement](#)

[Amazon CloudWatch Service Level Agreement](#)

[Amazon Cognito Service Level Agreement](#)

[Amazon Compute Service Level Agreement](#)

[Amazon Connect Service Level Agreement](#)

[Amazon DocumentDB \(with MongoDB compatibility\) Service Level Agreement](#)

[Amazon DynamoDB Service Level Agreement](#)

[Amazon EC2 Service Level Agreement](#)

[Amazon EFS Service Level Agreement](#)

[Amazon EKS Service Level Agreement](#)

[Amazon Elastic Container Registry Service Level Agreement](#)

[Amazon Elastic Load Balancing Service Level Agreement](#)
[Amazon Elastic Transcoder Service Level Agreement](#)
[Amazon ElastiCache Service Level Agreement](#)
[Amazon Elasticsearch Service - Service Level Agreement](#)
[Amazon EMR Service Level Agreement](#)
[Amazon FSx Service Level Agreement](#)
[Amazon GuardDuty Service Level Agreement](#)
[Amazon Inspector Service Level Agreement](#)
[Amazon Kinesis Service Level Agreement](#)
[Amazon Lightsail Instance and Block Storage Service Level Agreement](#)
[Amazon Lightsail Managed Databases Service Level Agreement](#)
[Amazon Macie Service Level Agreement](#)
[Amazon Machine Learning Language Service Level Agreement](#)
[Amazon Managed Blockchain Service Level Agreement](#)
[Amazon Messaging \(SQS, SNS\) Service Level Agreement](#)
[Amazon MQ Service Level Agreement](#)
[Amazon Neptune Service Level Agreement](#)
[Amazon QuickSight Service Level Agreement](#)
[Amazon RDS Service Level Agreement](#)
[Amazon Redshift Service Level Agreement](#)
[Amazon Rekognition Service Level Agreement](#)
[Amazon Route 53 Service Level Agreement](#)
[Amazon S3 Service Level Agreement](#)
[Amazon SageMaker Service Level Agreement](#)
[Amazon Simple Workflow Service Level Agreement](#)
[Amazon SimpleDB Service Level Agreement](#)
[Amazon User Engagement \(Pinpoint, SES\) Service Level Agreement](#)
[Amazon VPC NAT Gateway Service Level Agreement](#)
[Amazon WorkDocs Service Level Agreement](#)
[Amazon WorkLink Service Level Agreement](#)
[Amazon WorkMail Service Level Agreement](#)
[Amazon WorkSpaces Service Level Agreement](#)

[AWS Amplify Console Service Level Agreement](#)
[AWS AppSync Service Level Agreement](#)
[AWS Backup Service Level Agreement](#)
[AWS Budgets Service Level Agreement](#)
[AWS Certificate Manager Private Certificate Authority Service Level Agreement](#)
[AWS Client VPN Service Level Agreement](#)
[AWS CloudHSM Service Level Agreement](#)
[AWS Cloud Map Service Level Agreement](#)
[AWS CloudTrail Service Level Agreement](#)
[AWS CodeBuild Service Level Agreement](#)
[AWS CodeCommit Service Level Agreement](#)
[AWS CodeDeploy Service Level Agreement](#)
[AWS CodePipeline Service Level Agreement](#)
[AWS Config Service Level Agreement](#)
[AWS Cost Explorer API Service Level Agreement](#)
[AWS Database Migration Service Level Agreement](#)
[AWS Data Pipeline Service Level Agreement](#)
[AWS Device Farm Service Level Agreement](#)
[AWS Direct Connect Service Level Agreement](#)
[AWS Directory Service Level Agreement](#)
[AWS Elemental MediaConnect Service Level Agreement](#)
[AWS Elemental MediaConvert Service Level Agreement](#)
[AWS Elemental MediaLive Service Level Agreement](#)
[AWS Elemental MediaPackage Service Level Agreement](#)
[AWS Elemental MediaStore Service Level Agreement](#)
[AWS Elemental MediaTailor Service Level Agreement](#)
[AWS Firewall Manager Service Level Agreement](#)
[AWS GameLift Service Level Agreement](#)
[AWS Global Accelerator Service Level Agreement](#)
[AWS Glue Service Level Agreement](#)
[AWS Hybrid Storage and Data Transfer Service Level Agreement](#)
[AWS IoT 1-Click Service Level Agreement](#)

[AWS IoT Analytics Service Level Agreement](#)
[AWS IoT Core Service Level Agreement](#)
[AWS IoT Device Defender Service Level Agreement](#)
[AWS IoT Device Management Service Level Agreement](#)
[AWS IoT Greengrass Service Level Agreement](#)
[AWS Key Management Service Service Level Agreement](#)
[AWS Lambda Service Level Agreement](#)
[AWS OpsWorks Service Level Agreement](#)
[AWS PrivateLink Service Level Agreement](#)
[AWS RoboMaker Service Level Agreement](#)
[AWS Secrets Manager Service Level Agreement](#)
[AWS Security Hub Service Level Agreement](#)
[AWS Service Catalog Service Level Agreement](#)
[AWS Shield Advanced Service Level Agreement](#)
[AWS Site-to-Site VPN Service Level Agreement](#)
[AWS Step Functions Service Level Agreement](#)
[AWS Systems Manager Service Level Agreement](#)
[AWS Transit Gateway Service Level Agreement](#)
[AWS WAF Service Level Agreement](#)
[AWS X-Ray Service Level Agreement](#)

See the Supplier Terms document affiliated with this framework catalogue for additional information.

145.5 Trial Service Details

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

145.6 Data Restoration/Service Migration

Traditional enterprise backup and recovery strategies typically take an agent-based approach whereby the entire contents of a server are backed up over either the Local Area Network (LAN) or the Storage Area Network (SAN). Conventional architectures have required this approach because replacing failed components is complex, time consuming, and operationally intensive. This has, in turn, created a backup environment

that is complex to manage and resource intensive to operate, requiring technologies such as data de-duplication and virtual tape libraries to cope with ever-increasing workloads.

The AWS Cloud enables a far more lightweight approach to backup and recovery due, in part, to the following characteristics:

- Computers are now virtual abstract resources instantiated via code rather than being hardware based.
- Capacity is available at incremental cost rather than upfront cost.
- Resource provisioning takes place in minutes, lending itself to real-time configuration.
- Server images are available on demand, can be maintained by an organisation, and can be activated immediately.

These characteristics offer you opportunities to recover deleted or corrupted data, with less infrastructure overhead.

145.6.1.1 ***Protecting Configurations Rather Than Servers***

Amazon Elastic Compute Cloud (Amazon EC2) enables the backup and recovery of a standard server, such as a web server or application server, so that you can focus on protecting configuration and stateful data rather than on the server itself. This set of data is much smaller than the aggregate set of server data, which typically includes various application files, operating system files, temporary files, and so on. This change of approach means that regular nightly incremental or weekly full backups can take far less time and consume less storage space.

When a compute instance is started in Amazon EC2, it is based upon an AMI and can also connect to existing storage volumes (e.g., Amazon EBS). In addition, when launching a new instance, it is possible to pass user data to the instance so that it can be accessed internally as dynamic configuration parameters.

A sample workflow would include the following steps:

1. Launch a new instance of a web server, passing it the “identity” of the web server and any security credentials required for initial setup. The instance is based upon a prebuilt AMI that contains the operating system and relevant web server application (e.g., Apache, Internet Information Services [IIS]).
2. Upon start-up, a boot script accesses a designated and secured Amazon S3 bucket that contains the specified configuration file(s).
3. The configuration file(s) contains various instructions for setting up the server (e.g., web server parameters, locations of related servers, additional software to install, patch updates).
4. The server executes the specified configuration and is ready for service. An open-source tool for performing this process, called cloud-init, is already installed on Amazon Linux AMIs and is also available for a number of other Linux

distributions.

The figures below depict a traditional backup approach, and depicts an Amazon EC2 backup approach.




Figure 3 - Traditional Backup Approach

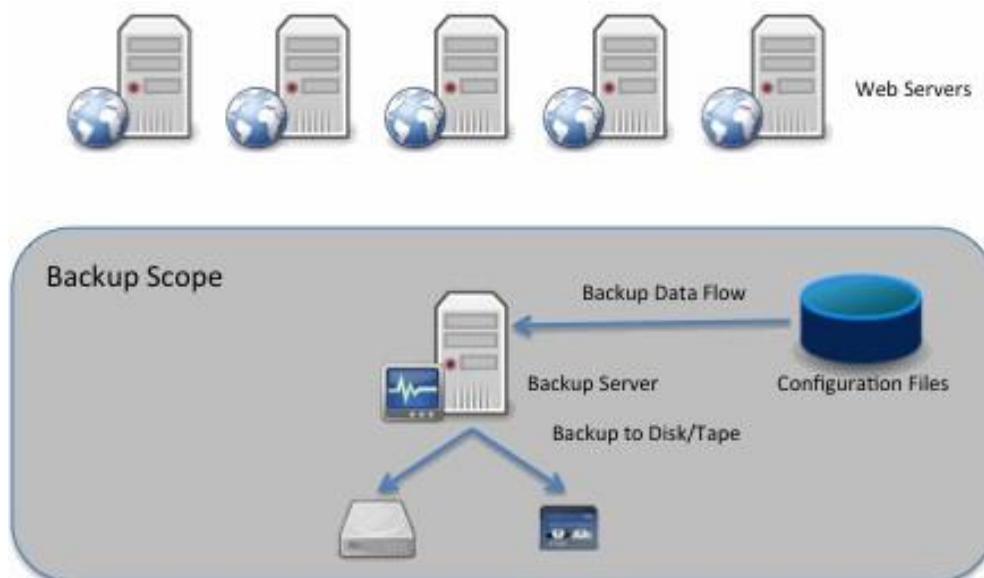


Figure 4 - Amazon EC2 Backup Approach

In the Amazon EC2 backup approach, there is no need to back up the server itself. The relevant configuration is contained in the combination of the AMI and the configuration file(s). Therefore, the only components requiring backup and recovery are the AMI and configuration file(s).

145.6.1.2 **Amazon Machine Images (AMIs)**

AMIs that you register are automatically stored in your account using Amazon EBS snapshots. These snapshots reside in Amazon S3 and are highly durable. This means that the underlying storage mechanism for the AMIs is protected from multiple failure scenarios.

It is also possible to share AMIs between separate AWS accounts. Consequently, you can create totally independent copies of the AMI by:

- Sharing the original AMI to another specified AWS account that you control.
- Starting a new instance based upon the shared AMI.
- Creating a new AMI from that running instance.
- The new AMI is then stored in the second account and is an independent copy of the original AMI. You also have the option to create multiple copies of the AMI within the same account.

145.6.1.3 **Configuration Files**

Customers use a variety of version management approaches for configuration files, and you can follow the same regime for the files used to configure your Amazon EC2 instances. For example, you could store different versions of configuration files in designated locations and securely control them like any other code. You could then back up these code repositories using the appropriate backup cycle (e.g., daily, weekly, monthly) and snapshots to protected locations. Furthermore, you could use Amazon S3 to store your configuration files, taking advantage of the durability of the service in addition to backing up the files to an alternate location on a regular basis.

145.6.1.4 **Database and File Servers**

Backing up data for database and file servers differs from the web and application layers. In general, database and file servers contain larger amounts of business data (tens of GB to multiple TB) that must be retained and protected at all times. In these cases, you can leverage efficient data movement techniques such as snapshots to create backups that are fast, reliable, and space efficient.

For databases that are built upon Redundant Array of Independent Disks (RAID) sets of Amazon EBS volumes (and have total storage less than 1 TB), an alternative backup approach is to asynchronously replicate data to another database instance built using a single Amazon EBS volume. While the destination Amazon EBS volume will have slower performance, it is not being used for data access and can be easily snapshotted to Amazon S3 using the Amazon EBS snapshot capability.

More information on AWS backup and recovery can be found at <http://aws.amazon.com/backup-storage/>. Refer to the [Storage Options on the AWS Cloud](#) whitepaper for additional information.

145.6.1.5 ***Using AWS for Disaster Recovery***

AWS provides the features and services that you can leverage for your disaster recovery processes. Using these AWS features, customers (oftentimes assisted by AWS partners) can build and operate agile and cost-effective disaster recovery solutions.

Traditional disaster recovery approaches involve the duplication of infrastructure to ensure the availability of spare capacity in a disaster scenario. This infrastructure needs to be procured, installed, and maintained so that it is ready to deal with the anticipated capacity requirements. Under normal operational circumstances, this infrastructure would typically be underutilised or over-provisioned.

With AWS, you can scale up your infrastructure on an as-needed basis, enabling faster disaster recovery of critical IT systems without incurring the infrastructure expense of a second physical site. This allows greater agility to change and optimise resources during a disaster recovery scenario. It also results in significant cost savings, because you only pay for what you use when using the highly scalable, reliable, secure, fast, and inexpensive AWS Cloud infrastructure.

The AWS Cloud supports many popular disaster recovery architectures, from “pilot light” environments that are ready to scale up at a moment’s notice to “hot standby” environments that enable rapid failover. Benefits of using AWS for disaster recovery include:

- **Fast Performance** – Readily access fast, disk-based storage and file retrieval.
- **No Tape** – Eliminate costs associated with transporting, storing, and retrieving tape media and associated tape backup software.
- **Compliance** – Avoid fines for missing compliance deadlines with fast retrieval of files.
- **Elasticity** – Quickly add any amount of data and easily expire and delete data without handling physical media.
- **Secure** – Trust our secure and durable technology platform that maintains industry-recognised certifications and audits.
- **Partners** – Connect with AWS solution providers and system integration partners to help with your deployment.

A business typically decides on an acceptable Recovery Time Objective (RTO) and Recovery Point Objective (RPO) based on the financial impact to the business when systems are unavailable. AWS can work with customers to plan disaster recovery solutions so that they cost-effectively provide system recovery based on the RPO within the timeline and service level established by the RTO. Additional information on using AWS for disaster recovery can be found at <http://aws.amazon.com/disaster-recovery/>.

We also recommend reviewing the [Using AWS for Disaster Recovery](#) whitepaper for additional information.

145.6.2 Service Migration

AWS migration services include:

- **AWS Application Discovery Service** – A service that helps you quickly and reliably plan application migration projects by automatically identifying applications running in on-premises data centres, their associated dependencies, and their performance profile.
- **AWS Database Migration Service (AWS DMS)** – A service that helps you migrate databases to AWS easily and securely. The source database remains fully operational during the migration, minimizing downtime to applications that rely on the database. AWS DMS can migrate your data to and from most widely used commercial and open-source databases.
- **AWS Server Migration Service (AWS SMS)** – An agentless service that makes it easier and faster for you to migrate thousands of on-premises workloads to AWS. AWS SMS allows you to automate, schedule, and track incremental replications of live server volumes, making it easier for you to coordinate largescale server migrations.
- **AWS Snowball Edge** – A 100 TB data transfer device with on-board storage and compute capabilities. You can use Snowball Edge to move large amounts of data into and out of AWS, as a temporary storage tier for large local datasets, or to support local workloads in remote or offline locations.

145.7 Customer Responsibilities

As cloud computing customers are building systems on top of cloud infrastructure, the security and compliance responsibilities are shared between the Cloud Service Provider (CSP) and cloud customers. In an Infrastructure as a Service (IaaS) model, customers control how they architect and secure their applications and data put on the infrastructure, while CSPs are responsible for providing services on a highly secure and controlled platform and providing a wide array of additional security features. The level of CSP and customer responsibilities in this shared responsibility model depends on the cloud deployment model (see the [NIST Definition of Cloud Computing](#) models). Customers should be clear as to their responsibilities in each model. AWS's shared responsibility model is depicted in **the Figure 5** below.

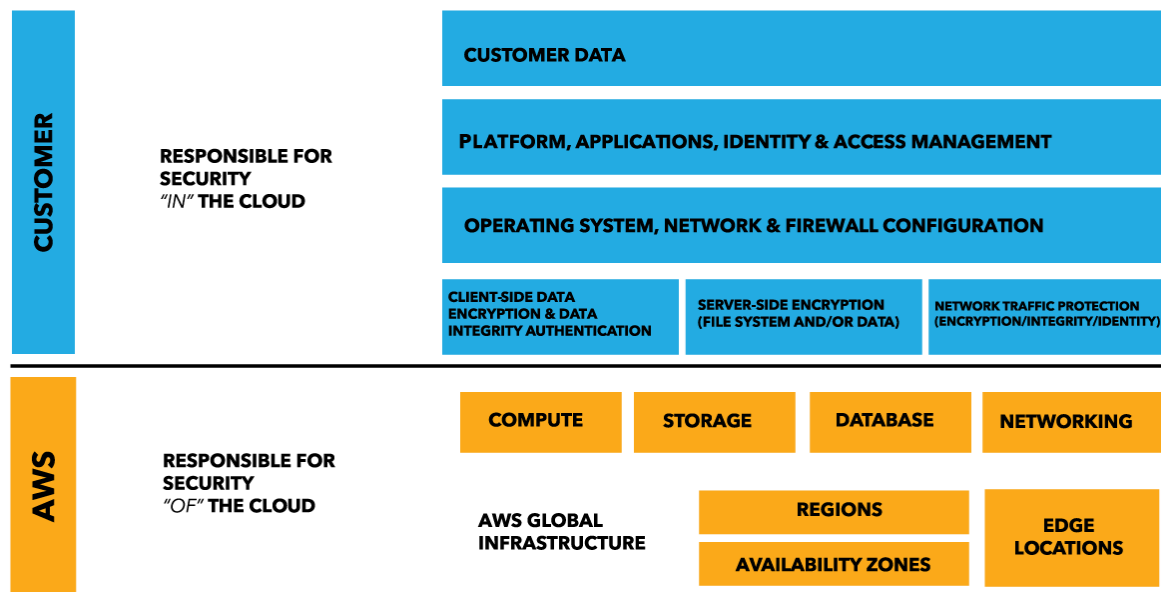


Figure 5 - AWS Shared Responsibility Model

AWS Responsibility – AWS operates, manages, and controls the infrastructure components, from the host operating system and virtualisation layer down to the physical security of the facilities in which the services operate.

Customer/Partner Responsibility – Customers/partners assume responsibility and management of the guest operating system (including updates and security patches), other associated application software, configuration of the AWS-provided security group firewalls, and other security, change management, and logging features.

AWS's shared responsibility model is further explained on the [AWS Compliance](#) webpage. AWS does not access customer data, and customers are given the choice as to how they store, manage, and protect their data. There are four important basics regarding data ownership and management in the shared responsibility model:

- Customers continue to own their data.
- Customers choose the geographic location(s) in which to store their data—it does not move unless the customer decides to move it.
- Customers can download or delete their data whenever they like.
- Customers should consider the sensitivity of their data and decide if and how to encrypt the data while it is in transit and at rest.