Table of Contents

1.0 Service Definition – Amazon Elastic Compute Cloud (Amazon EC2) 7
1.1 Service Overview 7
1.1.1 Instance types 8
1.2 Backup/Restore and Disaster Recovery 22
1.3 Pricing Overview 23
1.4 Service Constraints 23
1.5 Technical Requirements 23
2.0 Service Definition – AWS Web Application Firewall (AWS WAF) 24
2.1 Service Overview 24
2.2 Backup/Restore and Disaster Recovery 25
2.1 Pricing Overview 25
2.2 Service Constraints 25
2.3 Technical Requirements 25
3.0 Service Definition – AWS Server Migration Service (SMS) 26
3.1 Service Overview 26
3.2 Backup/Restore and Disaster Recovery 26
3.3 Pricing Overview 26
3.3.1 Additional Charges 26
3.4 Service Constraints 27
3.5 Technical Requirements 27
4.0 Service Definition – AWS Identity and Access Management (IAM) 28
4.1 Service Overview 28
4.2 Backup/Restore and Disaster Recovery 29
4.3 Pricing Overview 29
4.4 Service Constraints 29
4.5 Technical Requirements 29
5.0 Service Definition – AWS Directory Service 30
5.1 Service Overview 30
5.2 Backup/Restore and Disaster Recovery 31
5.3 Pricing Overview 31
5.4 Service Constraints 31
5.5 Technical Requirements 31
6.0 Service Definition – AWS Direct Connect 32

6.1 Service Overview 32
6.2 Backup/Restore and Disaster Recovery 33
6.3 Pricing Overview 33
6.4 Service Constraints 33
6.5 Technical Requirements 33
7.0 Service Definition – AWS Database Migration Service 34
7.1 Service Overview 34
7.2 Backup/Restore and Disaster Recovery 35
7.3 Pricing Overview 35
7.4 Service Constraints 35
7.5 Technical Requirements 35
8.0 Service Definition – AWS CloudFormation 36
8.1 Service Overview 36
8.2 Backup/Restore and Disaster Recovery 37
8.3 Pricing Overview 37
8.4 Service Constraints 37
8.5 Technical Requirements 37
9.0 Service Definition – Amazon WorkSpaces 39
9.1 Service Overview 39
9.2 Backup/Restore and Disaster Recovery 40
9.3 Pricing Overview 40
9.4 Service Constraints 40
9.5 Technical Requirements 40
10.0 Service Definition – Amazon Virtual Private Cloud (Amazon VPC) 41
10.1 Service Overview 41
10.2 Backup/Restore and Disaster Recovery 41
10.3 Pricing Overview 42
10.4 Service Constraints 42
10.5 Technical Requirements 42
11.0 Service Definition – AWS Storage Gateway 43
11.1 Service Overview 43
11.2 Backup/Restore and Disaster Recovery 44
11.3 Pricing Overview 44
11.4 Service Constraints 44
11.5 Technical Requirements 45

12.0 Service Definition – Amazon Simple Storage Service (Amazon S3) 46
12.1 Service Overview 46
12.2 Backup/Restore and Disaster Recovery 47
12.3 Pricing Overview 47
12.4 Service Constraints 48
12.5 Technical Requirements 48
13.0 Service Definition – Amazon Route 53 49
13.1 Service Overview 49
13.2 Backup/Restore and Disaster Recovery 51
13.3 Pricing Overview 51
13.4 Service Constraints 51
13.5 Technical Requirements 51
13.6 Service Overview – Amazon Route 53 Resolver 51
13.6.1 Backup/Restore and Disaster Recovery 53
13.6.2 Pricing Overview 53
13.6.3 Service Constraints 53
13.6.4 Technical Requirements 53
14.0 Service Definition – Amazon Relational Database Service (Amazon RDS) 54
14.1 Service Overview 54
14.2 Backup/Restore and Disaster Recovery 54
14.2.1 Database Instance Backups 55
14.2.2 Automated Backup 55
14.2.3 Point-In-Time Recovery 56
14.2.4 Automated Backups with Unsupported MySQL Storage Engines 56
14.2.5 DB Snapshots 56
14.3 Pricing Overview 56
14.4 Service Constraints 56
14.5 Technical Requirements 57
15.0 Service Definition – Amazon Inspector 58
15.1 Service Overview 58
15.2 Backup/Restore and Disaster Recovery 59
15.3 Pricing Overview 59
15.4 Service Constraints 59
15.5 Technical Requirements 59
16.0 Service Definition – Amazon Glacier 60

16.1 Service Overview ... 60
16.2 Backup/Restore and Disaster Recovery ... 60
16.3 Pricing Overview ... 60
16.4 Service Constraints ... 60
16.5 Technical Requirements ... 61
17.0 Service Definition – Amazon Elastic Block Store (Amazon EBS) ... 62
17.1 Service Overview ... 62
17.2 Backup/Restore and Disaster Recovery ... 62
17.3 Pricing Overview ... 63
17.4 Service Constraints ... 63
17.5 Technical Requirements ... 63
18.0 Service Definition – Amazon CloudWatch ... 64
18.1 Service Overview ... 64
18.2 Backup/Restore and Disaster Recovery ... 66
18.3 Pricing Overview ... 66
18.4 Service Constraints ... 66
18.5 Technical Requirements ... 67
19.0 Service Definition – AWS Backup ... 68
19.1 Service Overview ... 68
19.2 Backup/Restore and Disaster Recovery ... 69
19.3 Pricing Overview ... 69
19.4 Service Constraints ... 69
19.5 Technical Requirements ... 69
20.0 Service Definition – AWS Security Hub ... 70
20.1 Service Overview ... 70
20.2 Backup/Restore and Disaster Recovery ... 71
20.3 Pricing Overview ... 71
20.4 Service Constraints ... 71
20.5 Technical Requirements ... 71
21.0 Service Definition – Amazon Simple Queue Service (Amazon SQS) ... 72
21.1 Service Overview ... 72
21.2 Backup/Restore and Disaster Recovery ... 72
21.3 Pricing Overview ... 72
21.4 Service Constraints ... 72
21.5 Technical Requirements ... 73

22.0 Service Definition – Amazon GuardDuty 74
22.1 Service Overview 74
22.2 Backup/Restore and Disaster Recovery 75
22.3 Pricing Overview 75
22.4 Service Constraints 75
22.5 Technical Requirements 75
23.0 Service Definition – Amazon Simple Notification Service (Amazon SNS) 76
23.1 Service Overview 76
23.2 Backup/Restore and Disaster Recovery 76
23.3 Pricing Overview 76
23.4 Service Constraints 77
23.5 Technical Requirements 77
24.0 Service Definition – AWS Config 78
24.1 Service Overview 78
24.2 Backup/Restore and Disaster Recovery 78
24.3 Pricing Overview 78
24.4 Service Constraints 78
24.5 Technical Requirements 79
25.0 Service Definition – AWS CloudTrail 80
25.1 Service Overview 80
25.2 Backup/Restore and Disaster Recovery 81
25.3 Pricing Overview 81
25.4 Service Constraints 81
25.5 Technical Requirements 81
26.0 Service Definition – AWS Key Management Service (AWS KMS) 83
26.1 Service Overview 83
26.2 Backup/Restore and Disaster Recovery 84
26.3 Pricing Overview 84
26.4 Service Constraints 84
26.5 Technical Requirements 84
27.0 Service Definition – AWS Organizations 85
27.1 Service Overview 85
27.2 Backup/Restore and Disaster Recovery 86
27.3 Pricing Overview 86
27.4 Service Constraints 86

27.5 Technical Requirements 86
28.0 Service Definition – Amazon Elastic File System (Amazon EFS) 87
28.1 Service Overview 87
28.2 Backup/Restore and Disaster Recovery 88
28.3 Pricing Overview 88
28.4 Service Constraints 88
28.5 Technical Requirements 88
29.0 Service Definition – Amazon CloudFront 89
29.1 Service Overview 89
29.2 Backup/Restore and Disaster Recovery 90
29.3 Pricing Overview 90
29.4 Service Constraints 90
29.5 Technical Requirements 90

1.0 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

1.1 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.

1.1.1 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.

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

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

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

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

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

- 100 Gbps of network bandwidth on the largest instance size
- Requires HVM AMIs that include drivers for ENA and NVMe
- Powered by the AWS Nitro System, a combination of dedicated hardware and lightweight hypervisor
- Instance storage offered via EBS or NVMe SSDs that are physically attached to the host server
- With M5dn instances, local NVMe-based SSDs are physically connected to the host server and provide block-level storage that is coupled to the lifetime of the M5 instance

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

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

1.1.1.2 *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

1.1.1.3 *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

1.1.1.4 *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

1.1.1.5 *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

1.1.1.6 *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.

1.1.1.7 *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-xxtb1.metal), and up to 24TB RAM (u-24tb1.metal).

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

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

1.2 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.

1.3 Pricing Overview

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

1.4 Service Constraints

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

1.5 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.

2.0 Service Definition – AWS Web Application Firewall (AWS WAF)

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

2.1 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.

2.2 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/.

2.1 Pricing Overview

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

2.2 Service Constraints

Please see http://aws.amazon.com/waf/ for more information.

2.3 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

3.0 Service Definition – AWS Server Migration Service (SMS)

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

3.1 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.

3.2 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/

3.3 Pricing Overview

You can use the AWS Server Migration Service to migrate your on-premises workloads to AWS at no charge.

3.3.1 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.

3.4 Service Constraints

Please see https://aws.amazon.com/server-migration-service/faqs/ for more information on service constraints for AWS Server Migration Service (SMS).

3.5 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

4.0 Service Definition – AWS Identity and Access Management (IAM)

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

4.1 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.

4.2 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/.

4.3 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.

4.4 Service Constraints

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

4.5 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.

5.0 Service Definition – AWS Directory Service

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

5.1 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.

5.2 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/.

5.3 Pricing Overview

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

5.4 Service Constraints

Please see http://aws.amazon.com/directoryservice/ for more information.

5.5 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

6.0 Service Definition – AWS Direct Connect

The following subsections provide service definition information.

6.1 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.

6.2 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.

6.3 Pricing Overview

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

6.4 Service Constraints

Please see http://aws.amazon.com/documentation/direct-connect/ for more information.

6.5 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.

7.0 Service Definition – AWS Database Migration Service

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

7.1 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.

7.2 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/.

7.3 Pricing Overview

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

7.4 Service Constraints

Please see https://aws.amazon.com/dms/ for more information.

7.5 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

8.0 Service Definition – AWS CloudFormation

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

8.1 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.

8.2 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/.

8.3 Pricing Overview

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

8.4 Service Constraints

Please see https://aws.amazon.com/cloudformation/ for more information.

8.5 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

9.0 Service Definition – Amazon WorkSpaces

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

9.1 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.

9.2 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/.

9.3 Pricing Overview

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

9.4 Service Constraints

Please see https://aws.amazon.com/workspaces/ for more information.

9.5 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

10.0 Service Definition – Amazon Virtual Private Cloud (Amazon VPC)

The following subsections provide service definition information.

10.1 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.

10.2 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.

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/vpc/ for more information.

10.5 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.

11.0 Service Definition – AWS Storage Gateway

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

11.1 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.

11.2 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.

11.3 Pricing Overview

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

11.4 Service Constraints

Please see https://docs.aws.amazon.com/storagegateway/latest/userguide/Performance.html for more information.

11.5 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.

12.0 Service Definition – Amazon Simple Storage Service (Amazon S3)

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

12.1 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.

12.2 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.

12.3 Pricing Overview

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

12.4 Service Constraints

Please see http://aws.amazon.com/documentation/s3/ for more information.

12.5 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.

13.0 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

13.1 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.

13.2 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/.

13.3 Pricing Overview

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

13.4 Service Constraints

Please see https://aws.amazon.com/route53/ for more information.

13.5 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

13.6 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.

13.6.1 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/.

13.6.2 Pricing Overview

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

13.6.3 Service Constraints

Please see https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/DNSLimitations.html for more information.

13.6.4 Technical Requirements

Please see https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/resolver-getting-started.html for comprehensive technical documentation regarding Route 53 Resolver.

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

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

14.1 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.

14.2 Backup/Restore and Disaster Recovery

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

14.2.1 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.

14.2.2 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 8hour 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.

14.2.3 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.

14.2.4 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.

14.2.5 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.

14.3 Pricing Overview

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

14.4 Service Constraints

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

14.5 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.

15.0 Service Definition – Amazon Inspector

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

15.1 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.

15.2 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/.

15.3 Pricing Overview

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

15.4 Service Constraints

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

15.5 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

16.0 Service Definition – Amazon Glacier

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

16.1 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).

16.2 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.

16.3 Pricing Overview

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

16.4 Service Constraints

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

16.5 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.

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

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

17.1 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.

17.2 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.

17.3 Pricing Overview

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

17.4 Service Constraints

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

17.5 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

18.0 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

18.1 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.

18.2 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/.

18.3 Pricing Overview

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

18.4 Service Constraints

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

18.5 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

19.0 Service Definition – AWS Backup

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

19.1 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.

19.2 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.

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/aws-backup/latest/devguide/aws-backup-limits.html for more information.

19.5 Technical Requirements

Please refer to https://docs.aws.amazon.com/aws-backup/latest/devguide/whatisbackup.html for comprehensive technical documentation regarding AWS Backup.

20.0 Service Definition – AWS Security Hub

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

20.1 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.

20.2 Backup/Restore and Disaster Recovery

Findings are stored for 90 days within AWS Security Hub. You can export findings at any time.

20.3 Pricing Overview

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

20.4 Service Constraints

Please see https://docs.aws.amazon.com/securityhub/latest/userguide/securityhub_limits.html for more information.

20.5 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.

21.0 Service Definition – Amazon Simple Queue Service (Amazon SQS)

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

21.1 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.

21.2 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/.

21.3 Pricing Overview

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

21.4 Service Constraints

Please see https://aws.amazon.com/sqs/ for more information.

21.5 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

22.0 Service Definition – Amazon GuardDuty

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

22.1 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..

22.2 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/.

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://aws.amazon.com/guardduty/resources/ for more information.

22.5 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

23.0 Service Definition – Amazon Simple Notification Service (Amazon SNS)

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

23.1 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.

23.2 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.

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/sns/ for more information.

23.5 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

24.0 Service Definition – AWS Config

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

24.1 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.

24.2 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.

24.3 Pricing Overview

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

24.4 Service Constraints

Please see https://aws.amazon.com/config/ for more information.

24.5 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

25.0 Service Definition – AWS CloudTrail

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

25.1 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.

25.2 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.

25.3 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.

25.4 Service Constraints

Please see https://aws.amazon.com/cloudtrail/ for more information.

25.5 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

26.0 Service Definition – AWS Key Management Service (AWS KMS)

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

26.1 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.

26.2 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.

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://aws.amazon.com/kms/ for more information.

26.5 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

27.0 Service Definition – AWS Organizations

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

27.1 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 Organisations 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.

27.2 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/.

27.3 Pricing Overview

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

27.4 Service Constraints

Please see https://aws.amazon.com/organizations/ for more information.

27.5 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.

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

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

28.1 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.

28.2 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/.

28.3 Pricing Overview

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

28.4 Service Constraints

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

28.5 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

29.0 Service Definition – Amazon CloudFront

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

29.1 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.

29.2 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/.

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/cloudfront/ for more information.

29.5 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