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

May 2022

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

Table of Contents

1. Introduction

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

1.1. How to use the AWS Service Definition Documents

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

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

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

Notwithstanding that AWS has combined its service descriptions into a consolidated document for ease of review by Customers, to access the options through a Call-Off Contract the Customer must reference each individual Digital Marketplace Service ID within the Call-Off Contract in order to enable that service as an option that can be procured under their G-Cloud 13 Call-Off Contract. AWS recommends that Buyers list all of the Digital Marketplace Service ID's for every service described in this document in its Call-Off Contract to enable the option to switch between Services flexibly during the term. For a list of all AWS Digital Marketplace Service ID's, please contact an AWS account representative through aws-cloud@amazon.com.

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

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

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

2. AWS Security Assurance

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

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

British Standard 7858:2019

Buyers selecting AWS Services and expressly requiring AWS conformity to BS7858:2019 acknowledge that AWS scopes BS7858:2019 compliance to those AWS employees with physical access to the ‘data layer’ zones within datacentres and those who are directed by the Buyer to access Buyer Data such as Technical Account Managers (“TAMS”). A list of TAMs shall be provided to the Buyer by the Supplier prior to the Start date of the Call-Off Contract and the Buyer shall only contact the listed TAMs in relation to Buyer Data during the Term of the Call-Off Contract. Buyers are obliged in accordance with the Call-Off Contract to encrypt Buyer Data when using AWS Services. Buyer should note that the Supplier does not include Supplier Staff (as defined in the Call-Off Contract) responsible for operating the AWS Services or those with logical access to encrypted Buyer Data for the purposes of its BS7858:2019 compliance.

2.1. Information Assurance

The following subsections provide information relating to information assurance.

2.1.1. ISO 27001 Certification

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

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

2.1.2. NCSC UK Cloud Security Principles

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

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

2.2. GDPR and processing of Personal Data

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

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

3. AWS Mainframe Modernization

3.1. Service Overview

AWS Mainframe Modernization offers application intelligence, knowledge, and analysis for migration teams and developers to better understand large application portfolios.

3.1.1. Features

- **Analyzer:** Analyzer can help you assess, scope, and even plan a project. Once you've migrated mainframe workloads, the tool can also assist with application maintenance and ongoing modernization strategy by helping you evaluate change impacts, reduce risks, and strategize on enhancements or refactoring.
- **Developer:** AWS Mainframe Modernization offers an on-demand integrated development environment (IDE) so developers can write code quicker with smart editing and debugging, instant code compilation, and unit testing. You can also use the IDE to migrate mainframe application code, and to develop and enhance enterprise applications running on the Mainframe Modernization Managed Runtime.
- **Managed Runtime:** The AWS Mainframe Modernization managed execution environment continually monitors your clusters to keep enterprise workloads running with self-healing compute and automated scaling, so you can focus on application development. Mainframe Modernization Managed Runtime is built for business-critical enterprise applications, and offers high availability, reliability, and security.
- **Refactoring:** AWS Mainframe Modernization offers both automated and manual refactoring capabilities to accelerate the modernization of mainframe and legacy assets. You can use refactoring to convert legacy application programming languages, to create macroservices or microservices, and to modernize user interfaces (UIs) and application software stacks.
- **Continuous Integration and Delivery (CI/CD):** CI/CD introduces automation to boost agility and release velocity across every stage of the development and test pipeline. AWS Mainframe Modernization helps application development teams deliver code changes more frequently and reliably, which accelerates migration speed, increases quality, and helps reduce time-to-market for releasing new business functions.

3.1.2. Benefits

- **Migrate and Modernize:** Easily migrate and modernize your applications to eliminate the hardware and staffing costs of traditional mainframes.
- **Manage end-to-end migration:** Break up and manage your end-to-end migration with infrastructure, software, and tools to refactor and transform legacy applications.
- **Operate migrated applications:** Deploy, run, and operate migrated applications in the Mainframe Modernization environment with no upfront costs.

3.2. Backup/Restore and Disaster Recovery

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

3.3. Pricing Overview

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

3.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/m2/latest/userguide/what-is-m2.html
- **Service quotas:** https://docs.aws.amazon.com/general/latest/gr/m2.html
- **Service FAQs:** https://aws.amazon.com/mainframe-modernization/faqs/?nc=sn&loc=6&refid=ps_a134p000006gb2oaau&trkcampaign=acq_paid_search_brand

3.5. Technical Requirements

Please refer to https://docs.aws.amazon.com/m2/latest/userguide/what-is-m2.html and the following links for comprehensive technical documentation regarding this service.

- **User Guide:** Describes all AWS Mainframe Modernization concepts and provides instructions on using the various features with both the console and the command line interface.
- **API Reference**: Describes all the API operations for AWS Mainframe Modernization in detail.

4. AWS IoT TwinMaker

4.1. Service Overview

AWS IoT TwinMaker makes it easier for developers to create digital twins of real-world systems such as buildings, factories, industrial equipment, and production lines. AWS IoT TwinMaker provides the tools you need to build digital twins to help you optimize building operations, increase production output, and improve equipment performance. With the ability to use existing data from multiple sources, create virtual representations of any physical environment, and combine existing 3D models with real-world data, you can now harness digital twins to create a holistic view of your operations faster and with less effort.

4.1.1. Features

- **Data connectors:** AWS IoT TwinMaker provides built-in data connectors for the following AWS services: AWS IoT SiteWise for collecting, organizing, and storing equipment and time-series sensor data; Amazon Kinesis Video Streams for capturing, processing, and storing video data; and Amazon Simple Storage Service (S3) for storing visual resources (for example, CAD files) and data from enterprise applications. AWS IoT TwinMaker also provides a framework for you to easily create custom data connectors to use with other AWS or third-party data sources, such as Amazon Timestream, Snowflake, and Siemens Mindsphere. These data connectors allow your applications to only use the AWS IoT TwinMaker unified data access API to read from and write to the different data stores without needing to query each data source using their own individual API.
- **Model builder:** To model your physical environment, you can create entities in AWS IoT TwinMaker that are virtual representations of your physical systems, such as a furnace or an assembly line. You can also specify custom relationships between these entities to accurately represent the real-world deployment of these systems. You then connect these entities to your various data stores to form a digital twin graph, which is a knowledge graph that structures and organizes information about the digital twin for

easier access and understanding. As you build out this model of your physical environment, AWS IoT TwinMaker automatically creates and updates the digital twin graph by organizing the relationship information in a graph database.

- **Scene composer:** With AWS IoT TwinMaker, you build a 3D digital twin by using your existing and previously built 3D visual models, such as CAD files, Building Information Modeling (BIM) files, or point cloud scans. Using the AWS IoT TwinMaker scene composer and simple 3D tools, you import these visual assets into a scene and position them to match your physical environment—for example a factory and its equipment. You can then add interactive video and sensor data overlays from the connected data sources, insights from connected machine learning (ML) and simulation services, and maintenance records and operational documents to provide you with a regularly updated, spatially aware visualization of your operations.
- **Applications:** Once you've created the digital twin, AWS IoT TwinMaker provides a low-code experience for building a web application so your plant operators and maintenance engineers can access and interact with the digital twin. AWS IoT TwinMaker comes with a plug-in for Grafana, a popular open-source dashboard and visualization platform from Grafana Labs. The plug-in provides custom visualization panels, including a 3D scene viewer and dashboard templates, as well as a data-source component to connect to your digital twin data, allowing you to quickly create 3D-enabled applications for your specific needs. The plug-in can also be used to build applications with Amazon Managed Grafana, which is a fully managed service for open-source Grafana.

4.1.2. Benefits

- **No need to reingest:** Use your existing IoT, video, and enterprise application data where it already lives—without needing to reingest or move the data to another location.
- **Automatically generated knowledge graph:** Save time with an automatically generated knowledge graph that binds your data sources to virtual replicas of physical systems to accurately model real-world environments.
- **Immersive view:** Get an immersive 3D view of your systems and operations to optimize efficiency, increase production, and improve performance.

4.2. Backup/Restore and Disaster Recovery

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

4.3. Pricing Overview

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

4.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/iot-twinmaker/
- **Service quotas:** https://docs.aws.amazon.com/general/latest/gr/iot-twinmaker.html
- **Service FAQs:** https://aws.amazon.com/iot-twinmaker/faqs/?nc=sn&loc=5&refid=ps_a134p000006gb2oaau&trkcampaign=acq_paid_search_brand

4.5. Technical Requirements

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

- User Guide: Describes key concepts of AWS IoT TwinMaker and provides instructions on how to create digital twins.
- API Reference: Describes in detail all the API operations for AWS IoT TwinMaker. Also provides sample requests, responses, and errors.

5. AWS IoT FleetWise

5.1. Service Overview

With AWS IoT FleetWise, you can easily collect and organize data from the unique data format present in your vehicles (regardless of make, model, or options) and standardize the data format for easy analysis in the cloud without having to build custom data collection systems.

AWS IoT FleetWise helps you efficiently transfer data to the cloud in near-real time using the service's intelligent filtering capabilities. Reduce the amount of data transferred to the cloud by selecting what data to transfer and defining rules and events for when to transfer it based on parameters like weather conditions, location, or vehicle type.

Once the data is in the cloud, use it for tasks like remotely diagnosing issues in individual vehicles, analysing vehicle fleet health to help protect against warranty claims and recalls, and improving autonomous driving and advanced driver assistance systems with analytics and machine learning.

5.1.1. Features

- **Rules-based data collection:** AWS IoT FleetWise applies the rules you define for transferring only high-value data signals to the cloud. First, select which data to transfer, such as safety equipment data, camera data, or any other sensor-generated data. Then, define rules and events for when to transfer that data based on parameters such as weather, location, or vehicle type. This reduces the amount of unnecessary data transferred to the cloud, which lessens costs and gives access to more useful data.
- **Vehicle modeling:** Use AWS IoT FleetWise to build virtual representations of vehicles in the cloud and apply a common data format to structure and label vehicle attributes, sensors, and signals. AWS IoT FleetWise standardizes vehicle modeling using Vehicle Signal Specification (VSS) so that a signal like "fuel pressure" is always represented as fuel_pressure and measured in pound-force per square inch (PSI) and kilopascal (kPa). Once the vehicle is modeled, upload a standard CAN database (DBC) or AUTOSAR XML (ARXML) file so AWS IoT FleetWise can read the unique and proprietary data signals sent over a vehicle's Controller Area Network Bus (CAN Bus).
- **Edge Agent:** AWS IoT FleetWise Edge Agent facilitates communication between the vehicle and the cloud. While the vehicle is on, it continuously receives data collection schemes from AWS IoT FleetWise and collects data accordingly for cloud transfer. Install Edge Agent to supported vehicle hardware—currently the NXP S32G vehicle network processors—with plans for additional supported hardware. With Edge Agent, you control every step of the process from creation to installation, and maintain full data ownership and control of proprietary information.
- **Remote configuration deployment:** With AWS IoT FleetWise, you can deploy cloud-based data collection schemes to vehicles, ensure that vehicles can receive the

schemes, and take action when they do not receive them. For example, if a vehicle temporarily loses connectivity while in an underground parking structure, AWS IoT FleetWise can resend the message at regular time intervals until the vehicle responds.

- **Global signal catalog:** Select and standardize unique vehicle sensors and signals from a centralized repository of all your vehicle models.
- **Data engine:** AWS IoT FleetWise enriches collected vehicle data with metadata and vehicle attributes. For example, if you collect seatbelt data signals during hard braking events, AWS IoT FleetWise can enrich the data by appending the vehicle model and door count, making it easier to analyze the event data in the cloud.

5.1.2. Benefits

- **No need for custom data collection systems:** Access standardized fleet-wide vehicle data without the need to develop custom data collection systems.
- **Efficiency:** Reduce costs and enable more efficient data transfer with intelligent filtering that sends the exact data you need to the cloud.
- **Mitigate issues faster:** Surface vehicle health data in near-real time to detect and mitigate issues faster, help prevent potential recalls, and remotely assist customers.

5.2. Backup/Restore and Disaster Recovery

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

5.3. Pricing Overview

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

5.4. Service Constraints

Information on service features and constraints is available online here:

- **Service documentation:** https://docs.aws.amazon.com/iot-fleetwise/?id=docs_gateway
- **Service FAQs:** https://aws.amazon.com/iot-fleetwise/faqs/?nc=sn&loc=5&refid=ps_a134p000006gb2oaau&trkcampaign=acq_paid_search_brand

5.5. Technical Requirements

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

- Developer Guide: Describes key concepts of AWS IoT FleetWise and provides instructions for using the features of AWS IoT FleetWise.
- API Reference: Describes all the API operations for AWS IoT FleetWise in detail. Also provides sample requests, responses, and errors for the supported web services protocols.

6. Cross-Service Definitions

The following service definition topics are applicable to all AWS Service Offerings and are detailed once in a cross-service manner below.

6.1. Availability

AWS has the largest global infrastructure footprint of any provider, and this footprint is constantly increasing at a significant rate. The AWS Cloud infrastructure is built around AWS Regions and Availability Zones. A Region is a physical location in the world where we have multiple Availability Zones. Availability Zones consist of one or more discrete data centres, each with redundant power, networking, and connectivity and housed in separate facilities. These Availability Zones offer customers the ability to operate production applications and databases that are more highly available, fault tolerant, and scalable than would be possible with a single data centre.

AWS currently has 25 Regions and 81 Availability Zones throughout the world—including 14 Local Zones and 17 Wavelength Zones for ultralow latency applications. AWS Regions include: US East (N. Virginia), US East (Ohio), US West (Oregon), US West (N. California), AWS GovCloud (US-West), AWS GovCloud (US-East), Canada (Central), Europe (Ireland), Europe (Frankfurt), Europe (London), Europe (Milan), Europe (Paris), Europe (Stockholm), Middle East (Bahrain), Africa (Cape Town), Asia Pacific (Singapore), Asia Pacific (Tokyo), Asia Pacific (Sydney), Asia Pacific (Seoul), Asia Pacific (Mumbai), Asia Pacific (Osaka), Asia Pacific (Hong Kong), South America (Sao Paulo), China (Beijing), and China (Ningxia).

AWS has also announced plans for nine more AWS Regions in Australia, Canada, India, Indonesia, Israel, New Zealand, Spain, Switzerland, and United Arab Emirates (UAE).

Information about each Region can be found at the AWS Global Infrastructure page.

Figure 1 depicts the current AWS Regions and Availability Zones, along with the nine announced Regions.

Figure 1. AWS Regions and Availability Zones

The AWS products and services that are available in each Region are listed at the Region Table webpage.

Figure 2 illustrates the relationship between AWS Regions and Availability Zones.

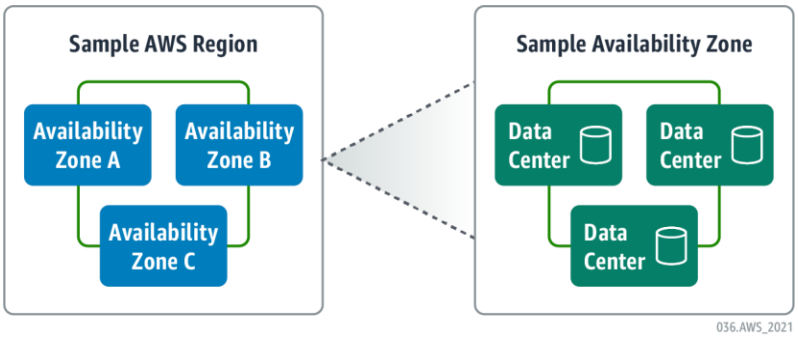

Figure 2. Relationship Between AWS Regions and Availability Zones

6.1.1. Region Availability

Exact service availability depends on a range of factors and choices made by customers when they architect and implement their solution.

The Services Offerings will be delivered from the AWS Region selected by the customer upon opening an AWS account. The customer may specify the AWS Region in which customer content will be stored. It is the customer's responsibility to select the relevant AWS Region in order to comply with its own security and governance requirements. AWS will not access or use customer content except as necessary to maintain or provide the Service Offerings, or as required by law or regulation. Customers acknowledge that AWS does not limit customers to any particular AWS Region. Note that not all AWS Cloud services are available in every AWS

Region; however, we are steadily expanding our service availability across AWS's global regions.

The full list of available AWS services, and their availability by region can be seen on our website at https://aws.amazon.com/about-aws/global-infrastructure/regional-product-services/.

6.1.2. Designing for Availability and Reliability

While AWS goes to great lengths to provide availability and reliability of the cloud, our customers share responsibility for ensuring availability and reliability within the cloud. Some best practices we recommend for building highly resilient systems in the AWS Cloud include designing for failure, automating failover and recovery, testing your recovery procedures, and accessing resources and reference architectures.

6.1.2.1. Design for Failure across Multiple Availability Zones

Although rare, failures can occur that affect the availability of resources that are hosted in the same Availability Zone. If you host all your resources in a single Availability Zone that is affected by such a failure, none of these resources would be available. It is therefore a best practice to architect across multiple Availability Zones in the same Region to achieve extremely high recovery time objectives (RTOs), recovery point objectives (RPOs), and service availability. Availability Zones are connected to each other with fast, private fibre-optic networking, enabling you to easily architect applications that automatically fail over between zones without interruption.

For mission-critical applications, it is a best practice to architect across Regions to handle the rare case of an entire Region failing—perhaps as a result of a major physical attack. You can do so using both private, high-speed networking and public internet connections to provide an additional layer of business continuity or to provide low-latency access across the globe.

6.1.2.2. Automate Failover and Recovery

By monitoring a system for key performance indicators (KPIs), you can trigger automation when a threshold is breached. These KPIs should be a measure of business value and not of the technical aspects of the service operation. This allows for automatic notification and tracking of failures, and automated recovery processes that work around or repair the failure. With sophisticated automation, it is possible to anticipate and remediate failures before they occur.

6.1.2.3. AWS Personal Health Dashboard

AWS Personal Health Dashboard provides alerts and remediation guidance when AWS is experiencing events that may impact you. The dashboard displays relevant and timely information to help you manage events in progress, and provides proactive notification to help you plan for scheduled activities. Alerts are triggered by changes in the health of AWS resources, giving you event visibility and guidance to help quickly diagnose and resolve issues.

The dashboard provides forward-looking notifications, and you can set up alerts across multiple channels, including email and mobile, so you receive timely and relevant information to help plan for scheduled changes that may affect you. In the event of AWS hardware maintenance activities that may impact one of your Amazon Elastic Compute Cloud (Amazon EC2) instances, for example, you would receive an alert with information to help you plan for, and proactively address, any issues associated with the upcoming change.

Personal Health Dashboard can integrate with Amazon CloudWatch Events, enabling you to build custom rules and select targets such as AWS Lambda functions to define automated remediation actions. The AWS Health API, which powers Personal Health Dashboard, allows

you to integrate health data and notifications with your existing in-house or third-party IT management tools.

6.1.3. Test Your Recovery Procedures

One of the benefits of the cloud is that you can test how your system fails, and you can validate your recovery procedures. You can use a test environment to simulate different failures or recreate scenarios that led to failures already. This exposes failure pathways that you can test and fix before a real failure scenario, reducing the risk of components that have not been tested before failing.

6.1.4. Resources and Reference Architecture

The resources described below help customers to understand AWS Cloud services and features and provide architectural guidance on designing and implementing systems that run on the AWS infrastructure:

- The AWS Well-Architected Framework codifies the experiences of thousands of customers, helping them assess and improve their cloud-based architectures and mitigate disruptions.
- The AWS Architecture Center is designed to provide customers with the necessary guidance and application architecture best practices to build highly scalable and reliable applications in the AWS Cloud.
- The AWS Outposts High Availability Design and Architecture Considerations whitepaper discusses architecture considerations and recommended practices to build highly available on-premises application environments with AWS Outposts.
- The AWS advanced continuous delivery best practices video.

6.2. On-Boarding/Off-Boarding Processes and Service Migration

AWS maintains a cadre of Getting Started Guides and schedules regular webinars. These guides and webinars cover a variety of topics, see http://aws.amazon.com/documentation/gettingstarted/ for more details.

AWS allows customers to move data as needed off AWS storage using the public internet or AWS Cloud services such as AWS Direct Connect, AWS Import/Export, and more.

With AWS, you can provision compute power, storage, and other resources, gaining access to a suite of elastic IT infrastructure services as your business demands them. With minimal cost and effort, you can move your application to the AWS Cloud and reduce capital expenses, minimise support and administrative costs, and retain the performance, security, and reliability requirements your business demands. To see a step-by-step migration strategy, refer to the Migrating Your Existing Applications to the AWS Cloud whitepaper.

6.3. Service Management Details

AWS Cloud services are driven by robust APIs that allow for a wide variety of monitoring, management and developer tools to integrate easily with AWS Cloud resources. Common tools from vendors such as Microsoft, VMware, BMC Software, Okta, RightScale, Eucalyptus, CA, Xceedium, Symantec, Racemi, and Dell are supported on AWS. This flexibility allows AWS customers to easily provision, manage, and monitor all of their IT resources through a “single pane of glass” with the tool that best fits their unique needs. This also means that a full

inventory of those resources is only a few clicks away. Below are various AWS-native management options.

6.3.1. AWS Management Console

The AWS Management Console is a single destination for managing all AWS resources, from Amazon Elastic Compute Cloud (Amazon EC2) instances to Amazon DynamoDB tables. Customers can use the AWS Management Console to perform any number of tasks, from deploying new applications to monitoring the health of applications. The AWS Management Console also enables customers to manage all aspects of their AWS account, including accessing monthly spending by service, managing security credentials, or even setting up new AWS Identity and Access Management (AWS IAM) users. The AWS Management Console supports all AWS Regions and lets customers provision resources across multiple AWS Regions.

6.3.2. AWS Developer Tools

AWS offerings are provided with a range of supporting components like management tools, networking services, and application augmentation services, with multiple interfaces to AWS Application Programming Interface (API)-based services, including Software Development Kits (SDKs), Integrated Development Environment (IDE) toolkits, and Command Line Tools. Browse by programming language for tools to develop and manage applications on AWS at our Tools to Build on AWS page.

AWS Developer Tools help you securely store and version control your application's source code and automatically build, test, and deploy your application to AWS or your on-premises environment. They are built to work with AWS, making it easier for your team to get set up and be productive.

AWS Developer Tools are designed to help you build software like Amazon. They facilitate practices such as continuous delivery and infrastructure as code for serverless, containers, and Amazon EC2.

6.3.3. Management and Governance Tools

In the past, organizations have had to choose between innovating faster and maintaining control over cost, compliance, and security. With AWS Management and Governance services, customers don’t have to choose between innovation and control—they can have both. With AWS, customers can enable, provision, and operate their environment for both business agility and governance control.

- **Scale** – AWS Management and Governance services are built to manage highly dynamic cloud resources at massive scale.
- **Simplicity** – AWS reduces complexity, offering a single control plane for customers to manage and govern their resources on AWS and on-premises.
- **Third-party solutions** – AWS offers the broadest partner ecosystem for customers to extend and augment their management and governance system.
- **Cost savings** – Customers can use AWS Management and Governance services to assess their resource utilization and identify ways to reduce costs.

6.4. Service Levels and Service Credits

AWS currently provides SLAs, with a corresponding Service Credit regime, for several products. Due to the rapidly evolving nature of AWS's product offerings, SLAs are best reviewed directly on our website via the links below:

https://aws.amazon.com/api-gateway/sla/

https://aws.amazon.com/appstream2/sla/

https://aws.amazon.com/athena/sla/

https://aws.amazon.com/rds/aurora/sla/

https://aws.amazon.com/braket/sla/

https://aws.amazon.com/chime/sla/

https://aws.amazon.com/cloud-directory/sla/

https://aws.amazon.com/cloudfront/sla/

https://aws.amazon.com/cloudsearch/sla/

https://aws.amazon.com/cloudwatch/sla/

https://aws.amazon.com/cognito/sla/

https://aws.amazon.com/compute/sla/

https://aws.amazon.com/connect/sla/

https://aws.amazon.com/it/detective/sla/

https://aws.amazon.com/devops-guru/sla/

https://aws.amazon.com/documentdb/sla/

https://aws.amazon.com/dynamodb/sla/

https://aws.amazon.com/ecs/anywhere/sla/

https://aws.amazon.com/efs/sla/

https://aws.amazon.com/ecr/sla/

https://aws.amazon.com/eks/sla/

https://aws.amazon.com/elasticloadbalancing/sla/

https://aws.amazon.com/elastictranscoder/sla/

https://aws.amazon.com/emr/sla/

https://aws.amazon.com/eventbridge/sla/

https://aws.amazon.com/finspace/sla/

https://aws.amazon.com/forecast/sla/

https://aws.amazon.com/fraud-detector/sla/

https://aws.amazon.com/fsx/sla/

https://aws.amazon.com/guardduty/sla/

https://aws.amazon.com/healthlake/sla/

https://aws.amazon.com/inspector/sla/
https://aws.amazon.com/ivs/sla/
https://aws.amazon.com/kendra/sla/
https://aws.amazon.com/keyspaces/sla/
https://aws.amazon.com/kinesis/sla/
https://aws.amazon.com/lightsail/sla-lightsail-instances-and-block-storage/
https://aws.amazon.com/lightsail/sla-lightsail-managed-databases/
https://aws.amazon.com/location/sla/
https://aws.amazon.com/lookout-for-equipment/sla/
https://aws.amazon.com/machine-learning/language/sla/
https://aws.amazon.com/macie/sla/
https://aws.amazon.com/managed-blockchain/sla/
https://aws.amazon.com/grafana/sla/
https://aws.amazon.com/msk/sla/
https://aws.amazon.com/managed-workflows-for-apache-airflow/sla/
https://aws.amazon.com/memorydb/sla/
https://aws.amazon.com/messaging/sla/
https://aws.amazon.com/monitron/sla/
https://aws.amazon.com/amazon-mq/sla/
https://aws.amazon.com/neptune/sla/
https://aws.amazon.com/nimble-studio/sla/
https://aws.amazon.com/elasticsearch-service/sla/
https://aws.amazon.com/personalize/sla/
https://aws.amazon.com/qldb/sla/
https://aws.amazon.com/quicksight/sla/
https://aws.amazon.com/rds/proxy/sla/
https://aws.amazon.com/redshift/sla/
https://aws.amazon.com/rekognition/sla/
https://aws.amazon.com/rds/sla/
https://aws.amazon.com/route53/sla/
https://aws.amazon.com/s3/sla-rtc/
https://aws.amazon.com/sagemaker/sla/
https://aws.amazon.com/s3/sla/
https://aws.amazon.com/swf/sla/

https://aws.amazon.com/simpledb/sla/
https://aws.amazon.com/textract/sla/
https://aws.amazon.com/timestream/sla/
https://aws.amazon.com/pinpoint/sla/
https://aws.amazon.com/vpc/ipam-sla/
https://aws.amazon.com/vpc/sla/
https://aws.amazon.com/workdocs/sla/
https://aws.amazon.com/worklink/amazon-worklink-service-level-agreement/
https://aws.amazon.com/workmail/amazon-workmail-service-level-agreement/
https://aws.amazon.com/workspaces/sla/
https://aws.amazon.com/amplify/sla/
https://aws.amazon.com/application-migration-service/sla/
https://aws.amazon.com/appsync/sla/
https://aws.amazon.com/audit-manager/sla/
https://aws.amazon.com/backup/sla/
https://aws.amazon.com/aws-cost-management/aws-budgets/sla/
https://aws.amazon.com/certificate-manager/private-certificate-authority/sla/
https://aws.amazon.com/vpn/client-vpn-sla/
https://aws.amazon.com/cloud-map/sla/
https://aws.amazon.com/cloudhsm/sla/
https://aws.amazon.com/cloudtrail/sla/
https://aws.amazon.com/codeartifact/sla/
https://aws.amazon.com/codebuild/sla/
https://aws.amazon.com/codecommit/sla/
https://aws.amazon.com/codedeploy/sla/
https://aws.amazon.com/codepipeline/sla/
https://aws.amazon.com/compute-optimizer/sla/
https://aws.amazon.com/config/sla/
https://aws.amazon.com/aws-cost-management/aws-cost-explorer/sla/
https://aws.amazon.com/datapipeline/sla/
https://aws.amazon.com/dms/sla/
https://aws.amazon.com/device-farm/sla/
https://aws.amazon.com/directconnect/sla/
https://aws.amazon.com/directoryservice/sla/

https://aws.amazon.com/disaster-recovery/sla/
https://aws.amazon.com/mediaconnect/sla/
https://aws.amazon.com/mediaconvert/sla/
https://aws.amazon.com/medialive/sla/
https://aws.amazon.com/mediapackage/sla/
https://aws.amazon.com/mediastore/sla/
https://aws.amazon.com/mediatailor/sla/
https://aws.amazon.com/firewall-manager/sla/
https://aws.amazon.com/global-accelerator/sla/
https://aws.amazon.com/glue/sla/
https://aws.amazon.com/ground-station/sla/
https://aws.amazon.com/transfer/sla/
https://aws.amazon.com/iot-1-click/sla/
https://aws.amazon.com/iot-analytics/sla/
https://aws.amazon.com/iot-core/sla/
https://aws.amazon.com/iot-device-defender/sla/
https://aws.amazon.com/iot-device-management/sla/
https://aws.amazon.com/iot-events/sla/
https://aws.amazon.com/greengrass/sla/
https://aws.amazon.com/iot-sitewise/sla/
https://aws.amazon.com/iot-things-graph/sla/
https://aws.amazon.com/kms/sla/
https://aws.amazon.com/lambda/sla/
https://aws.amazon.com/migration-hub/sla/refactor-spaces/
https://aws.amazon.com/network-firewall/sla/
https://aws.amazon.com/opsworks/sla/
https://aws.amazon.com/privatelink/sla/
https://aws.amazon.com/resilience-hub/sla/
https://aws.amazon.com/robomaker/sla/
https://aws.amazon.com/secrets-manager/sla/
https://aws.amazon.com/security-hub/sla/
https://aws.amazon.com/servicecatalog/sla/
https://aws.amazon.com/shield/sla/
https://aws.amazon.com/vpn/site-to-site-vpn-sla/

https://aws.amazon.com/step-functions/sla/

https://aws.amazon.com/systems-manager/sla/

https://aws.amazon.com/transit-gateway/sla/

https://aws.amazon.com/waf/sla/

https://aws.amazon.com/xray/sla/

https://aws.amazon.com/elasticache/sla/

See the Supplier Terms document affiliated with this framework catalogue for additional information.

6.5. Trial Service Details

The AWS Free Tier provides customers the ability to explore and try out AWS services free of charge up to specified limits for each service. The Free Tier is comprised of three different types of offerings, a 12-month Free Tier, an Always Free offer, and short term trials. Services with a 12-month Free Tier allow customers to use the product for free up to specified limits for one year from the date the account was created. Services with an Always Free offer allow customers to use the product for free up to specified limits as long as they are an AWS customer. Services with a short-term trial are free to use for a specified period of time or up to a one-time limit depending on the service selected.

Details on the limits and services provided for free are detailed in each card on the Free Tier page. If your application use exceeds the free tier limits, you simply pay standard, pay-as-you-go service rates (see each service page for full pricing details). Restrictions apply; see offer terms for more details.

6.6. Data Restoration/Service Migration

6.6.1. AWS Backup

Backing up your data is an important step towards protecting your applications and ensuring that you meet your business and regulatory backup compliance requirements. Even durable resources are susceptible to threats (e.g., bugs in your application) that could cause accidental deletions or corruption. Building and managing your own backup workflows across all your applications in a compliant and consistent manner can be complex and costly. AWS Backup removes the need for costly, custom solutions or manual processes.

AWS Backup is a fully managed backup service that makes it easy to centralize and automate the back up of data across AWS Cloud services, as well as on premises using the AWS Storage Gateway. With AWS Backup, you can create automated backup policies called backup plans, such as how frequently to back up your data and how long to retain those backups. Together with AWS Organizations, AWS Backup enables you to centrally deploy data protection (backup) policies to configure, manage, and govern your backup activity across your organization's AWS accounts and AWS resources, such as:

- Amazon Elastic Compute Cloud (Amazon EC2) instances
- Amazon Elastic Block Store (Amazon EBS) volumes
- Amazon Relational Database Service (Amazon RDS) databases, including Amazon Aurora clusters
- Amazon DynamoDB tables

- Amazon Elastic File System (Amazon EFS) file systems
- Amazon FSx for Lustre
- Amazon FSx for Windows File Server
- AWS Storage Gateway volumes

You can use AWS Backup to create backup policies that automate backup schedules and retention management. As illustrated in **Figure 3**, AWS Backup provides a fully managed, policy-based solution, simplifying backup management and enabling you to meet business and regulatory compliance requirements.

Figure 3. AWS Backup provides a fully managed, policy-based backup solution

6.6.2. Disaster Recovery

Businesses and public sector organizations of all sizes are using AWS to enable faster and cheaper disaster recovery (DR) of their critical IT systems. Having your own DR site ready and on standby in the cloud, without having to pay for the IT infrastructure, makes the AWS Cloud a perfect solution for DR. With the AWS Cloud, not only can you recover quickly from a disaster and ensure business continuity while keeping costs down, you also can make it easy, secure, and reliable.

With its fault-tolerant architecture and 200+ service offerings, AWS supports many DR architectures—from those built for smaller workloads to enterprise solutions that enable rapid failover at scale. These architectures include “pilot light” environments that are ready to scale up rapidly to “hot standby” environments that fail over at a moment’s notice. As a customer, you have the flexibility to choose the right approach for your DR strategy, depending on your recovery time objective (RTO) and recovery point objective (RPO) goals and budget.

Additionally, AWS offers a business continuity solution called CloudEndure Disaster Recovery that minimizes downtime and data loss by providing fast, reliable, cloud-based DR. The solution continuously replicates applications from physical, virtual, or cloud-based infrastructure to a low-cost staging area that is automatically provisioned in any target AWS Region of your choice. During failover or testing, an up-to-date copy of an application can be spun up on demand and be fully functioning in minutes.

6.6.2.1. High Availability (HA) in the AWS Cloud

A discussion of DR is not complete without addressing high availability (HA). While DR focuses on bringing systems back online once disaster strikes, HA focuses on ensuring that there is no single point of failure in your architecture from the outset. The AWS Cloud global infrastructure,

with its construct of Regions and Availability Zones, is designed for HA. Using AWS Regions (geographically isolated components of the global infrastructure) and Availability Zones (fully fault-tolerant clusters of data centres with redundant power, networking, and connectivity), you can build inherent HA and fault tolerance into your architecture. The AWS global infrastructure allows us to deliver the highest network availability of any cloud provider, with 7x fewer downtime hours than the next largest cloud provider.[1]

Additionally, to better isolate any issues and achieve HA, you can partition applications across multiple Availability Zones in the same AWS Region. AWS control planes and the AWS Management Console are distributed across Regions and include regional application programming interface (API) endpoints. These endpoints are designed to operate securely for at least 24 hours, if isolated from the global control plane functions, without requiring customers to access the Region or its API endpoints via external networks during any isolation.

Both HA and DR rely on some of the same best practices, such as monitoring for failures, deploying to multiple locations, and automatic failover. However, HA focuses on components of the workload, whereas DR focuses on discrete copies of the entire workload. DR has different objectives from HA, measuring time to recovery after the larger scale events that qualify as disasters. You should first ensure your workload meets your availability objectives, as a highly available architecture will enable you to meet customers' needs in the event of availability impacting events. Your DR strategy requires different approaches than those for availability and should focus on deploying discrete systems to multiple locations so that you can failover the entire workload if necessary.

6.6.2.2. Benefits of Using AWS for DR

With AWS, you can eliminate the need for additional physical infrastructure, offsite data replication, and upkeep of spare capacity. Availability Zones are a key AWS feature as they make partitioning applications for high availability easy. This enables customers to protect applications from the failure of a single location, resulting in significant cost savings and increased agility to change and optimize resources during a DR scenario.

AWS provides fine-grained control and building blocks to create the appropriate DR solution in the cloud, given a customer's unique resiliency requirements, recovery objectives (RTO and RPO as shown in **Figure 4**), and budget. With AWS, customers can build highly resilient applications while taking advantage of flexible, cost-effective infrastructure solutions.

[1] Based on downtime hours from 1/1/18 to 12/31/18 pulled directly from the public service health dashboards of the major cloud providers.

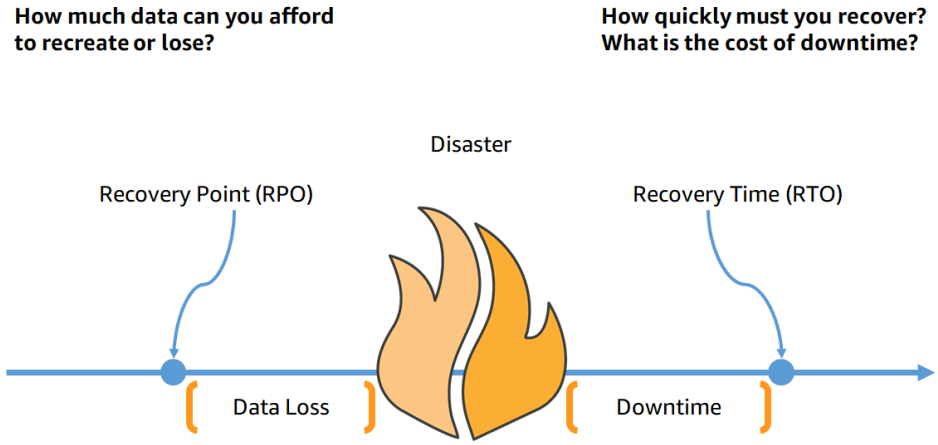

Figure 4. RPO and RTO. RPO is related to data loss in the event of disaster; RTO is related to the amount of time systems are down in the event of disaster.

Figure 5 shows a spectrum of scenarios—multi-site, warm standby, pilot light, and backup and restore—arranged by how quickly a system can be available to users after a DR event. Typically, the shorter the recovery, the higher the cost of the solution.

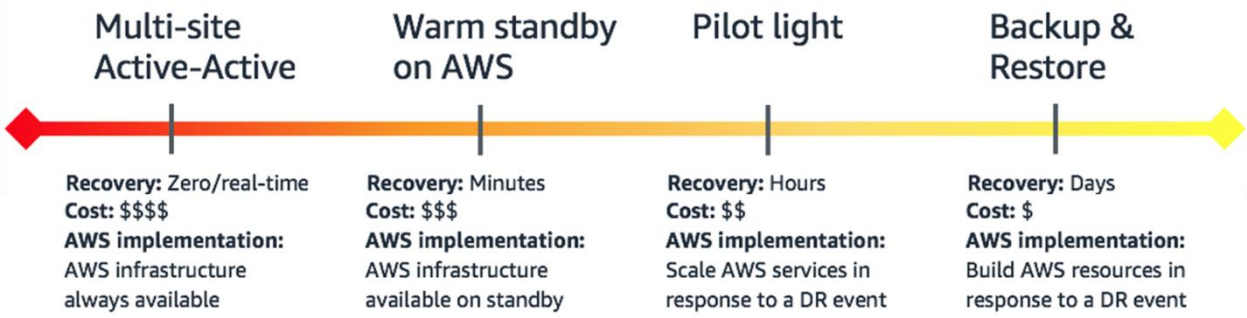

Figure 5. Spectrum of DR Options. Customers can choose a preferred DR option based on preferences for cost and the time it takes to recover systems.

Each DR option is discussed in more detail below.

Backup and Restore

In most traditional environments, data is backed up to tape and sent offsite regularly. Recovery time will be the longest using this method, and lack of automation leads to increased costs. Amazon Simple Storage Service (Amazon S3) is ideal for backup data, as it is designed to provide 99.999999999% durability of objects over a given year. Amazon S3 offers a range of storage classes—including options for one-zone infrequent access and archive—to help you save on costs. Transferring data to and from Amazon S3 is typically done via the network, and it is therefore accessible from any location. To centralize and automate your backup, AWS Backup allows you to configure backup policies and monitor backup activity.

Pilot Light for Simple Recovery into AWS Warm Standby Solution

The idea of the pilot light is an analogy that comes from gas heating. In that scenario, a small flame that's always on can quickly ignite the entire furnace to heat up a house. In this DR approach, you simply replicate part of your IT structure for a limited set of core services so that the AWS cloud environment seamlessly takes over in the event of a disaster. A small part of your infrastructure is always running simultaneously syncing mutable data (as databases or

documents), while other parts of your infrastructure are switched off and used only during testing. Unlike a backup and recovery approach, you must ensure that your most critical core elements are already configured and running in AWS (the pilot light). When the time comes for recovery, you can rapidly provision a full-scale production environment around the critical core.

Warm Standby Solution on AWS

The term “warm standby” is used to describe a DR scenario in which a scaled-down version of a fully functional environment is always running in the cloud. It further decreases recovery time because, in this case, some services are always running. By identifying business-critical systems, you could fully duplicate these systems on AWS and have them always on. These servers can be running a minimum-sized fleet of the smallest Amazon Elastic Compute Cloud (Amazon EC2) instances possible. This solution is not scaled to handle a full-production load, but it is fully functional. It can be used for non-production work, such as testing, quality assurance, and internal use. In a disaster, you can scale out the system quickly to handle the production load by adding more instances. You can automate this process using Amazon EC2 Auto Scaling and Elastic Load Balancing.

Multi-Site Solution Deployed on AWS and Onsite

A multi-site solution runs in AWS as well as on your existing on-site infrastructure, in an active-active configuration. The data replication method that you employ will be determined by the recovery point that you choose. You can use a Domain Name System (DNS) service that supports weighted routing, such as Amazon Route 53, to route production traffic to different sites that deliver the same application or service. A proportion of traffic will go to your infrastructure in AWS, and the remainder will go to your on-site infrastructure.

In an on-site disaster situation, you can adjust the DNS weighting and send all traffic to the AWS servers. The capacity of the AWS Cloud service can be rapidly increased to handle the full production load. You can use Amazon EC2 Auto Scaling to automate this process. You might need some application logic to detect the failure of the primary database services and cut over to the parallel database services running in AWS.

The cost of this scenario is determined by how much production traffic is handled by AWS during normal operation. In the recovery phase, you pay only for what you use for the duration that the DR environment is required at full scale. You can further reduce cost by purchasing Amazon EC2 Reserved Instances for your “always on” AWS servers.

6.6.2.3. AWS Elastic Disaster Recovery

AWS Elastic Disaster Recovery enables organizations to minimize downtime and data loss with fast, reliable recovery of on-premises and cloud-based applications. This service minimizes downtime and data loss with fast, reliable recovery of on-premises and cloud-based applications using affordable storage, minimal compute, and point-in-time recovery.

AWS Elastic Disaster Recovery continuously replicates your source servers to your AWS account without impacting performance. It reduces costs compared to traditional on-premises disaster recovery solutions by removing idle recovery site resources, and instead leverages affordable AWS storage and minimal compute resources to maintain ongoing replication. If you need to recover applications on AWS, you can do so within minutes.

During recovery, you can choose the most up-to-date server state as a recovery point, or choose to recover an operational copy of your applications from an earlier point in time. Point in time recovery is helpful for recovery from data corruption events such as ransomware. After

issues are resolved in your primary environment, you can use AWS Elastic Disaster Recovery to fail back your recovered applications.

Figure 6 illustrates how AWS Elastic Disaster Recovery operates.

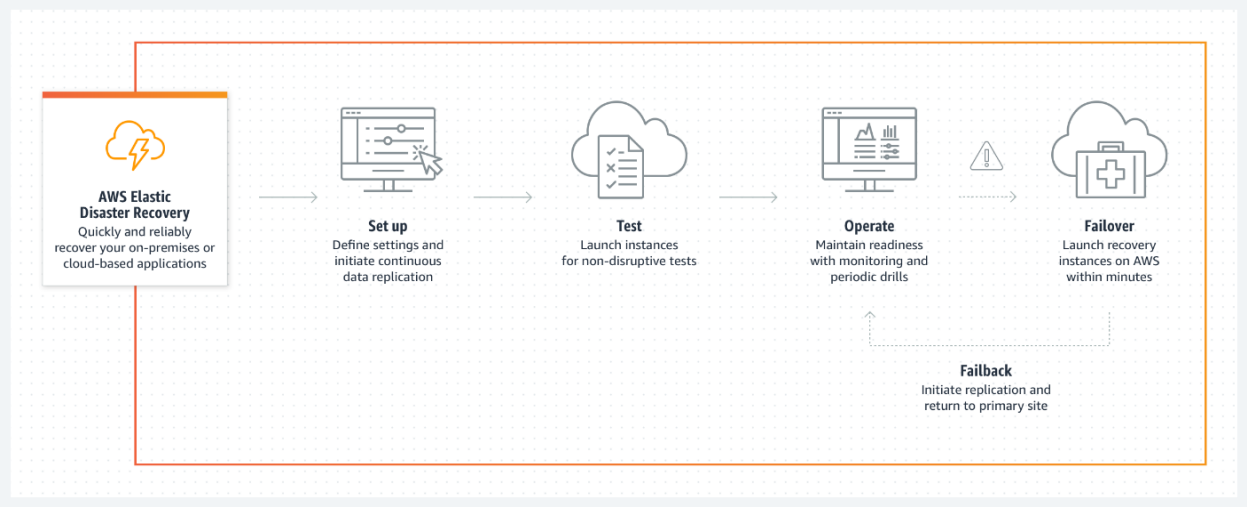

Figure 6. AWS Elastic Disaster Recovery Architecture.This service simplifies recovery of a wide range of applications on AWS and uses a unified process for drills, recovery, and failback.

AWS Elastic Disaster Recovery is the recommended service for disaster recovery to AWS. It provides similar capabilities as CloudEndure Disaster Recovery (see **Section Error! Reference source not found.**) and is operated from the AWS Management Console. This enables seamless integration between AWS DRS and other AWS services, such as AWS CloudTrail, AWS Identity and Access Management (IAM), and Amazon CloudWatch.

6.6.2.4. Complementary AWS Cloud Services for Backup/DR

In addition to CloudEndure Disaster Recovery, many other AWS Cloud services can be used to provide complementary capabilities to enable a successful DR strategy. These services include Amazon EC2, Amazon S3, AWS Storage Gateway, Amazon Relational Database Service (Amazon RDS), Amazon CloudFront, and others.

For example, Amazon RDS is a fully managed relational database service that supports automated backups by default. It can be set up using a Multi-AZ configuration to reduce risk from outages and minimize loss of valuable data. Amazon S3 is designed for 99.999999999% (11 9s) durability and stores data for millions of applications for companies around the world. Amazon EC2 and AWS Storage Gateway enable you to back up key data stores in the AWS Cloud. AWS Backup provides a common way to manage backups across several AWS Cloud services, both in the cloud and on premises. Amazon Elastic File System (Amazon EFS) and Amazon EBS provide HA and durability for file and block storage so that your backups are protected and available when needed.

More details can be found at these links:

- Learn to Build on AWS: Backup and Recovery
- AWS DR Resources
- Backup & Restore Services with AWS

6.6.2.5. DR Resources

There are multiple resources to help organizations start using AWS for a DR solution:

- Read the whitepaper Affordable Enterprise-Grade Disaster Recovery Using AWS.
- Read about CloudEndure Disaster Recovery, which offers fast, cost-effective business continuity for your mission-critical workloads.
- Read the eBook Leverage the Cloud for your Disaster Recovery Strategy.
- Download the IT Disaster Recovery Plan Checklist to ensure you are on track.

6.6.3. Resiliency Planning

An organization's resiliency and continuity plan outlines a range of disaster scenarios and the steps the organization will take to return to a regular state. Plans are written ahead of time, by key staff and multiple departments, with the goal of creating contingencies that minimize potential harm and negative impacts to the organization.

A strong resilience strategy combines both operational and cultural resilience. Operational resilience includes five pillars:

1. Remote workforce enablement
2. Constituent engagement
3. Operational continuity
4. Real-time analytics
5. Process and systems modernization

Figure 7 below illustrates the AWS Organizational Resiliency Framework and displays which core foundational technologies can assist when creating a strong resiliency plan, for each of the five pillars.

Figure 7. AWS Organizational Resiliency Framework. This framework can improve an organization's odds of minimizing the impact of an emergency on employees, customers, and partners.

A resilience plan also addresses cultural resiliency best practices that prepare your team for any disruption. Beyond pandemics or natural disasters, organizations can be adversely impacted by changing market conditions, mergers, or general restructuring, which can also be classified as emergencies. Regardless of the type of disruption, an employee's alignment to organizational objectives is most at stake during times of disruption. Other elements that influence employees during crisis include the existing organizational culture, leadership, and one's work environment—which can include technologies, tools, physical space, and processes. Your resiliency plan should identify employee risks, just as it identifies technology, customer, and financial risks.

Organizations should start by building a baseline resiliency plan. No resiliency plan is bullet-proof; however, having a plan in place drastically improves your organization's odds of minimizing the impact of an emergency on your employees, customers, and partners.

6.6.3.1. Business Continuity Plan (BCP)

Your DR plan should be a subset of your organization's business continuity plan (BCP). There is no point in maintaining aggressive DR targets for restoring a workload if that workload's business objectives cannot be achieved because of the disaster's impact on elements of your business other than your workload.

For organizations serving the public, disasters and emergencies can derail missions and critical emergency response, public safety, and public health services. Many organizations move from crisis to crisis with short-term fixes, without addressing the underlying lack of a long-term strategy for resilience and business continuity. Even with an existing resiliency plan in place in one department, the plan may not cover process re-engineering requirements and dependencies in other departments, like IT, finance, human resources, legal, communications, operators, or others.

Key disruption areas to address with a long-term BCP include the most common and high-impact disruptors such as the following:

- Ensuring IT infrastructure durability, availability, and security
- Supporting employees and tapping expertise
- Access to data for real-time situational awareness
- Lack of financial agility
- Lack of planning

Business Impact Analysis

As part of a BCP, you should carry out a business impact analysis to quantify the business impact of a disruption to your workloads. It should identify the impact on internal and external customers of not being able to use your workloads and the effect that has on your business. The analysis should help to determine how quickly the workload needs to be made available and how much data loss can be tolerated. However, it is important to note that recovery objectives should not be made in isolation; the probability of disruption and cost of recovery are key factors that help to inform the business value of providing DR for a workload.

Business impact may be time dependent. You may want to consider factoring this into your DR planning. For example, disruption to your payroll system is likely to have a very high impact to

the business just before everyone gets paid, but it may have a low impact just after everyone has already been paid.

Risk Assessment

A risk assessment of the type of disaster and geographical impact along with an overview of the technical implementation of your workload will determine the probability of disruption occurring for each type of disaster.

For highly critical workloads, you may consider high availability across multiple Regions with continuous backups in place to minimize business impact. For less critical workloads, a valid strategy may be not to have any DR in place at all. And for some disaster scenarios, it is also valid not to have any DR strategy in place as an informed decision based on a low probability of the disaster occurring. Remember that Availability Zones within an AWS Region are already designed with meaningful distance between them, and careful planning of location, such that most common disasters should only impact one zone and not the others. Therefore a multi-AZ architecture within an AWS Region may already meet your risk mitigation needs.

The cost of the DR options should be evaluated to ensure that the DR strategy provides the correct level of business value considering the business impact and risk.

With all of this information, you can document the threat, risk, impact and cost of different disaster scenarios and the associated recovery options. This information should be used to determine your recovery objectives for each of your workloads.

6.6.4. Migration

Many customers seek to migrate their workloads to the AWS Cloud to benefit from IT costs savings. These workloads include applications, websites, databases, storage, physical or virtual servers, or entire data centers—all residing in the on-premises environment, hosting facility, or other public cloud.

The AWS Cloud allows you to provision elastic compute power, storage, and other resources in the cloud on demand. Moving your applications to the AWS Cloud can help you reduce upfront expenses by using our pay-as-you-go pricing structure. You can also minimize support and administrative costs through our shared responsibility model and use of AWS Support—all while retaining performance, security, and reliability requirements. At AWS, our research and experience have led us to identify the following key drivers that compel businesses to migrate to the cloud.

- Substantial IT costs savings
- Digital transformation
- Improvements in staff productivity and business agility
- Improved security and operational resilience
- Data center consolidation
- Approaching hardware/software end-of-life
- Going global, mergers, and acquisitions
- Exploring new technologies (e.g., artificial intelligence/machine learning and Internet of Things)

Migrating to the AWS Cloud can enable your organization to reduce operational costs by up to 51% and bring products and services to market 18.8% faster. AWS has helped thousands of organizations, including enterprises such as GE, the Coca-Cola Company, BP, Enel, Samsung, NewsCorp, and Twenty-First Century Fox migrate to the cloud, freeing-up resources by lowering IT costs and improving productivity, operational resiliency, and business agility.

To ensure that you have a holistic view of the transformation initiative that is required for an effective move to the cloud and to create a foundation for your cloud strategy that returns ongoing measurable value to your organization, AWS offers the Cloud Adoption Framework. The AWS Cloud Adoption Framework enables you to analyze your environment through different perspectives: Business, People, Governance, Platform, Security, and Operations. This gives you a complete view of which areas to improve before moving forward with a large migration effort.

6.6.4.1. Migration Patterns

A strong migration plan starts with a deeper understanding of the interdependencies between applications, and it evaluates migration strategies to meet your business case objectives. **Figure 8 – Seven Common Migration** depicts the seven most common migration patterns that help you create a modern migration strategy.

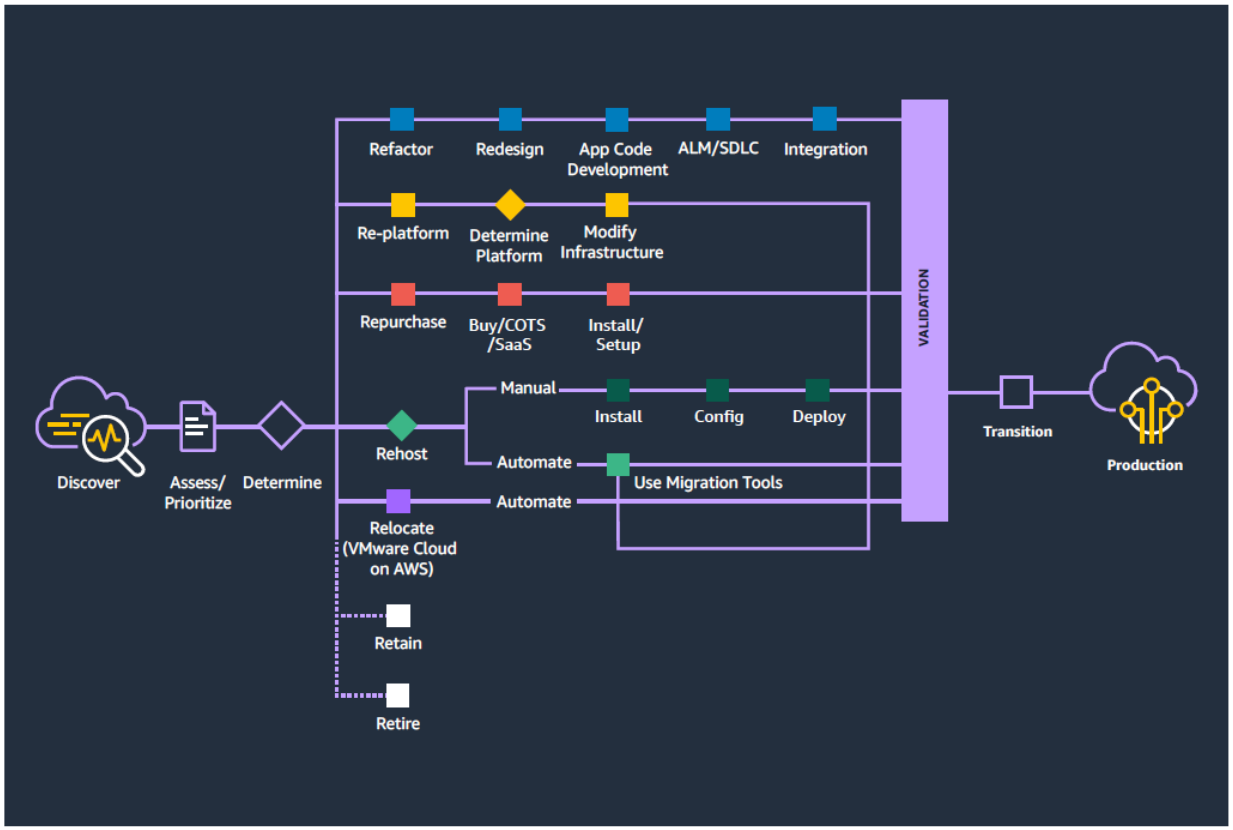

Figure 8 – Seven Common Migration Patterns

The seven migration patterns include:

1. **Rehost** (i.e., “lift and shift”) involves moving applications without changing them.

The AWS Application Migration Service (AWS MGN) enables you to rehost a large number of physical, virtual, or cloud servers to AWS without compatibility issues, performance disruption, or long cutover windows. Features such as automatic server conversion and continuous replication, combined with non-disruptive tests, ensure a smooth cutover for your most critical databases and applications, such as SAP CRM, Oracle E-Business Suite, and Microsoft SharePoint. Using AWS MGN, you can rehost applications from VMware vSphere, Microsoft Hyper-V, Amazon Elastic Compute Cloud (Amazon EC2), Amazon Virtual Private Cloud (Amazon VPC), and other clouds to AWS.

If your preferred AWS Region is not currently supported by AWS MGN, or, if the operating system on which your applications run is not currently supported by AWS MGN, you can automate rehosting using CloudEndure Migration, also designed for enterprise workloads such as SQL Server, Oracle, and SAP.

If you cannot or do not want to install an agent on your servers, you can use the AWS Server Migration Service (AWS SMS). AWS SMS offers agentless capabilities to migrate thousands of on-premises workloads to AWS. AWS SMS utilizes incremental, snapshot-based replication of the existing servers and enables cutover windows measured in hours.

2. **Replatform** (i.e., “lift, tinker, and shift”) involves making a few cloud optimizations to achieve a tangible benefit, but without changing the core architecture of the application. For example, our fully managed Amazon MQ service can easily replace a messaging broker without rewriting your applications or paying for third-party software licenses. The Amazon FSx for Windows File Server can help you migrate a Windows-based application that requires file storage.
3. **Refactor** (i.e., re-architect) enables you to re-imagine how the application is architected and developed by using cloud-native features. Refactoring is driven by a strong business need to add features, scale, or improve performance that would otherwise be difficult to achieve in the application’s existing environment. Many enterprises use the migration effort to modernize their businesses and refactor their legacy technology portfolio.
4. **Relocate** involves moving VMware vSphere®-based applications to AWS without application changes. Common migration patterns usually follow one of the other seven basic patterns, but when you migrate to AWS, you gain this option. For details, see VMware Cloud on AWS.
5. **Repurchase** (i.e., “drop and shop”) involves replacing your current environment with either a newer version of software or an entirely new solution. The AWS Marketplace offers a curated digital catalog of software solutions that support each phase of migration.
6. **Retain** involves retaining portions of your IT portfolio that you are not ready to migrate or believe are best kept on-premises. For applications that remain on-premises, AWS Outposts bring the same hardware, software, services, application programming interface (APIs), management tools, support, and operating model that is used in the AWS Cloud to your data center, co-location space, or on-premises facility.
7. **Retire** involves decommissioning or archiving portions of your IT portfolio that are no longer useful.

For more information, refer to the AWS eBook: 8 Migration Business Drivers and migration-related whitepapers available on our website.

6.7. Customer Responsibilities

Security and compliance responsibilities are shared between AWS and the customer. This shared responsibility model can help relieve customers' operational burdens as AWS operates, manages, and controls the components from the host operating system and virtualization layer down to the physical security of the facilities in which the service operates.

Customers—and in some cases, our AWS Partner Network (APN) Partners who work with those customers—control how they architect and secure their applications and data in the AWS Cloud. AWS provides a wide array of security and compliance services; a customer's responsibilities will vary depending on the services used, the integration of those services into their IT environment, and applicable laws and regulations.

As shown in **Figure 9** below, this differentiation of responsibility is commonly referred to as security "in" the cloud versus security "of" the cloud.

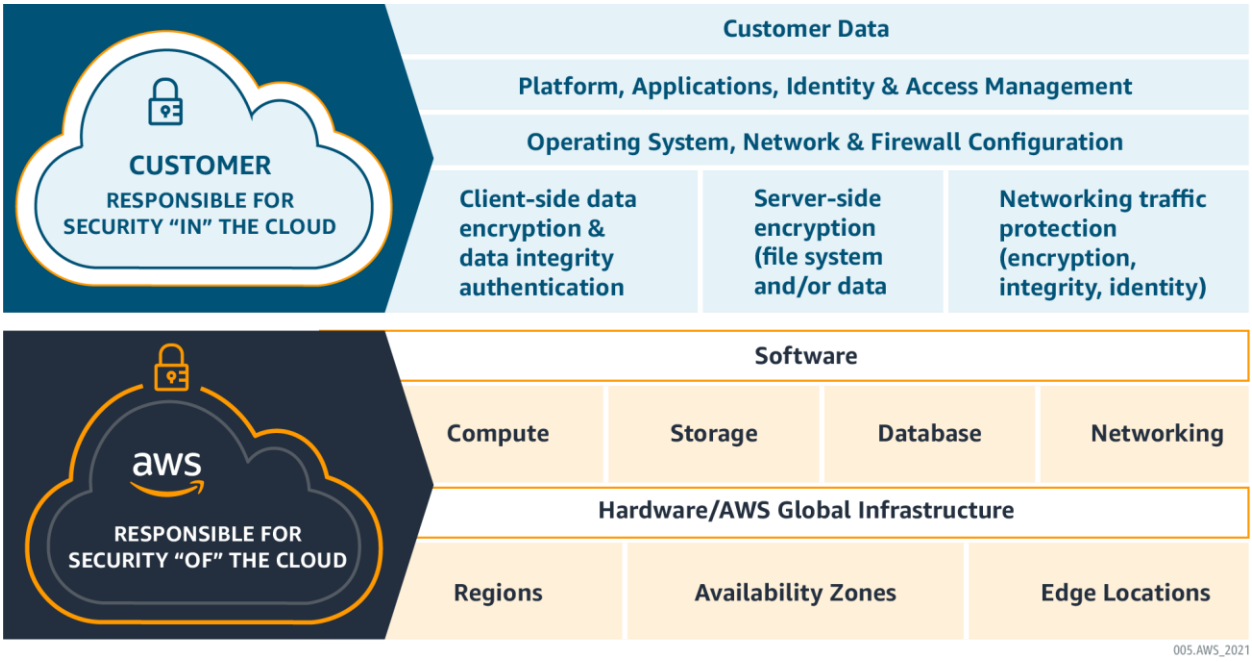

Figure 9. The Shared Responsibility Model

6.7.1. AWS and Customer Responsibilities

As described above, under the shared responsibility model, security and compliance responsibilities are shared between AWS and the customer.

6.7.1.1. AWS Responsibilities (Security of the Cloud)

AWS is responsible for protecting the infrastructure that runs all of the services offered in the AWS Cloud. This infrastructure is composed of the hardware, software, networking, and facilities that run AWS Cloud services. AWS operates, manages, and controls the infrastructure components that customers build upon.

We are vigilant about our customers' security and have implemented sophisticated technical and physical measures against unauthorized access. Customers can validate the security controls in place within the AWS environment through our certifications and reports, including the AWS System & Organization Control (SOC) 1, 2 and 3 reports, International Organization

for Standardization (ISO) 27001, 27017, 27018 and 9001 certifications, and Payment Card Industry (PCI) Data Security Standard (DSS) compliance reports.

Our ISO 27018 certification demonstrates that AWS has a system of controls in place that specifically address the privacy protection of customer content. These reports and certifications are produced by independent third-party auditors and attest to the design and operating effectiveness of AWS security controls. More information on AWS compliance certifications, reports, and alignment with best practices and standards can be found at http://aws.amazon.com/compliance/.

6.7.1.2. Customer Responsibilities (Security in the Cloud)

Customer responsibility is determined by the AWS Cloud services that they select. This determines the amount of configuration work the customer must perform as part of their security responsibilities. This determines the amount of configuration work the customer must perform as part of their security responsibilities.

For example, a service such as Amazon Elastic Compute Cloud (Amazon EC2) requires the customer to perform all of the necessary security configuration and management tasks. Customers that deploy an Amazon EC2 instance are responsible for management of the guest operating system (including updates and security patches), any application software or utilities installed on the instance, and the configuration of the AWS-provided firewall (called a *security group*) on each instance.

For abstracted services, such as Amazon Simple Storage Service (Amazon S3) and Amazon DynamoDB, AWS operates the infrastructure layer, the operating system, and platforms, and customers access the endpoints to store and retrieve data. Customers are responsible for managing their data (including encryption options), classifying their assets, and using AWS Identity and Access Management (IAM) tools to apply the appropriate permissions.

Customers are responsible for managing their data (including encryption options), classifying their assets, and using IAM tools to apply the appropriate permissions.

6.7.2. Shared Responsibility and IT Controls

The AWS shared responsibility model also extends to IT controls. Just as the responsibility to operate the IT environment is shared between AWS and its customers, so is the management, operation, and verification of IT controls. AWS can help relieve the customer burden of operating controls by managing those controls associated with the physical infrastructure deployed in the AWS environment (**inherited controls**).

Some controls apply to both the infrastructure layer (AWS responsibility) and customer layers (customer responsibility), but in completely separate contexts or perspectives (**shared controls**). With shared controls, AWS provides the requirements for the infrastructure, and the customer must provide their own control implementation within their use of AWS Cloud services. Examples of these shared controls include the following:

- **Patch Management**: AWS is responsible for patching and fixing flaws within the infrastructure, but customers are responsible for patching their guest operating system and applications.
- **Configuration Management**: AWS maintains the configuration of our infrastructure devices, but customers are responsible for configuring their own guest operating systems, databases, and applications.

- **Awareness and Training**: AWS trains AWS employees, but customers must train their own employees.

Some controls are the sole responsibility of the customer based on the application(s) they are deploying within the AWS Cloud (**customer-specific controls**). Examples include Service and Communications Protection or Zone Security, which may require a customer to route or zone data within specific security environments. Customers also control how they use their account and what content moves into and out of the account.

6.7.3. Shared Responsibility and Customer Data

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

6.7.3.1. Customer Content

Customers maintain ownership of their content, and they select which AWS Cloud services can process, store, and host it. We do not access or use customer content for any purpose without a customer's consent. We never use customer content or derive information from it for marketing or advertising.

We define customer content as software (including machine images), data, text, audio, video, or images that a customer or any end user transfers to us for processing, storage, or hosting by AWS Cloud services in connection with that customer's account; and any computational results that a customer or any end user derives from the foregoing, through their use of AWS Cloud services. For example, customer content includes content that a customer or any end user stores in Amazon S3. Customer content does not include account information, which we describe below. The terms of the AWS Customer Agreement and the AWS Service Terms apply to customer content.

There are five important basics concepts regarding customer content in the shared responsibility model:

1. Customers continue to own their content.
2. Customers choose the geographic location(s) in which to store their content—it does not move unless the customer decides to move it.
3. Customers can download or delete their content whenever they like.
4. Customers can "crypto-delete" their content by deleting the master encryption keys that are required to decrypt the data keys, which are, in turn, required to decrypt the data.
5. Customers should consider the sensitivity of their content and decide if and how to use various techniques such as encryption, tokenization, data decomposition, and cyber deception to protect their content.

Legal Requests for Customer Content

We are vigilant about the privacy of our customers. We do not disclose customer content unless we're required to do so to comply with the law, or with a valid and binding order of a governmental or regulatory body. Governmental and regulatory bodies need to follow the applicable legal process to obtain valid and binding orders. We review all orders and object to overbroad or otherwise inappropriate ones.

Unless prohibited from doing so, or if there is clear indication of illegal conduct in connection with the use of Amazon products or services, Amazon notifies customers before disclosing customer content so they can seek protection from disclosure. It is also important to point out

that our customers can encrypt their customer content, and we provide customers with the option to manage their own encryption keys.

We know that transparency matters to our customers, so we regularly publish a report about the types and volume of information requests we receive. The report is available on the Law Enforcement Information Requests webpage.

6.7.3.2. Account Information

We define account information as information about a customer that a customer provides to us in connection with the creation or administration of a customer account. For example, account information includes names, usernames, phone numbers, email addresses, and billing information associated with a customer account. The information practices described in the AWS Privacy Notice apply to account information.

6.7.3.3. Customer Virtual Instances

Customer virtual instances are solely controlled by the customer. AWS personnel do not have the ability to log into customer instances. AWS customers manage the creation and deletion of their data on AWS, maintain control of access permissions, and manage appropriate data retention policies and procedures. Controls in place limit access to systems and data, and provide access to systems or data that is restricted and monitored. Refer to the AWS SOC 1 audit report (available under AWS Nondisclosure Agreement [NDA]) for more information and validation of the control testing related to access permissions and data deletion. Refer to the AWS Payment Card Industry Data Security Standard (PCI DSS) Compliance Package (available under AWS NDA) for testing performed to confirm data deletion. Both the AWS SOC 1 audit report and the AWS PCI Compliance Package can be requested at http://aws.amazon.com/compliance/contact/.

6.7.4. Additional Resources

AWS Website:

- Overview of the shared responsibility model: https://aws.amazon.com/compliance/shared-responsibility-model/
- Detail on customer responsibilities relating to customer data: https://aws.amazon.com/compliance/data-privacy-faq/
- How the shared responsibility model interacts with the GDPR: https://aws.amazon.com/blogs/security/the-aws-shared-responsibility-model-and-gdpr/

AWS Whitepapers:

- Using AWS in the Context of Common Privacy & Data Protection Considerations (Detailed sections on security “of” and “in” the cloud)
- Risk and Compliance (Deep dive on risk and compliance in the context of the shared responsibility model)
- Overview of Security Processes (AWS security measures to deliver security “in” the cloud)
- Data Residency (How customers can manage and protect their data within the shared responsibility model)