

AWS Whitepaper

Next-Generation OSS with AWS



Next-Generation OSS with AWS: AWS Whitepaper

Copyright © 2024 Amazon Web Services, Inc. and/or its affiliates. All rights reserved.

Amazon's trademarks and trade dress may not be used in connection with any product or service that is not Amazon's, in any manner that is likely to cause confusion among customers, or in any manner that disparages or discredits Amazon. All other trademarks not owned by Amazon are the property of their respective owners, who may or may not be affiliated with, connected to, or sponsored by Amazon.

Table of Contents

| | |
|--|-----------|
| Abstract and introduction | i |
| Introduction | 1 |
| OSS Concepts | 3 |
| Business-centric view | 3 |
| Network-centric view | 3 |
| Network Characteristics | 4 |
| Fault Management | 5 |
| Configuration Management | 5 |
| Performance Management | 6 |
| Security Management | 7 |
| OSS Architecture on AWS | 9 |
| Governance & Network | 10 |
| Compute & Storage | 10 |
| Data Movement, Ingestion, Analysis, and Storage | 11 |
| Integration enablement | 12 |
| OSS Deployment Architecture | 13 |
| Domain Management | 14 |
| Service Assurance | 16 |
| Service Fulfillment | 18 |
| Service Orchestration | 20 |
| Network Analytics | 21 |
| Edge Analytics | 23 |
| Data Unification | 24 |
| Security | 25 |
| Security of the OSS Solution | 25 |
| Security of the network functions | 26 |
| Connectivity | 26 |
| Digital Transformation and DSP Enablement | 27 |
| Conclusion | 28 |
| Contributors | 29 |
| Further reading | 30 |
| Document history | 31 |
| Glossary | 32 |
| Notices | 35 |

AWS Glossary 36

Next-Generation OSS with AWS

Publication date: **September 21, 2021** ([Document history](#))

An Operations Support System (OSS) is key to enabling Communication Service Providers' digital transformation journey. Building OSS applications on Amazon Web Services (AWS) enables operational efficiency, cost reduction, elasticity, and innovation. This whitepaper outlines the best practices for developing OSS applications on the AWS Cloud platform, and offers reference architectures to guide organizations in the delivery of OSS solutions spanning domain management, service assurance, service fulfillment, service orchestration, and network analytics.

Introduction

Communication Service Providers (CSPs) are constantly looking for ways to improve the efficiency of their network operations, reduce their time to market and their operating costs, and adapt to technological evolution. CSPs view their network as an array of network components and network services no longer separated by arbitrary functional lines. CSPs require an operational agility that enables forthcoming opportunities such as Mobile Virtual Network Operator (MVNO), business plans for new network slices (for several use cases like IoT), consumers application (such as ultra-low latency gaming), enterprise application (such as private network), and virtual reality (VR) and augmented reality (AR) business-to-consumer applications (which are unlocked by the introduction of 5G technology).

CSPs are looking for OSS to enable operational agility, and want an OSS that informs on what is being serviced and where, what is being provisioned and where, and how their entire operations perform, predict faults, and self-heal. CSPs want an OSS that allows them to programmatically deploy new network services and functions and have them readily available, and to dynamically reconfigure their network while reducing the complexity of their OSS stack.

The traditional view of an OSS stack that is comprised of multiple independent and functionally-separated network management functions doesn't match the dynamic behavior of Network Function Virtualization (NFV) and doesn't take advantage of the adaptability that the NFV cloud architecture enables. CSPs OSS solutions need to evolve with the network.

This whitepaper describes the benefits of OSS on AWS. It includes an OSS reference architecture, an overview of OSS functions and requirements based on characterizing the network OSS manages, operational characteristics, use cases, and best practices for architecting OSS on AWS

(which includes high-availability, scalability, security, performance, and operational excellence). Information contained in this document will enable you to develop a next-generation OSS solution on AWS, which will provide a cost-efficient and agile path to CSPs in their digital transformation journey to becoming Digital Service Providers (DSPs).

OSS Concepts

There are two conceptual views to define an OSS solution: a business-centric view and a network-centric view. Before diving into the recommended OSS solution architecture on AWS, the following section deconstructs OSS concepts and associates them with technical requirements that are the foundations of the proposed OSS solution on AWS.

Business-centric view

The business-centric view of OSS aligns with the enhanced Telecom Operations Map (eTOM), where OSS applications provide capabilities in fulfillment, assurance, and operations support and readiness. This logical representation helps define which outcome an OSS component is aligned with, which helps DSPs to ensure their ability in delivering their services.

Service fulfillment enables the service provider to plan, build, provision, and activate an end-to-end service. These are OSS applications that help with the creation of an Ultra Reliable Low Latency Communications (URLLC) network slice dedicated to self-driving cars, the activation of a cell site, customer equipment, the provisioning of new optical transport equipment, and more high-level services like voice over 5G.

Service assurance enables the service provider to optimize the Quality of Service (QoS) associated with its networks and services. These OSS applications ensure that end users are provided with the best quality of experience by enabling the monitoring of Service Level Agreement compliance and providing clear information on existing problems while predicting future problems and enabling preemptive correction.

Operations support and readiness (OSR) equips service providers with the tools, architecture, and environment to execute fulfillment and assurance processes efficiently. Next-generation OSS is geared towards extreme automation across all business processes by enabling self-planning, self-optimizing, and self-healing capabilities. For example, users can cross-correlate the data to improve customer satisfaction (e.g. by enabling proactive fault notifications), or recommend specific services based on existing service consumption.

Network-centric view

The network-centric view of OSS aligns with the Open Systems Interconnection (OSI) network management model where OSS applications are grouped into the following five categories: Fault, Configuration, Accounting, Performance, and Security (FCAPS).

- **Fault management (FM)** encompasses software that helps a CSP understand whether a network component, network subsystem, and/or network is healthy or in a faulty mode. It helps with detection, correction, isolation, and recovery of faults, and generation, handling, distribution, and clearance of alarms.
- **Configuration management** includes software that helps a CSP define the characteristics of their network components. It helps with network/system discovery, inventory cataloging, capacity and resource availability, provisioning, rollback, and configuration life cycle management.
- **Accounting** helps with network/system usage, managing quotas, audits, costing, and applicable fraud management measures. Software within this category supports Business System Support (BSS) systems as well.
- **Performance Management (PM)** includes software that helps a CSP understand its network performance of its network, network functions, and up to the smallest components of their network functions. It handles the network/system utilization, report generation, metrics that are specific to system performance, and capacity planning.
- **Security** helps with access management, audits, incident reporting and management, and compliance. It allows CSPs to define who accesses what areas of their network components, security of the software supporting the network operations, encryption of data residing and transiting through their network, etc.

By tackling both network and business views, the proposed OSS Solution on AWS can help CSPs to reduce the complexity of their operations while helping them in uniformizing their operations.

Network Characteristics

For an OSS solution to provide value, it should align with network characteristics while fulfilling DSPs' business outcomes. At its core, an OSS solution's intent is to facilitate the management and the enablement of a given network.

A telecommunication network is comprised of numerous Network Functions (NFx) that are categorized as wireless and wireline network functions. Such functions span physical and software equipment such as routers, firewalls, Mobility Management Entity (MME), Evolved Node B (eNodeB), Access and Mobility management Function (AMF), User Plane Function (UPF), small cells, microwave backhaul, 5G New Radio (5G NR), and many more. Each of these NFx must be provisioned, monitored, analyzed, configured, and secured. An OSS architecture on AWS can help

you reduce the complexity inherent to managing a complex network comprised of many NFx, often from many Network Equipment Providers (NEPs).

Fault Management

Telecommunication networks are inherently critical to our lives. For example, their network architecture must align with regulatory constraints that require high availability services (e.g. emergency services numbers like 9-1-1).

Note

NFx and General Architectural guidelines are not discussed in this whitepaper. An OSS architecture on AWS can help you identify when a network component is at fault, and provide or act on that information in near real-time.

Different NFx can have different impacts on a given service provided to a consumer. Based on the impact and the type of information NFx can generate, a CSP may require an alarm/event to be treated on the fly or stored for future analysis. An OSS architecture on AWS can help process an alarm with or without storing it. Similarly, it supports numerous storage options.

A given NFx can become faulty at any time for numerous reasons: loss of power, faulty software, faulty hardware, faulty operations, natural disasters, etc. These events follow a random pattern. They could be isolated to a unique NFx, generalized to a few NFx, or impacting a large amount of NFx. An OSS architecture on AWS can help you scale up and down based on the flow of alarms coming from the network.

Given the complexity of the network and the plurality of components in a given service path, a CSP's OSS solution should help in identifying which faulty component results in network degradation (of any form) while filtering noise. For example, a down S1 interface will result in both eNodeB and MME alarms. An OSS architecture on AWS can help you identify which element is faulty through alarm correlations, configuration knowledge, logical and physical inventory information, learning historical behaviors, and using predictive models.

Configuration Management

Configuration Management (CM) is a central part of an OSS solution. CM includes applications that help manage NFx configuration and provides an inventory layer to subsequent OSS applications

as well as BSS applications. Note that in this document, we purposely didn't expand on Inventory Management (IM), which is a critical component to CM. At its core, CM is a database (or many databases) that is constantly updated with network CM data and constantly requested by network CM data consumers. It provides a centralized view of the network.

The amount of network changes a CM solution should support varies based on use cases. A DSP may want to optimize an NFX or push a new service to thousands of NFX. Similarly, events can occur on one NFX or many NFX that cause NFX configuration to self-change their behavior and configurations. To efficiently support this behavior, an OSS architecture on AWS can help you scale up and down based on the amount of network changes requested, ingested, or processed, whether programmatically or by users.

When a change event (or a service state change) occurs on an NFX, a DSP may want to trigger a script that fine tunes the NFX itself or other NFX without negatively impacting the NFX, other NFX, and the overall service delivery. An OSS architecture on AWS can help you be responsive and support the positioning of the CM workload to reduce latency when needed.

DSPs have to handle a new layer of complexity with the introduction of new architectures and new concepts, such as 5G. Each of these new NFX and new architectures bring more data for the CM solution to handle. With financial pressure, limited knowledgeable engineer availability, and time-to-market requirements, DSPs are looking for a CM solution that incorporates machine learning concepts to gain insight from their network configuration, identify patterns and anomalies, and support automation. An OSS architecture on AWS can help you leverage data lake concepts to enable an intelligent and self-operating network.

Performance Management

PM encompasses applications that help manage NFX performance. PM allows a DSP to enable observability with a specific granularity on a given NFX, multiple NFX, and network wide. An OSS architecture on AWS provides DSPs with the network analytics capabilities that give them insight on their network performance while their network evolves.

An OSS architecture on AWS can help DSPs to automatically discover their NFX datasets as their network evolves, while having the ability to ingest large amounts of information in a scalable manner. The reference architecture provided in this whitepaper helps your solution to be cost-efficient by reducing data duplication. This is achieved by building a PM solution on top of a data lake to simplify the navigation through a large amount of data, and providing the required data governance.

Our proposed OSS Architecture on AWS can help DSPs reduce their operational costs by providing programmatic data insights and enabling the development of workflow based on performance data. It provides DSPs with the ability to identify which components have the most negative impact on their overall network performance while learning what series of configuration, changes, faults, and KPIs result in performance degradation.

The proposed OSS Architecture on AWS enables data enrichment and learning from past events. It supports constant innovation by providing an architecture that allows for quick experimentation and quickly building applications (graphical, programmatic, storage quarriable, etc.).

Security Management

Security Management is a critical component to an OSS solution. It can be viewed in two parts:

- Security of the OSS solution.
- Security of the network it manages.

First, as defined in [AWS Well-Architected Framework](#), our proposed AWS Architecture on AWS is built on the seven design principles for security in the cloud:

- Implement a strong identity foundation
- Enable traceability
- Apply security at all layers
- Automate security best practices
- Protect data in transit and at rest,
- Keep people away from data
- Prepare for security events

The proposed OSS architecture on AWS helps you identify operational mistakes, identify risk vectors before they appear, predict operational mistakes, and prevent malicious parties from having prime access to a DSP's network.

Second, the proposed OSS architecture on AWS simplifies securing the network it manages. This is achieved by providing auditing capabilities to determine if an NFX's configuration deviates from a secure gold standard, and by providing automation that addresses identified security issues. Similarly, the proposed OSS architecture on AWS can help you simplify the distribution of security

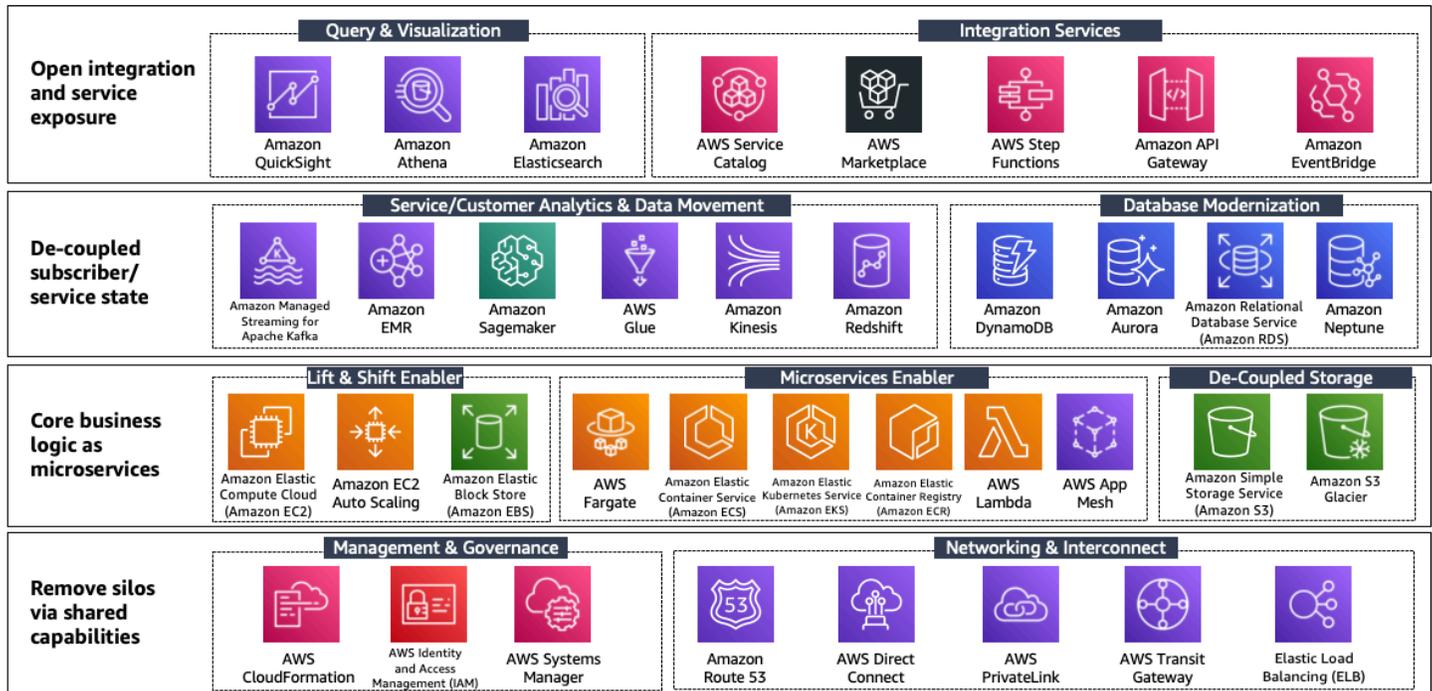
updates to both physical and virtual NFx, providing the scalability to adequately upgrade a large portion of the network while minimizing network impact. Finally, the proposed architecture can help you to simplify the encryption of network traffic and Operations, Administration and Maintenance (OAM) traffic while reducing the encryption cost.

By having an OSS solution that enables securing the network while providing DSPs with the security governance to protect a given network, DSPs can benefit from enhanced security compliance.

The proposed next-generation OSS architecture on AWS can help you align with the BSS solutions of the future to eliminate traditional business and operational silos. It's critical to view network data that enables business logics as easily-made-accessible by the OSS stack, while also providing the framework to enrich data based on business logic. This convergence of BSS and OSS stacks is a key enabler to CSPs becoming DSPs. This combination of BSS and OSS capabilities in a unified layer supporting the business and operations of DSPs is also called Digital Support System (DSS).

OSS Architecture on AWS

This section introduces an OSS architecture framework on AWS that aligns with network characteristics and business drivers described in the previous chapter. The following reference architecture illustrates key enablers that apply to the entirety of the OSS stack, providing guidance in architecting principles to enable a next-generation OSS solution.



Key Enablers for Telco Digital Support System

Topics

- [Governance & Network](#)
- [Compute & Storage](#)
- [Data Movement, Ingestion, Analysis, and Storage](#)
- [Integration enablement](#)
- [OSS Deployment Architecture](#)
- [Domain Management](#)
- [Service Assurance](#)
- [Service Fulfillment](#)
- [Service Orchestration](#)

- [Network Analytics](#)
- [Edge Analytics](#)
- [Data Unification](#)
- [Security](#)

Governance & Network

Whereas a typical OSS application stack comprises of multi-vendor implementation, often with duplicated and overlapping tooling across different network domains, an OSS architecture on AWS simplifies:

- The management and governance of OSS applications with services such as [AWS CloudFormation](#), [AWS Identity and Access Management](#) (IAM), and [AWS Systems Manager](#).
- Network and interconnecting of OSS, and IT and network workloads with services such as [AWS Direct Connect](#) (Direct Connect), [AWS Transit Gateway](#) and [Elastic Load Balancing](#).

Many of these shared capabilities can be deployed as part of a secure, multi-account [AWS Landing Zone](#) that is part of the same [AWS Organization](#) (for example, a shared security account or shared networking account). The creation, management, and governance of these landing zones is simplified and automated by leveraging [AWS Control Tower](#). These services help you reduce the complexity inherent to managing a complex network comprised of multiple technology domains, NFX, and Independent Software Vendors (ISVs).

Compute & Storage

[Amazon Elastic Compute Cloud](#) (Amazon EC2) provides the broadest and deepest compute platform with choices of processor, storage, networking, and purchase model, enabling the migration of legacy OSS workload into the cloud. Similarly, leveraging Auto Scaling and [AWS Graviton Processor](#), AWS provides OSS Solution developers with the ability to optimize the performance, and to get a price performance that is up to 40% better over comparable current generation x86-based instances. [Amazon EC2 Auto Scaling](#) provides the flexibility to scale manually, on-demand, on a schedule, and predictively. For example, when expanding a live mobility network, service fulfillment applications can be directed to scale up at the start of the maintenance window and scale down at the end of it.

[Amazon Elastic Kubernetes Service](#) (Amazon EKS) enables you to run Kubernetes-compliant OSS applications on AWS without the need to install and operate your own Kubernetes control plane. One Amazon EKS cluster can support up to 10 node groups, where each node group can support up to 100 nodes. This enables you to reduce the overall complexity of your OSS stack and limit the traditional control plane overhead associated with on-premises workloads. [Amazon Elastic Container Service](#) (Amazon ECS) provides you with a fully-managed container orchestration service and leverages serverless technologies from [AWS Fargate](#) to provide your OSS application with dynamic scaling ability without the typical provisioning, configuration, and scaling overhead of managing infrastructure fleets. [AWS App Mesh](#) provides you with the observability and control to govern how OSS services communicate with each other and what metrics, traces, and logs to capture.

[Amazon Elastic Block Store](#) (EBS) provides a high-performance, block-storage service designed for use with Amazon EC2, which enables OSS applications that are throughput and transaction intensive, such as running a custom graph database representing network topologies in real-time, or running a database to support a legacy OSS application. [Amazon Simple Storage Service](#) (Amazon S3) provides the scalability, data availability, security, and performance required to store network performance data and network configuration exports, and share across the OSS stack, enabling to scale the size of the network supported by an OSS solution. Similarly, Amazon S3 provides various storage classes, inclusive of [Amazon S3 Glacier](#) (that provides low-cost data archiving for stringent data storage regulatory requirements). [Amazon Elastic File System](#) (EFS) provides a simple, serverless, set-and-forget, elastic file system that enables the sharing of configuration data, temporary files, logs, etc., across the entire OSS application stack, without having to provision or manage the storage. This enables CSPs to leverage cloud storage benefits for legacy and new OSS applications.

Data Movement, Ingestion, Analysis, and Storage

AWS provides you with managed services to help you ingest network data at scale, and move, analyze, and store the data. [Amazon Managed Streaming for Apache Kafka](#) (Amazon MSK) provides you with a path to migrate your Kafka Streams applications to AWS Cloud. Amazon MSK provides you with scaling capabilities while eliminating the effort taken to self-manage Apache Kafka brokers and its associated components.

[Amazon Kinesis Data Streams](#) (KDS) provides you with a serverless, scalable, and durable real-time data streaming service, allowing your OSS solution to ingest network events such as alarms, configuration changes, and signaling events. A Kinesis stream is comprised of one or more shard,

where the latter is a uniquely-identified sequence of data records in a stream. The rate of data flowing through the stream is a function of the number of shards in a stream. Using prediction models, defined schedules, or monitored KPIs, you can perform *resharding* on a stream to maintain a data rate when a network condition obliges. For example, when a hurricane generates a large number of network alarms and service failures, *resharding* allows you to scale your stream to maintain the rate of data and support your Service Assurance applications.

[AWS Glue](#) is a serverless data integration service that enables you to discover and prepare data to support your OSS application. For example, using AWS Glue, you can transform the format of data ingested from a newly-integrated network element into a format that is suitable for your Service Assurance, Domain Management, and/or Network Analytic solution. AWS Glue helps you build applications that automatically discover network elements, network services, and north-south-east-west application inputs.

AWS provides you with [purpose-built, managed database services](#) to support your OSS data structures and transactions needs. For example, [Amazon Neptune](#) is a fully-managed graph database service that enables you to represent complex network service relationships, enabling your Service Assurance applications to detect network anomalies and misconfigurations, and provides your network engineering teams with recommendations. [Amazon Aurora](#) provides you with a MySQL and PostgreSQL-compatible relational database built for the cloud. Amazon Aurora is up to five times faster than standard MySQL databases and three times faster than standard PostgreSQL databases, providing you with the performance to enable your next-generation OSS solution.

Integration enablement

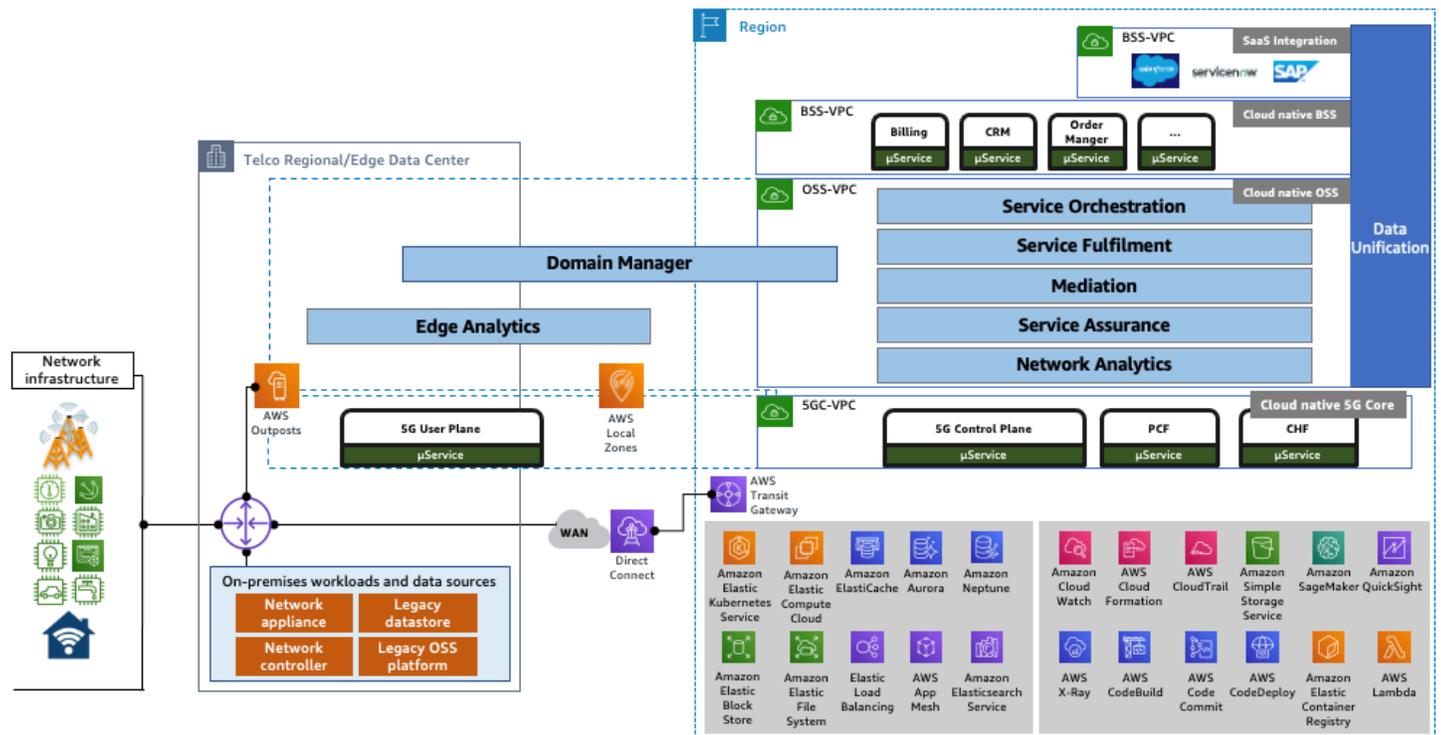
Integration enablement aligns with the increasing focus of implementing modular OSS applications exposed via Open APIs, and the need to support dynamic exposure of network-as-a-service constructs towards enterprise or ecosystem partners. Collectively, this is enabled by [Amazon API Gateway](#), which simplifies the development and deployment lifecycle of both internal and external APIs.

The integration framework is enhanced by using [AWS Step Functions](#), which enables a coordinated state-machine implementation, and [Amazon EventBridge](#) (EventBridge), which enables event-driven integration with other BSS/external applications as part of the OSS value-chain transaction. To enable a consistent, common, consumption layer, network operation or product/business teams can extract business insights via [Amazon QuickSight](#) (QuickSight), [Amazon Athena](#) and [Amazon OpenSearch Service](#). [Amazon S3](#) can be leveraged as common object storage across all

OSS applications that can handle different network data formats, and serve as a foundation of a shared data unification layer.

OSS Deployment Architecture

While we acknowledge the broad industry definition of what constitutes OSS, the proposed OSS architecture framework focuses on the most critical modules commonly observed in OSS deployments, and their deployment architecture on AWS. The following reference architecture illustrates the deployment strategy of OSS modules on [AWS Regions](#) and [AWS Outposts](#) based on their latency needs. For example, a PM/CM Event Collector can run on AWS Outposts, and feed a real-time Self-Organizing Network (SON) application while aggregating events for processing in [AWS Regions](#) for applications that do not have a stringent latency budget. Similarly, you can also leverage [AWS Local Zones](#) for latency-sensitive applications while having access to the elasticity, scalability, and security benefits of the cloud.



OSS Deployment Architecture on AWS

With the exception of Edge analytics and a few Domain Manager functions, the majority of OSS functions can be deployed in centralized AWS Cloud regions given these OSS workloads do not directly participate in network data-lane or network control plane. The ability to deploy in a region means these workloads can be deployed across multi-[Availability Zones](#), can leverage the

full suite of cloud enablers, and can benefit from the full set of cloud elasticity and purchasing options available (including pricing for [Reserved Instance](#) and [Spot Instance](#)). Depending on the Recovery Time Objective (RTO)/Recovery Point Objective (RPO) target and specific disaster recovery requirement, this architecture can be further extended to a multi-region design.

As illustrated in the previous reference architecture, the proposed OSS deployment on AWS enables a Data Unification layer between OSS, BSS, Software as a Service (SaaS) platforms, and NFX. By leveraging Data Lake and Lake House concepts, DSPs have the ability to optimize data flows between applications, and limit traditional data duplication.

A cloud-centric model also represents a clear opportunity to implement a shared services model, in which multiple OSS modules (potentially managed by different teams or even supplied by different ISV partners) can be deployed into a consistent, centralized, managed cloud environment (an experience that is enabled by [AWS Organizations](#) and [AWS Control Tower](#)), which adheres to a common set of security and audit standards and leverages from a common set of cloud-native capabilities for management and operations. This [landing zone](#) environment can also be extended to network or BSS workloads to maximize synergies across the broader Telco organization.

Our proposed OSS deployment on AWS allows you to place workloads closer to the network when needed. As such, edge analytics capabilities (such as what's prescribed by Open Radio Access Network (ORAN) RAN Intelligent Controller (RIC)) can be collocated with network functions, enabling near real-time application.

The following sub-sections outline the foundational concepts to build an OSS architecture on AWS for domain management, service assurance, service fulfillment, service orchestration, network analytics, and edge analytics.

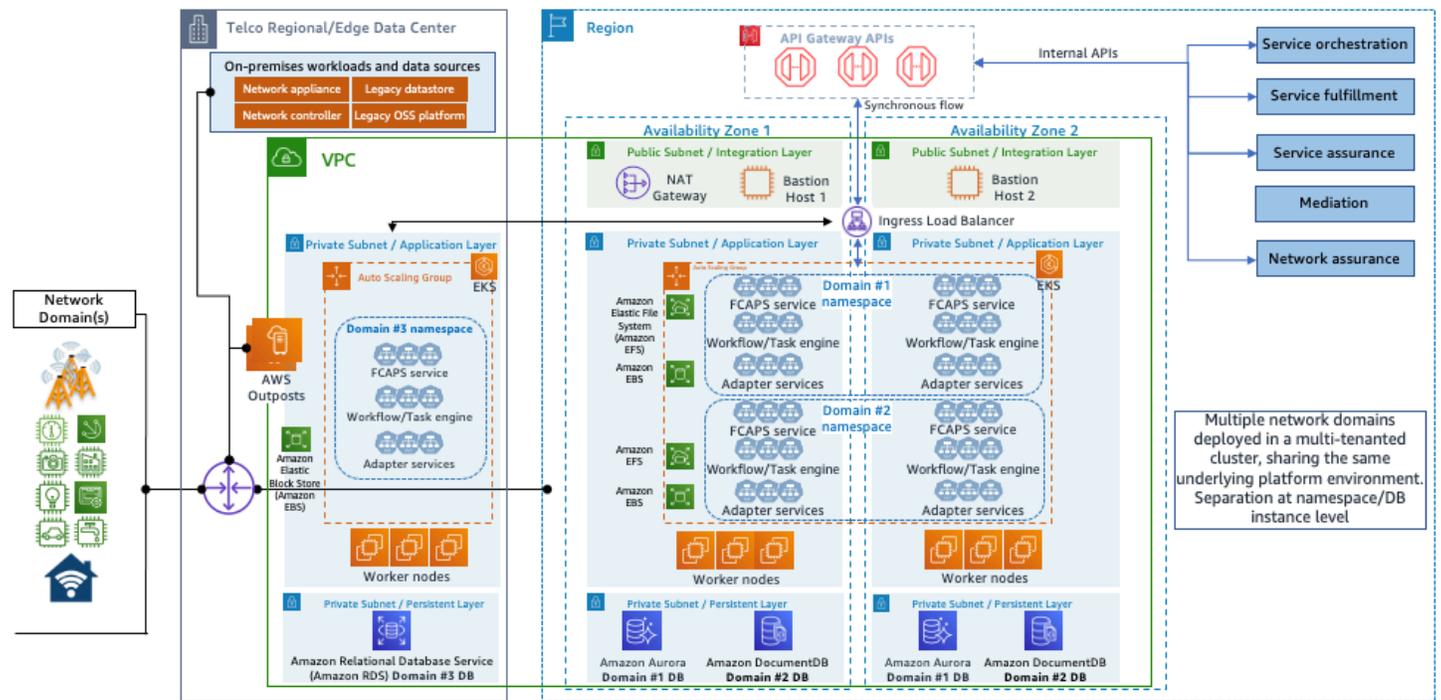
Domain Management

Domain management is undergoing a shift towards a more open approach to reduce network complexity, enable CSPs to either forego domain managers or evolve into open domain managers, and have a multi-vendors approach to domain management. Domain management on AWS helps you reduce the size of your OSS stack, and helps you eliminate infrastructure complexity associated with the operations of domain managers (per NFX, per vendor, and for a given capacity).

The following reference architecture outlines an example domain management implementation leveraging AWS Outposts for functions requiring low-latency budget. AWS Regions are leveraged for mediation and domain-specific applications that enable engineering, operations, and planning groups to efficiently perform their tasks.

Amazon S3 provides a scalable solution to host network configuration exports and mediate performance data, providing you with the control to apply Life Cycle Management (LCM) policies that are specific to your needs. Amazon Elastic File System (EFS) provides you with scalable and elastic file storage. You can mount EFS on your on-premise legacy OSS systems using standard Linux commands for mounting a file system via the NFSv4.1 protocol. This enables you to take advantage of the AWS Cloud, even for legacy systems, and enables CSPs to move away from complex and costly hardware expansions.

Similarly, AWS enables CSPs (and DSPs) to migrate to cloud databases using services such as the [AWS Schema Conversion Tool \(SCT\)](#) and [AWS Database Migration Service \(DMS\)](#), providing you with the tools to automate schema conversion and data movement. The process of developing APIs is simplified by [AWS API Gateway](#) to expose domain management functions and build ones that spawn domains, NFX, and technologies.



Domain Management Architecture on AWS

[Amazon EKS](#) provides you with both Kubernetes namespace capabilities and [AWS Auto-Scaling group](#) to reduce infrastructure costs of domain management. CSPs and DSPs can run domain-specific as well as multi-domain domain managers on different namespaces, simplifying their operations. By separating domain managers across namespace and via role-based access control integration with [AWS IAM](#), it's possible to control per-domain-level access to the Kubernetes API for compute-level isolation between domains. Further networking and storage-level isolation is also possible via network policies and service mesh, and via volume-defined, per-storage classes.

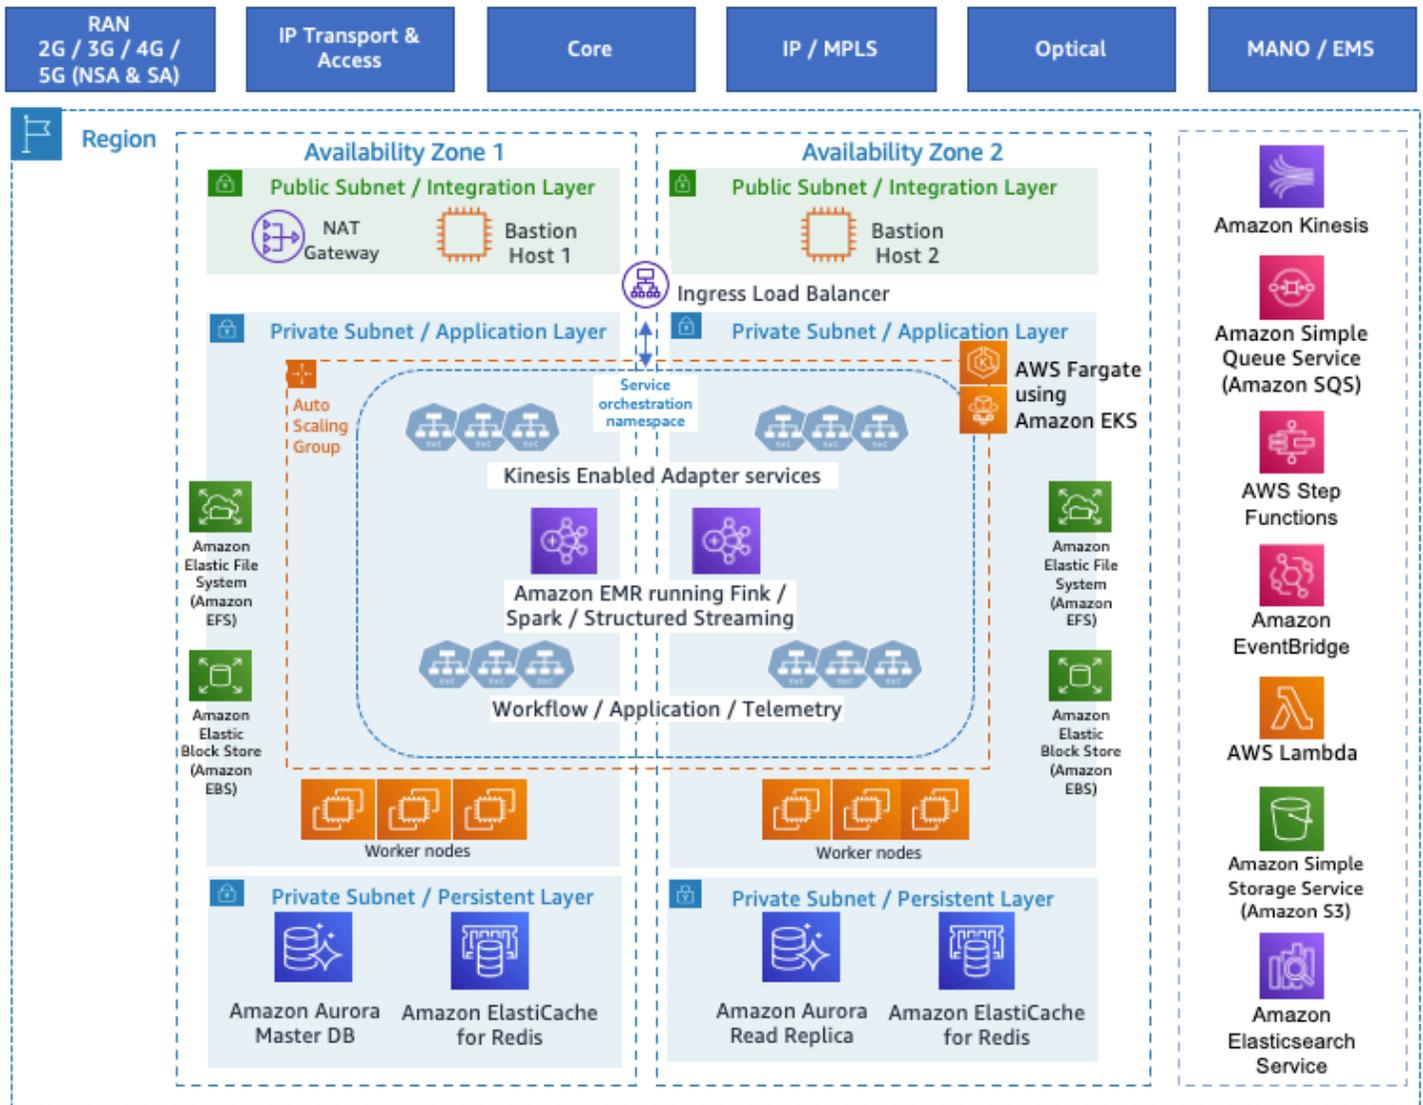
This enables CSPs and DSPs to eliminate the infrastructure complexity of on-premise domain managers, allowing them to take advantage of the AWS Cloud benefits such as elasticity.

AWS services such as [Amazon CloudWatch](#) (CloudWatch) and Kinesis can be leveraged to manage OAM data from traditional network elements, given those elements are running on Linux. For example, CloudWatch agents can be installed on an NFX to collect standard metrics such as CPU utilization, as well as process custom metrics using StatsD or collectd protocols. The [Kinesis Client Library](#) (KCL) provides an easy-to-use programming model for processing data. This enables the processing of real-time configuration events and alarm events from NFX. With Prometheus Server Grafana Agent, you can also collect metrics from NFX, which provides Domain Manager with the ability to expose a real-time dashboard for analysis and view of the network it manages. [AWS Systems Manager](#) provides you with the capability to automate operational tasks across on-premise NFX as well as towards legacy OSS systems. Operators and ISVs can leverage System Manager Agent (SSM Agent) to apply security patches, create automated responses, etc.

The proposed architecture enables you to migrate legacy domain managers from on-premise to AWS Cloud, and provides you with a path to leverage AWS services natively for OAM data.

Service Assurance

This section presents a service assurance architecture on AWS that provides you with the scalability, flexibility, reliability, and innovation to enable a fully-automated network that identifies issues and heals itself. The following reference architecture depicts this as well as illustrates key services enabling such automation.



Service Assurance Architecture on AWS

[Kinesis](#) can be leveraged to ingest network events from all DSPs NFX. Kinesis provides you with the scalability to ingest alarms and configuration changes when they occur. It also enables you to integrate with AWS services to perform operations based on the ingested event. For example, a network event from NFX, ingested through Kinesis, can trigger an [AWS Step Functions](#) that orchestrates a workflow; this workflow could correlate three things to trigger a corrective action to the network:

- KPIs available through a data lake query
- Network configuration validation using an [AWS Lambda](#) function
- Leveraging a prediction model using [Amazon SageMaker](#)

The workflow and the result of the workflow can be achieved using a low-code visual workflow designer, [Workflow Studio for AWS Step Functions](#). The previously described services and the type of workflow enables you to identify when a network component/service is degraded, provide that status across the OSS stack, and act on that status to correct the situation.

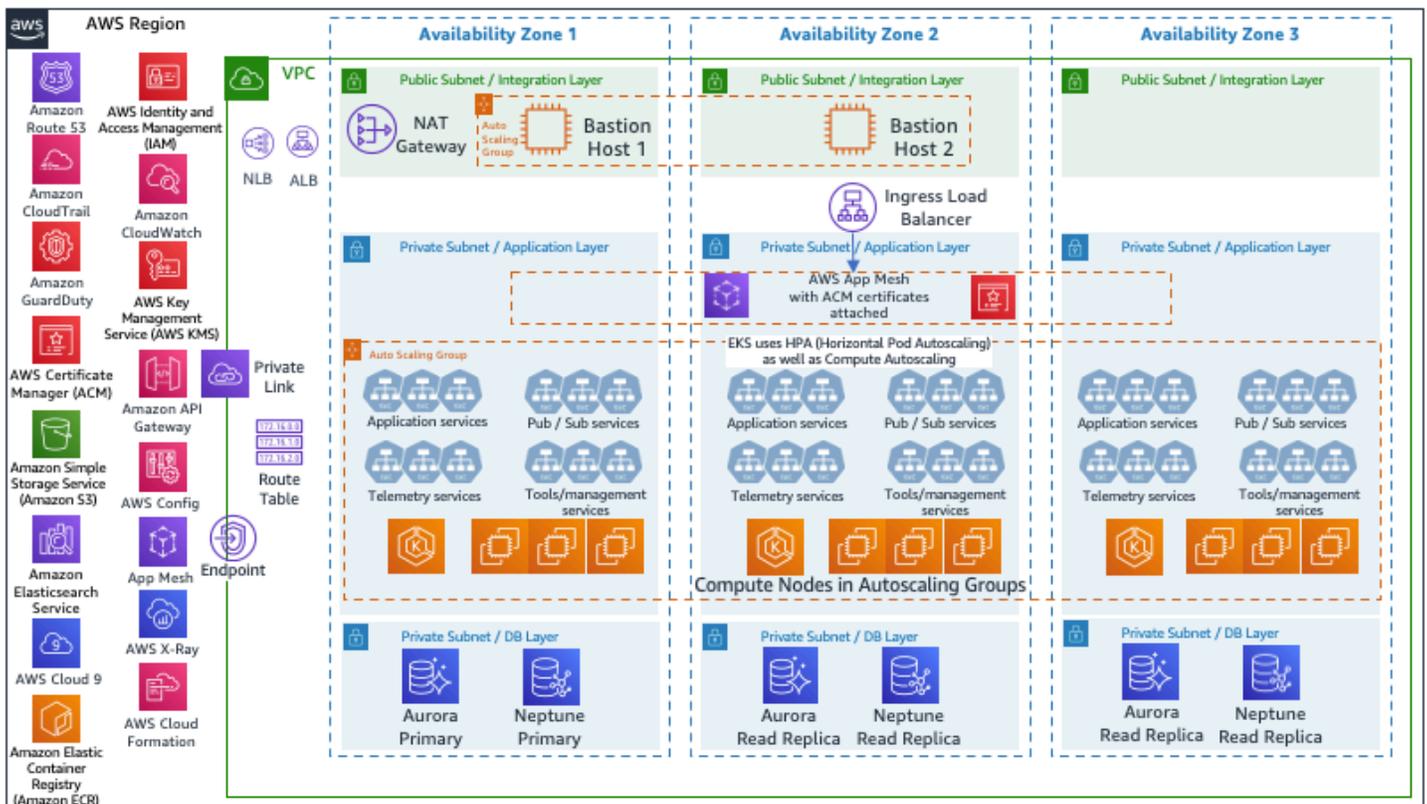
[Amazon EMR](#) or partner products can be used to store a vast amount of data. EMR simplifies the processing of data and removes the operational complexities associated with provisioning and tuning of big data clusters. The flexibility that EMR provides you with enables you to process trained data at a low cost. Mobility use cases often require 2-5 weeks of network data, and EMR enables you to only spin up the required compute capacity once training data is available.

By decoupling persistent data layer from the application layer through usage of managed services (e.g., [Amazon Aurora](#) and [Amazon Neptune](#)), data is easily consumed by applications within the OSS stack, enabling closed loop use cases. For example, real-time state of inventory can be used to input service orchestration and fulfillment.

Finally, [AWS Auto Scaling](#) is used to scale up or down based on demand, and thus optimal usage of infrastructure is achieved and total cost of ownership is optimized. An OSS solution on AWS scales with the network as network events occur.

Service Fulfillment

This section presents a service fulfillment architecture on AWS that provides flexibility, scalability, reliability, and the integration points to enable the customer journey. As depicted by the following reference architecture, the proposed architecture enables you to seamlessly integrate with order management, service orchestration, service assurance, and domain managers across multiple networks and service domains. Services such as [Amazon API Gateway](#), [AWS App Mesh](#), and [Amazon Virtual Private Cloud](#) (Amazon VPC) provide you with the ability to develop service fulfillment applications that are fully integrated with service orchestration applications and service assurance applications, enabling you to eliminate functional duplication and choose delimitation based on a given technology, network, and service type. For example, [Amazon Aurora](#) provides you with a MySQL and PostgreSQL-compatible relational database that can host inventory data used both in service fulfillment and orchestration. While your fulfillment can act as a central location of the inventory data for network services, Amazon Aurora enables you to define read-replicas. This enables you to support low-latency reading of network services for your service orchestration to speed up decision-making while providing APIs to govern provisioning requests at the fulfillment level.



Service Fulfilment Architecture on AWS

[Amazon EKS](#) provides you with the ability to run Kubernetes applications that scale. To achieve high availability and resiliency, the pods are distributed across multiple [AZs](#). [AWS' purpose-built database services](#) for Graph DB, NoSQL, and RDBMS, can further help you achieve your goals. An ingress gateway fronts the communication to and between your applications and uses the AWS native service [AWS App Mesh](#) to provide application networking, and offer end-to-end visibility and high availability.

AWS Step Functions allows you to seamlessly integrate service fulfillment applications with order management applications, allowing you to execute a multitude of events (such as dependency verification, dates, location validation, breaking tasks into sub-tasks, and executing the configuration on an individual network function).

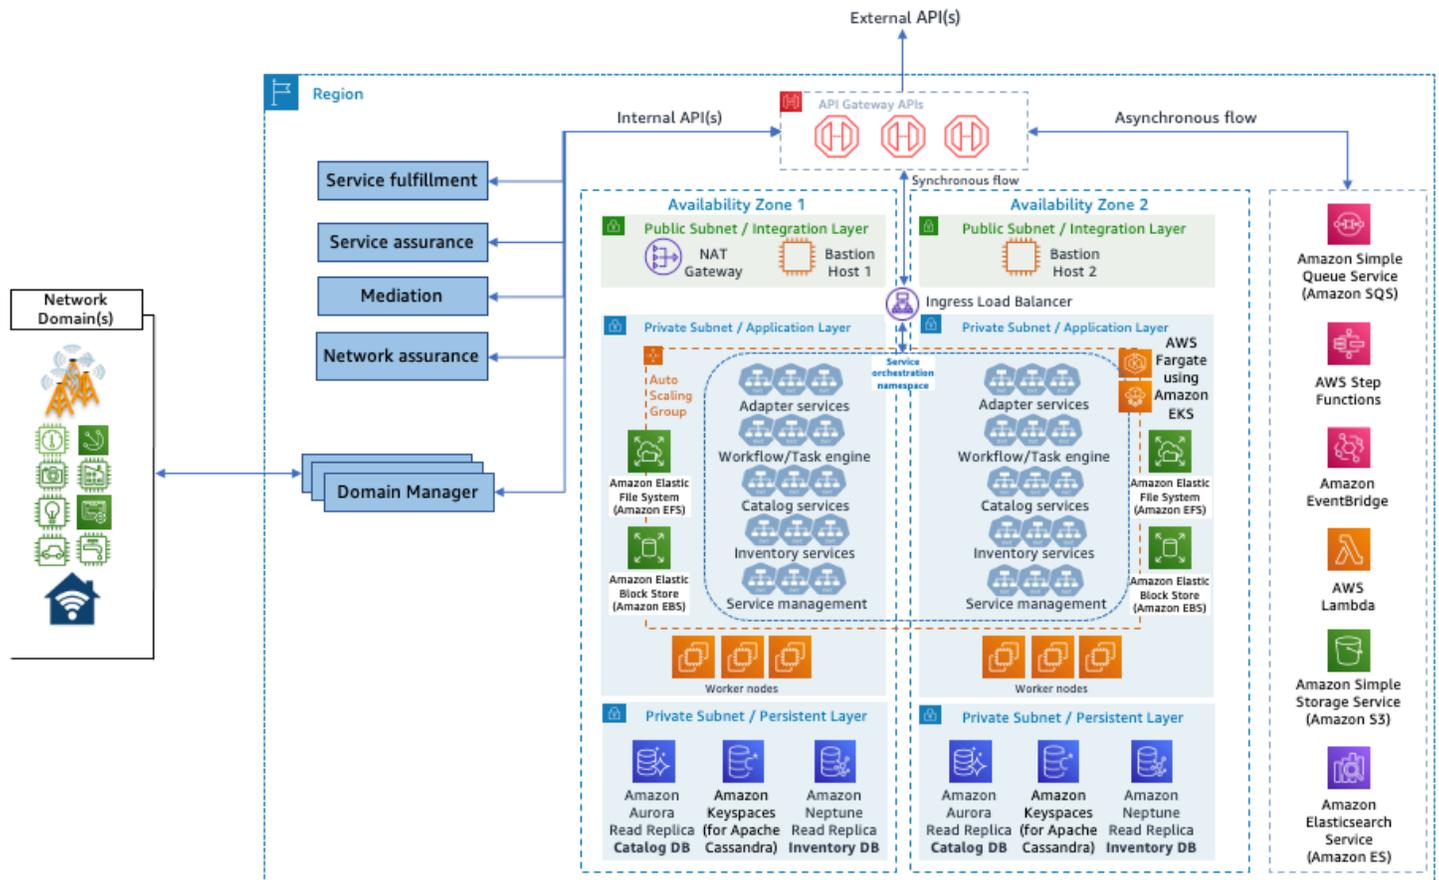
Dynamic service inventory management is done by using [Amazon Relational Database Service](#) (RDS) and [graph databases](#) to depict the relationship of a service, its status, and the underlying provisioned resources.

Fulfilment tracking in near real-time is enabled by [Amazon Simple Notification Service](#) (Amazon SNS) and [Amazon Simple Queue Service](#)(SQS). It provides you with the mechanisms to control

how fulfillment operations interact within the service fulfillment stack as well as between other applications in the OSS Stack.

Service Orchestration

AWS services enable an event-driven Service Orchestration (SO) architecture, which leverages AWS serverless services such as [AWS Lambda](#) and [AWS Step Functions](#). The following reference architecture enables you to build SOs that are aligned with European Telecommunications Standards Institute (ETSI) Management and Orchestration (MANO) while enabling cost optimization, security, scalability, and innovation. Similarly, AWS Partners' ecosystems enable you to leverage ready-to-be-deployed SO solution.



Service Orchestration Architecture on AWS

This architecture highlights the exposure/integration layer of a service orchestration platform, which leverages [Amazon API Gateway](#) (and integration services such as [AWS Step Functions](#)/[EventBridge](#) as mentioned in a previous section) for both internal API interactions (such as the other OSS modules highlighted, and BSS), as well as external API integrations (e.g., integrating with AWS Partners' orchestrator or B2B digital platform).

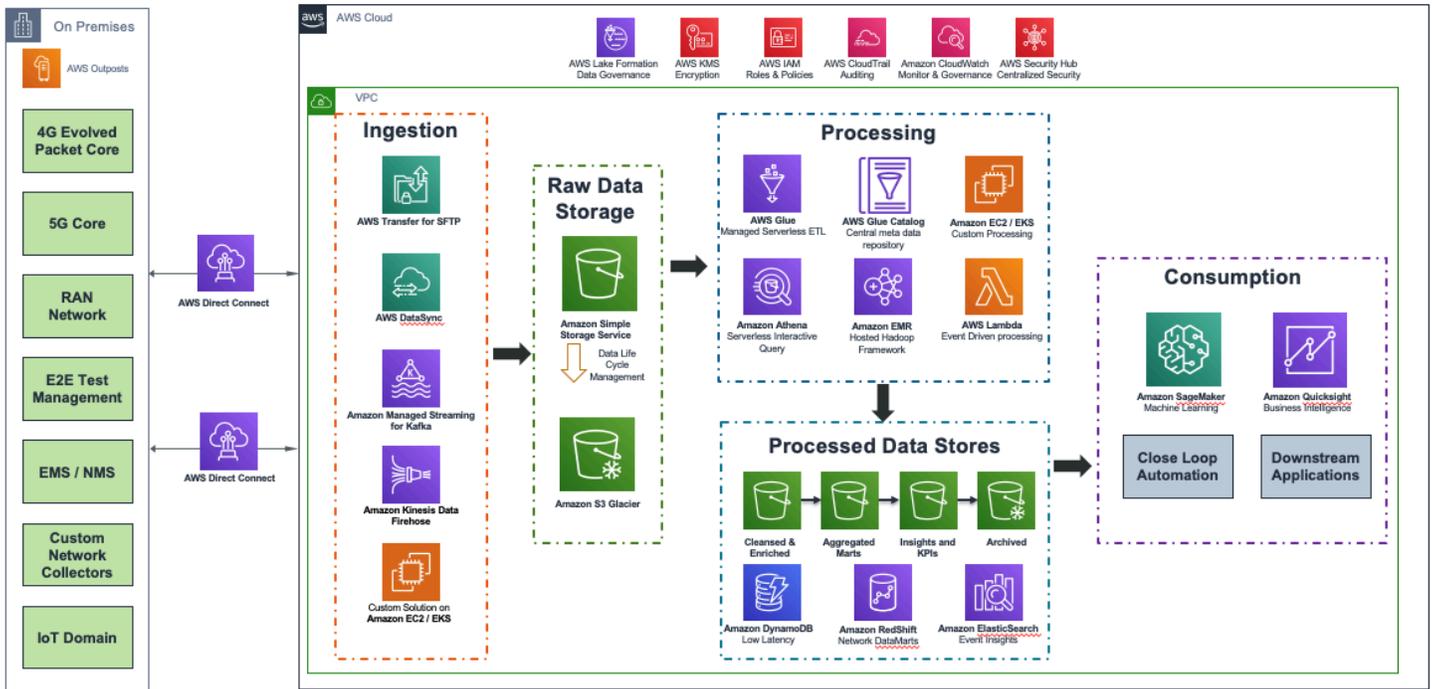
Core orchestration capabilities highlighted in the application layer are comprised of adapters, workflow engine, catalog services, inventory services, and service management, and can be deployed as containers leveraging Amazon EKS in a shared cluster/namespace. With these functional capabilities implemented as Kubernetes jobs, they can be scheduled and scaled, dynamically leveraging [AWS Fargate](#) as a serverless compute engine, thus maximizing resource utilization within the cluster.

Stateful/persistent data (required for drive service orchestration such as inventory, config, and catalog data) are de-coupled from the orchestration process and offloaded to managed database services such as [Amazon RDS](#), [Amazon Aurora](#), [Amazon Neptune](#), and [Amazon Keyspaces \(for Apache Cassandra\) \(Amazon Keyspaces\)](#). As a result, this persistent layer can be scaled independently from the orchestration logic layer, and since it has been de-coupled you can easily enable cross-domain service orchestration use cases.

To support emerging event-driven/policy-driven architecture patterns, the proposed architecture also enables asynchronous flows to be handled by a serverless, event-driven, layer-utilizing native service such as [Amazon SQS](#) for queuing, as well as the following serverless, implementation-leveraging functions: [AWS Step Functions](#), [Amazon EventBridge](#), and [AWS Lambda](#). Over time, we expect more service-orchestration transactions could be moved towards an event-driven approach, which serves to further reduce the deployment footprint of the service orchestration platform.

Network Analytics

This section presents a network analytics architecture on AWS that provides flexibility, scalability, and innovation through Machine Learning (ML) integration. The components to a network analytics solution can be divided in four categories: ingestion, storage, processing and analysis, and consumption. The following reference architecture illustrates the AWS services that support the proposed architecture.



Network Analytics Architecture on AWS

Data can be ingested through [AWS Transfer for Secure File Transfer Protocol \(SFTP\)](#) to periodically collect data from NFX, Domain Managers, Custom Edge collectors, and legacy network performance analytics solutions. Similarly, you can leverage [Kinesis](#) and/or [Amazon MSK](#) to inject real-time performance data such as events-driven messages (for example, UE attach). Kinesis supports real-time data streaming where data collected is available in milliseconds to enable real-time analytics use cases.

Amazon S3 provides flexible, scalable, and performant storage. Amazon S3 enables DSPs to manage data and access controls, query-in-place for analytics, and provide a wide range of cost-effective storage classes. [AWS Lake Formation](#) (Lake Formation) provides an effective, simple way to secure the data lake supporting your network analytics solution. You can use one single data lake for your data, whether it is untransformed network performance data or enriched performance data. You can govern access to the data by allowing read instructions from an operations team to a given table while allowing a development team the ability to alter it. Data lakes provide you with the ability to reduce data duplication by governing what can be consumed and how it can be consumed, and providing one viewpoint of DSP’s performance (and configuration) data.

An [AWS Glue Crawler](#) crawls into your data lake to identify the format and create the tables (or updates) in your Data Catalog. It creates the structure that allows you to query your data. For

example, if an operator initiates ingestion and loads data into the [Amazon S3](#) buckets for a new NFX, DSPs can define the AWS Glue Crawler that will go through the NFX performance data and identify its metadata. Once the [AWS Glue Data Catalog](#) is built using the AWS Glue Crawler, DSPs have the ability to easily query their data using [Amazon Athena](#) (a serverless interactive query service that allows you to analyze data in Amazon S3). DSPs can access their network data on the fly and perform complex SQL queries.

[EMR](#) can be leveraged to process the vast amount of network data. EMR makes it easy for the operator to set up, operate, and scale their big data environment by automating time-consuming tasks (like provisioning capacity and tuning clusters). Similarly, DSPs can leverage [Kinesis](#) to ingest real-time data and run an [AWS Lambda](#) function to transform the ingested data.

DSPs can leverage [Amazon Redshift \(Redshift\)](#) as a data warehouse solution to create specialized views and procedures, and support their network analytics needs. [AWS Glue ETL](#) jobs can be leveraged to create a database schema in Redshift and copy data from Amazon S3 to Redshift.

[Amazon QuickSight](#) makes it easy for DSPs to build dashboards showing the performance of their network, share that information across engineering and leadership groups, and support quick integration with ML-powered insights. QuickSight reads from Redshift, from Amazon S3 through Athena, [etc.](#), making it a great Business Intelligence (BI) tool to correlate data at various stages of a given analysis path.

AWS services integrate easily with existing DSPs' in-house consumption solutions by providing the tools, APIs, and security necessary. For example, DSPs can perform SQL queries towards Redshift to feed into their legacy reporting systems using the same SQL queries used in their current set of queries.

Edge Analytics

Edge Analytics is an extension of the DSP Network Analytics solution. It provides the required analytics capabilities for low latency use cases by providing network functions with the capabilities to make real-time decisions based on users' behaviors and network conditions. NFX such as near real-time RIC and distributed Network Data Analytics Function (NWDAF) require proximity with the network, the ability to scale, and support for innovation such as the deployment of new optimization algorithms.

[AWS Outposts](#) (Outposts) is a service that addresses Edge Analytics use cases. It is a fully managed service that offers the same AWS infrastructure, AWS services, APIs, and tools to DSPs' on-premises edge location. It provides you with a consistent development experience for your network analytics

solution. For example, an [EMR](#) cluster can be spun up on Amazon Outposts that contains 5G Core network functions, and it can feed off performance events created by each NFX to provide a near-real time recommendation to the network.

Data Unification

An OSS Architecture on AWS enables telecommunication providers to consolidate their OSS and BSS stacks, allowing for reduction of the data duplication inherent to legacy architectures. For example, traditional, on-premise architecture required a network inventory to be copied across the entire stack, which often existed in hundreds of systems. Following the best practices and architecture defined in this document, you can consolidate the Inventory Management (IM) solution and have a consumers-based model, allowing you to develop applications without duplicating that data. With each duplication coming at a cost, this reduces your overall solution costs. AWS services provide the flexibility and scale to develop your OSS & BSS applications. The following figure illustrates the idea behind data unification:



Data Unification – Telco Data Lake

In this figure, the concept of Telco Data Lake is introduced. AWS enables you to unify your configuration, performance, and inventory data. This improves your overall operational efficiency, reduces your costs, and enables you to innovate faster. For example, a BSS Analytics solution may require information on the network availability to enable a service request requiring additional QoS at additional costs. That same information may also be used by the Service Assurance solution to build a prediction model supporting network optimization. AWS services allow you to consolidate that information, scale it, and govern its access.

Moreover, with the advent of 5G networks, the classical division between OSS and BSS is blurred. The mechanisms to monetize and configure the network now need to happen in seconds, dynamically, and automatically. To realize the benefits of 5G, the OSS and BSS architectures should align with data unification concepts.

Security

AWS services provide the necessary framework to secure your OSS solution and the network it manages. This section discusses the AWS services that can help you secure your solution.

Security of the OSS Solution

[The security pillar of the AWS Well-Architected Framework](#) provides guidance in developing secure applications and providing the best practices and AWS services recommendations to achieve security excellence.

[Amazon VPC](#) allows the creation of private networks and control access to the OSS Solutions using subnets, security groups that are stateful, and Network Access Control Lists (NACL) that are stateless. This enables the isolation of OSS applications from one another, from network elements, and from business and IT applications, ensuring only specific access is allowed.

OSS application developers can leverage [AWS Key Management Service](#) (KMS) to create and manage cryptographic keys for data-at-rest encryption for the AWS services discussed previously (such as [Amazon S3](#), [Amazon EBS](#), [Amazon RDS](#), [Redshift](#), [Amazon ElastiCache](#) (ElastiCache), etc.).

Similarly, OSS applications can leverage [AWS Directory Service](#) to integrate and federate with existing corporate directories to reduce administrative overhead and improve end-user experience. This simplifies CSPs and DSPs' desired Single Sign On (SSO) for their entire application spectrum, inclusive of network workloads such as OSS.

[AWS CloudTrail](#) (CloudTrail) provides a history of AWS API calls, allowing for identification of source IPs for attempted AWS services access. [CloudWatch](#) Logs allows for a centralized view of all OSS

application logs. It makes it easy to search for specific error codes or patterns while providing a highly-scalable service, and it helps you identify operational mistakes.

Security of the network functions

Traditional OSS solutions provide the Public Key Infrastructure (PKI) necessary for the encryption of OAM and network traffic. Monolithic applications from different ISVs required a high level of operational overhead: Many disparate PKIs existed and a complex hierarchical relationship of the various PKIs. [KMS](#) makes it easy to create and manage cryptographic keys, and provides native integration with AWS CloudTrail to provide you with logs of all key usage. This allows the operator to know what application is being used, and what organization and what users leverage a given key. Various options are available, and they are inclusive of the ability to import your own 256-bit symmetric key. This simplifies your ability to, and increases your control in, encrypting data at rest and in transit, such as configuration data in Amazon S3.

[AWS Certificate Manager](#) (ACM) is a service that simplifies the provisioning, management, and deployment of Secure Sockets Layer (SSL) / Transport Layer Security (TLS) certificates. ACM Private Certificate Authority (CA) enables telecommunication service providers to create a complete CA hierarchy, allowing for a common root and sub-hierarchy for different organizations, traffic-related encryption, and non-traffic data encryption. For example, one sub-CA can be used for encryption of S1U interfaces, while another sub-CA can be used for encrypting domain manager FM interfaces. This reduces the number of CAs managed by a DSP, reducing the cost paid for CAs, supports API-based automation for programmatic deployment, and simplifies the management of Certificate Revocation List (CRL).

Connectivity

[Direct Connect](#) makes it easy to establish a dedicated connection from a DSP on-premise network to its AWS VPCs, inclusive of VPCs running their OSS workloads. This provides a consistent network experience to support the transfer of network OAM data. DSPs can combine Direct Connect with [AWS VPN](#) to provide an end-to-end secure IPsec connection.

Amazon VPC supports VPC sharing across accounts, allowing you to isolate OSS workloads from network workloads, and enabling the creation, modification, and deletion of OSS applications, in a collocated manner, to network workload without the ability to view, modify, or delete network resources. Network topologies are simplified by interconnecting shared Amazon VPCs using connectivity features, such as [AWSPrivateLink](#), transit gateways, and VPC peering.

Digital Transformation and DSP Enablement

With connectivity ubiquity, CSPs are venturing in new spaces, and connecting things in support of Industry 4.0, delivery of content, and improving their users' experience. This is particularly compelling with the advent of 5G as it tackles latency limitations by using cloud technologies concepts and leveraging computing advances. This enables a new wave of services to be provided to end users such as AR/VR gaming, remote surgeries, and connected cars. CSPs will no longer solely provision basic network connectivity services; they will provision a wider range of services such as remote surgery, intelligent driving, smart factories, etc., becoming **DSPs**.

To enable the transformation from CSP to DSP, traditional business and operational processes need to evolve. This is achieved through the following key enablers:

- **Digitization of the customer touchpoints** to transform all of CSPs' customer touchpoints (retail and enterprise), and adoption of a digital-first strategy. In short, all of the CSP services being offered are accessible through digital mediums and exposed over Machine to Machine (M2M) interfaces for consumption and automation.
- **Offer digital services to customers** to augment CSPs' core traditional Telco offering of voice, messaging, and data with not just advanced media services and applications, but also exposed network services to industry/enterprises to enable innovative services and monetize 5G networks. This requires a multi-service OSS automation platform to manage and operate this ecosystem.
- **Modernization of the underlying network** to enable the elasticity and scalability required for offering such digital services with high efficiency and low cost. The architectural innovation in 5G allows that, but it requires full automation of the operating tools and processes.
- **Data-driven network** to enable Artificial Intelligence (AI) and ML-based business and network operations. There is a huge amount of data generated by all the network, IT, and enterprise workloads, which provides tremendous possibilities to drive efficiencies and improve customer experience, as well as open new revenue streams. This would require setting up a secure unified data lake at the center of the network.

Our proposed OSS Architecture on AWS helps you deliver on these key enablers. It also helps CSPs to adopt an automation framework that can manage both the IT and telecom workloads, across all the deployment models and vendors, to provide **a single unified view** and **a single unified API** for their entire network.

Conclusion

AWS provides the capabilities for you to build a next-generation OSS solution that attains scalability, elasticity, and high-availability, all while providing a framework to improve operational efficiency, reduce OSS costs, and unlock the deployment of 5G services. Several customers are using AWS and AWS Partner Network (APN) partners to transform their OSS stack and digitalize their operations. AWS has the broadest and strongest partners in the ecosystem (especially for partner solutions) available through AWS Marketplace and the APN Partner Central for each of OSS functions presented in this paper.

The reference architectures and best practices provided in this whitepaper can help you develop OSS workloads on AWS, providing you with the framework on which to build solutions that are agile, elastic, scalable, cost-efficient, secure, and that allow you to innovate faster.

Developing and running OSS solutions on AWS helps CSPs transition to DSPs by providing the framework to develop a programmatic network, fully automated and fully integrated, and by enabling end-to-end 5G services that meet the specificity of your end users.

To learn more about how telecommunications companies are leveraging AWS services, visit [Telecommunications on AWS](#).

Contributors

Contributors to this document include:

- Aymen Saidi, Principal Partner Solutions Architect, WW Telco Partner, Amazon Web Services
- Visu Sontam, Senior Partner Solutions Architect, WW Telco Partner, Amazon Web Services
- Owen Law, Senior Partner Solutions Architect, WW Telco Partner, Amazon Web Services
- Srikrishna Komatineni, Principal Solutions Architect, Telco Accounts, Amazon Web Services
- Mohit Gupta, Principal Portfolio Lead BSS/OSS, WW Telco Partner, Amazon Web Services

Further reading

For additional information, see:

- [Introduction to AWS Security](#)
- [5G Network Evolution with AWS](#)
- [Carrier-Grade Mobile Packet Core Network on AWS](#)
- [A Modern and Simple Approach to Address CSP's Network Performance Analytics Challenges Using AWS](#)
- [Continuous Integration and Continuous Delivery for 5G Networks on AWS](#)

Document history

To be notified about updates to this whitepaper, subscribe to the RSS feed.

| Change | Description | Date |
|-------------------------------------|-----------------------------|--------------------|
| Minor update | Fix non-inclusive language. | April 6, 2022 |
| Initial publication | Whitepaper first published. | September 21, 2021 |

Note

To subscribe to RSS updates, you must have an RSS plug-in enabled for the browser that you are using.

Glossary

- **5G NR:** 5G New Radio
- **AI:** Artificial Intelligence
- **AMF:** Access and Mobility management Function
- **API:** Application Programming Interface
- **AZ:** Availability Zone
- **BI:** Business Intelligence
- **BSS:** Business Support Systems
- **CA:** Certificate Authority
- **CDMA:** Code Division Multiple Access
- **CM:** Configuration Management
- **CRL:** Certificate Revocation List
- **CSPs:** Communication Service Providers
- **DSPs:** Digital Service Providers
- **eNodeB:** E-UTRAN Node B, also known as Evolved Node B (abbreviated as eNodeB)
- **ETL:** Extract, Transform, Load
- **eTOM:** enhanced Telecom Operations Map
- **ETSI MANO:** European Telecommunications Standards Institute (ETSI) Management And Orchestration (MANO)
- **FCAPS:** Fault management, Configuration management, Accounting, Performance management and Security
- **FM:** Fault Management
- **GSM:** Global System for Mobile Communications
- **IM:** Inventory Management
- **ISV:** Independent Software Vendor
- **KPIs:** Key Performance Indicators
- **LCM:** Life Cycle Management
- **LTE:** Long Term Evolution
- **M2M:** Machine to Machine

- **ML:** Machine Learning
- **MME:** Mobility Management Entity
- **MVNO:** Mobile Virtual Network Operator
- **NACL:** Network Access Control List
- **NEPs:** Network Equipment Providers
- **NFV:** Network Function Virtualization
- **NFx:** Network Functions
- **NWDAF:** Network Data Analytics Function
- **OAM:** Operations, Administration and Maintenance
- **ORAN RIC:** Open Radio Access Network (ORAN) RAN Intelligent Controller (RIC)
- **OSI:** Open Systems Interconnection
- **OSR:** Operations Support and Readiness
- **OSS:** Operational Support System
- **PKI:** Public Key Infrastructure
- **PM:** Performance Management
- **QoS:** Quality of Service
- **RDBMS:** Relational Database Management System
- **RTO / RPO:** Recovery Time Objective / Recovery Point Objective
- **SaaS:** Software as a Service
- **SNMP:** Simple Network Management Protocol
- **SO:** Service Orchestration
- **SON:** Self-Organizing Network
- **SQL:** Structured Query Language
- **SSL / TLS:** Secure Sockets Layer (SSL) / Transport Layer Security (TLS)
- **SSO:** Single Sign On
- **TDMA:** Time-division Multiple Access
- **UMTS:** Universal Mobile Telecommunications System
- **UPF:** User Plane Function
- **URLLC:** Ultra Reliable Low Latency Communications
- **VPC:** Virtual Private Cloud

- **VR / AR:** Virtual Reality / Augmented Reality
- **WiMAX:** Worldwide Interoperability for Microwave Access

Notices

Customers are responsible for making their own independent assessment of the information in this document. This document: (a) is for informational purposes only, (b) represents current AWS product offerings and practices, which are subject to change without notice, and (c) does not create any commitments or assurances from AWS and its affiliates, suppliers or licensors. AWS products or services are provided “as is” without warranties, representations, or conditions of any kind, whether express or implied. The responsibilities and liabilities of AWS to its customers are controlled by AWS agreements, and this document is not part of, nor does it modify, any agreement between AWS and its customers.

© 2021 Amazon Web Services, Inc. or its affiliates. All rights reserved.

AWS Glossary

For the latest AWS terminology, see the [AWS glossary](#) in the *AWS Glossary Reference*.