

# AWS Architecture Monthly



5G  
June 2021



## Editor's note

In our 5G issue, we discuss the evolution, implementation and trends for 5G. We've included some case studies from Verizon and Ericsson, and take a look at Ribbon SBC SWe and network-slicing with Cloudify. There are detailed Reference Architectures for you to check out, and videos on DISH, Verizon, AWS Wavelength, VR, Automotive, and Telefónica.

We'd like to thank Young Jung, our Expert, for his contribution to this issue. We hope you find it useful and informative.

Please give us your feedback! Include your comments on the [Amazon Kindle](#) page. You can [view past issues](#) and reach out to [aws-architecture-monthly@amazon.com](mailto:aws-architecture-monthly@amazon.com) anytime with your questions and comments.

*Jane Scolieri, Managing Editor*

## Table of Contents:

- [Ask an Expert](#): Young Jung, Ph.D., Principal Solutions Architect, 5G & NFV
- [Case Study](#): Verizon Accelerates 5G Network Capability Launch Using AWS
- [Quick Start](#): Ribbon SBC SWe on AWS
- [Blog](#): Implementing 5G Network Slicing with Cloudify on AWS
- [Whitepaper](#): 5G Network Evolution with AWS
- [Reference Architecture](#): Deploying 5G Core on AWS
- [Case Study](#): Ericsson Delivers Low-Latency Communication with 5G at the Edge Using AWS Outposts
- [Reference Architecture](#): Deploying E2E 5G Network with AWS
- [Blog](#): Open source mobile core network implementation on Amazon Elastic Kubernetes Service
- [Whitepaper](#): Continuous Integration and Continuous Delivery for 5G Networks on AWS
- [Videos](#): DISH; Verizon; AWS Wavelength; VR Adoption; All Things Automotive, Telefónica

## 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.

# Ask an Expert:

## Young Jung, Ph.D., Principal Solutions Architect, 5G & NFV

### ***What are the general architecture pattern trends for 5G in the cloud?***

We've seen the telecom industry evolve their networks from the initial physical bare-metal system to a virtual network function. Now, most companies operate container network functions within a network function virtualization. But industry standards are always progressing. As they move towards 5G networking, we see telecom shifting to using function decomposition (such as Open-RAN and Control-User Plane Separation), distributed edge (because of high-capacity demand), and service-based architecture that uses RestAPI rather than legacy protocols. 5G architecture trends embrace cloud-native microservice architecture. Kubernetes-based containers and NoSQL databases for stateless architecture are commonly implemented in telecom networks. This pattern consistently applies to much of 5G in the telecom industry, such as operations support systems/business support systems, core networks, and radio access networks. Another architecture pattern in 5G is to have more logical and virtual separation of the network. This means having multiple logical network slices for all the various enterprise specific use cases.

### ***When putting together an AWS architecture to solve business problems specifically for 5G customers, what are some of the considerations?***

Because the 5G network provides important services such as voice calling and mobile broadband, networks and applications require much more high availability, security, scalability, and ease of operation and automation than other industries.

AWS global infrastructure, built around Regions and Availability Zones, provides a baseline high availability of platform and infrastructure layer. To take advantage of this benefit, we recommend using stateless, active-active, or active-standby architectures over multiple AZs and Regions. Our data centers provide customers a network that protects 5G subscribers' important information, identities, applications, and devices. The AWS [shared responsibility model](#) of security applies to 5G network and IT applications on AWS, ensuring our customers follow our [security principles](#).

While high availability and security are required for telecom network and IT workloads, scalability and zero-touch automation are also becoming crucial in the 5G era. This is because of the demand for high-capacity networks, flexible network slicing, and complexity of microservice architecture. To help meet this demand, AWS provides hyperscaling compute, networking, and storage-like infrastructures as well as services for programming-based orchestration and automation tools.

### ***Do you see different trends in 5G in cloud versus on-premises?***

Yes, functionality, scalability, flexibility, and automation can be constrained by underlying infrastructures and capabilities of network function providers in on-premises. However, in AWS, communication service provider (CSP) s can build a more automated 5G network and service by using various AWS services for additional programming-model-based, closed-loop orchestration, such as [AWS Lambda](#), [AWS Step Functions](#), [Amazon CloudWatch Events](#), and [Amazon API Gateway](#), as well as [CI/CD DevOps tools](#).

## **What's your outlook for 5G, and what role will cloud play in future development and deployment efforts?**

In the past, telecom companies used to mostly serve mobile voice and broadband internet services. The 5G network will serve enhanced mobile broadband and ultra-reliable low latency and massive machine-type communication. It will have multiple network slices, which requires the network to be more agile, scalable, and flexible.

AWS helps CSP customers focus on new services by preparing, installing, and staging compute servers, storage, and networking devices. Moreover, AWS provides an API programming model-based orchestration. You can create a virtual data center environment and deploy automated network functions inside of it. CSP customers can fully take advantage of these services to bring agility and fast innovation. In addition, Regions and Availability Zones, [AWS Outposts](#), [AWS Local Zones](#), and [AWS Wavelength](#) can provide a benefit of hyper-scalability in their network vertically and horizontally, which helps you monetize your 5G network for various enterprise and consumer use cases.

## **Anything else you'd like to add?**


5G networks will be moving toward closed-loop automation powered by Artificial Intelligence/ Machine Learning (AI/ML). AWS offers multiple tools that our telecom customers can use to quickly create machine learning solutions. You can build a data lake to collect metrics, alarm, fault, and other input information. Our data lake solution automatically configures the core AWS services necessary to easily tag, search, share, transform, analyze, and govern specific subsets of data across all network domains. This collected data can be trained as an input for machine learning. It can then provide a trigger to the orchestration to perform predictive and proactive measurements. When you consider the complexity in the 5G network (because of its various use cases and slice and microservice architecture), this path is likely inevitable. AWS will provide momentum for an AI/ML driven network evolution for 5G and also for the future - beyond the 5G network.

### **Young Jung**

Dr. Young Jung is a Principal Partner Solutions



Architect in the AWS Worldwide Global Telecom Partner Alliance team. He specializes in the NFV space, and works with various global telecom partners. He has been leading cloud-adoption of telco industry's 4G/5G implementation on the AWS environment.



# Case Study:

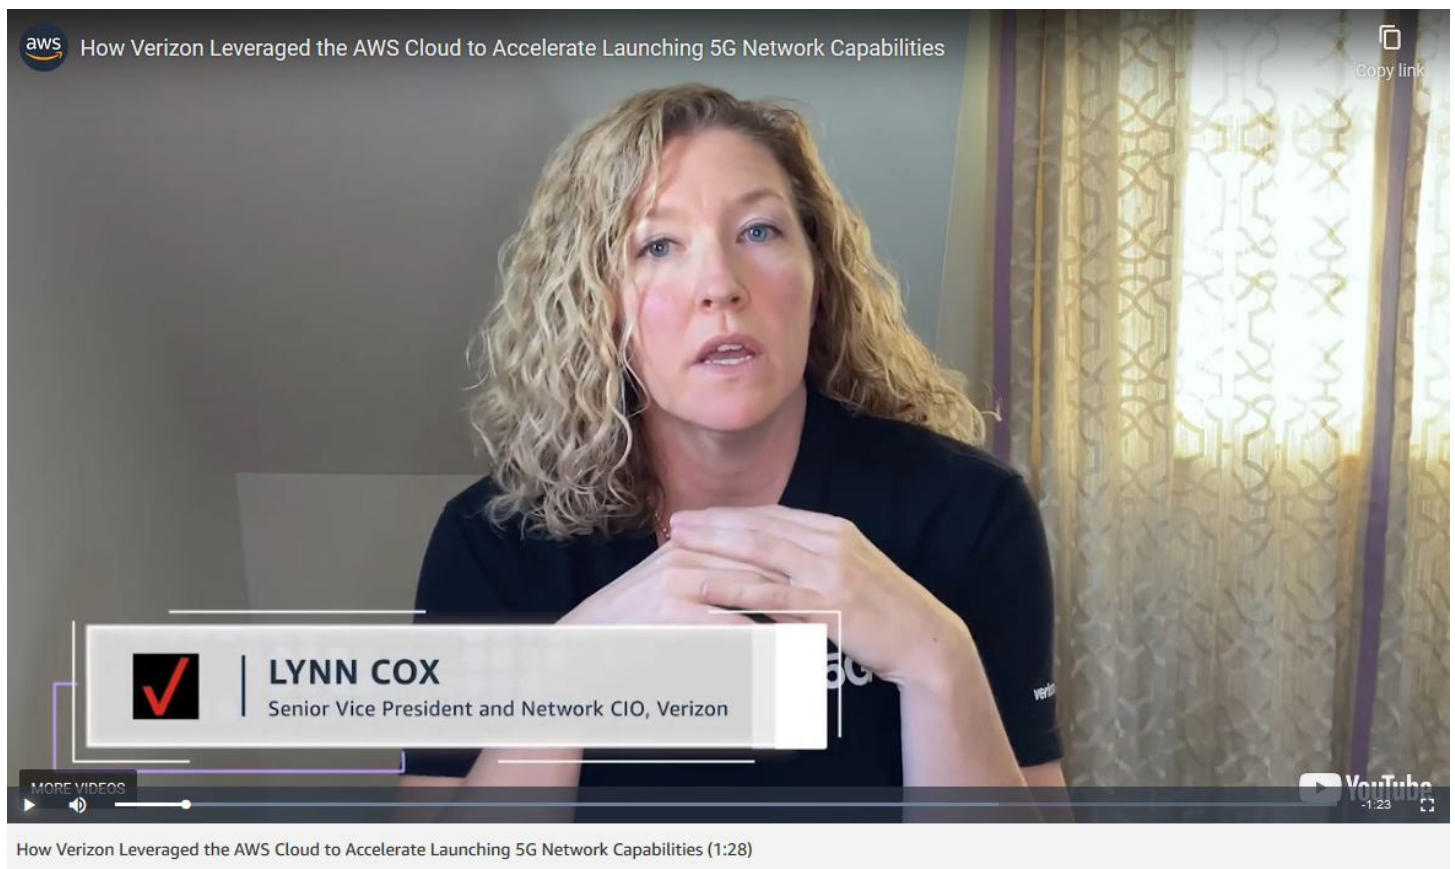
## Verizon Accelerates 5G Network Capability Launch Using AWS

2021

Verizon Communications, Inc. is the first company in the world to launch a commercial 5G mobile network with a commercially available 5G-enabled smartphone. In this video, Lynn Cox, Verizon senior vice president and network CIO, explains how leveraging Amazon Web Services (AWS) to support its 5G network rollout helped the company maximize its decision-making velocity and take advantage of large-scale compute capabilities of the cloud. "If we had done it the old way, which was with in-house servers, the cost to do it would have been astronomical, but also the time," says Cox. "That's where we turned to our partnership with AWS to maximize the speed with which we could make decisions and run these scenarios to come up with the best plan for our 5G network."

Learn more about AWS for Telecom at [aws.amazon.com/telecom](https://aws.amazon.com/telecom).

[View case study online](#)



The image shows a YouTube video player interface. At the top left, the AWS logo is followed by the video title: "How Verizon Leveraged the AWS Cloud to Accelerate Launching 5G Network Capabilities". In the top right corner, there is a "Copy link" icon. The main video frame shows Lynn Cox, a woman with blonde curly hair, wearing a black Verizon polo shirt, speaking with her hands clasped. A lower-third graphic identifies her as "LYNN COX, Senior Vice President and Network CIO, Verizon". At the bottom left of the video frame, there is a "MORE VIDEOS" button. The bottom of the player shows a progress bar and the YouTube logo. Below the video frame, the text "How Verizon Leveraged the AWS Cloud to Accelerate Launching 5G Network Capabilities (1:28)" is displayed.

# Quick Start:

## Ribbon SBC SWe on AWS

*Secure multimedia session control and efficient, fault-tolerant real-time communications*

This Quick Start deploys the Ribbon Session Border Controller Software Edition (SBC SWe) on the Amazon Web Services (AWS) Cloud.

Use Ribbon SBC SWe to secure real-time communications, including unified communications, conferencing and collaboration, and contact center services. SBC SWe provides robust interworking and normalization for multiple signaling and media protocols, call admission control to manage traffic levels, and security features for both signaling and media to protect privacy and to help ensure regulatory compliance.

Ribbon SBC SWe consists of the following components:

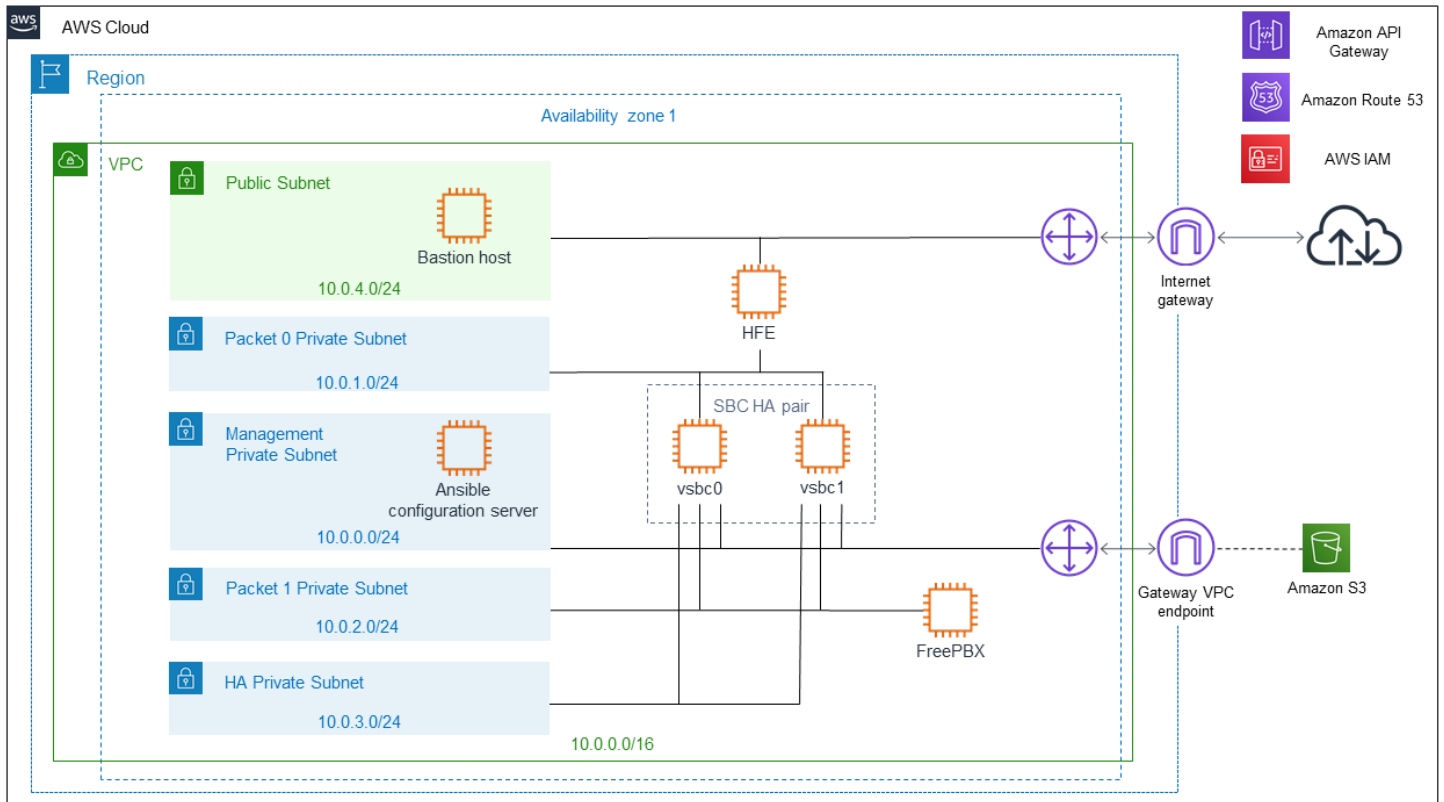
- An integrated session border controller as a Session Initiation Protocol (SIP) back-to-back user agent (B2BUA).
- A high-availability front-end (HFE) node to assist with improved media restoration times in the event of an SBC instance failure.
- A FreePBX application server that acts as a SIP registrar.
- An Ansible control node to launch playbooks that configure the session border controller through a Representational State Transfer (REST) application programming interface (API).

Deploying the Ribbon SBC SWe on AWS delivers a solution that is elastic and fault-tolerant. It uses compute resources efficiently for cost control, and it dynamically adjusts to varying loads of traffic.

Use this Quick Start to automatically set up the following environment on AWS:

- A highly available architecture capable of handling application-level failure within a single Availability Zone.\*
- A virtual private cloud (VPC) configured with public and private subnets according to AWS best practices, to provide you with your own virtual network on AWS.\*
- In the public subnet:
  - A Linux bastion host that allows inbound Secure Shell (SSH) access to the SBC instances in the management private subnet.\*
  - A managed network address translation (NAT) gateway to allow access to the Amazon Elastic Compute Cloud (Amazon EC2) API.\*
- In the private subnets:
  - SBC traffic management, which includes an Ansible configuration server to launch playbooks that configure the SBC.
  - A high-availability (HA) subnet for mirroring and synchronization of traffic between the two SBC instances.
  - A core (trusted media) subnet to handle signaling and media between the SBC and an application server or registrar (FreePBX).
  - An access (untrusted media) subnet to handle signaling and media between SBC and the HFE node.
- A gateway VPC endpoint to allow Amazon Simple Storage Service (Amazon S3) access from the EC2 instances in a private subnet.\*
- Two SBC instances that form an SBC HA pair with interfaces for the private subnets.
- An HFE node to improve SBC instance failover performance.
- A FreePBX application server that also acts as an SIP registrar.
- An S3 bucket that contains playbooks and scripts to configure HFE, SBC, and FreePBX.
- VPC endpoints to access the S3 bucket and other services over the AWS private network.

*\* The template that deploys the Quick Start into an existing VPC skips the components marked by asterisks and prompts you for your existing VPC configuration.*



[See the source code for this Quick Start](#)

[View Deployment Guide](#)

*This Quick Start was developed by Ribbon in collaboration with AWS. Ribbon is an APN Partner.*



# Blog:

## Implementing 5G Network Slicing with Cloudify on AWS

by Rabi Abdel and Nati Shalom

Recent advances in 4G and 5G cellular technologies have led to the adoption of virtualization of network functions, also known as Network Function Virtualization (NFV). Moving away from monolithic and vendor-specific equipment to a separation between hardware and software, 4G and 5G network functions can be hosted on general-purpose hardware (on-premises or in the cloud). Additionally, cloud-native principles of microservices-based architectures where network service elements can scale out independently based on demand is becoming the new paradigm. This trend has opened up the telco partner ecosystem to new players from ISVs, System Integrators (SI), and third-party managed services providers.

Management and orchestration (MANO) is telco industry-defined terminology and a key element of the [ETSI Network Functions Virtualization \(NFV\) architecture](#). It is responsible for allocating and managing network resources for the lifecycle management of virtual network functions (VNFs), Cloud Native Network Functions (CNFs), and End-to-End Network Services (NS). MANO consists of NFV Orchestrator (NFVO), VNF Manager (VNFM), and Virtual Infrastructure Manager (VIM). The need for MANO solutions has risen with the evolution of 5G due to its service-based architecture, cloud-native components, and the scalability requirements that come with it.

In order for MANO solutions to deliver on its expected business outcomes for 5G, they need to oversee the entire operator's network: Radio Access Network (RAN), Edge, Core, Data Network, and Centralized Cloud. This makes the AWS Cloud a perfect solution for 5G MANO due to AWS having a wide range of integrated services that are targeting different parts of an operator's network such as:

- AWS Snow Family for RAN and Edge deployment
- AWS Outposts for Edge in addition to Core deployments
- AWS Wavelength for Data Network and IT Data Centers
- AWS Managed Services in AWS Regions to provide a single pane of glass for monitoring, management, and automation

One of the prime use cases of 5G that requires a great level of management and orchestration is 5G network slicing. 5G network slicing enables the operator to deploy multiple independent virtual E2E networks over the same infrastructure. Each slice can be customized for different services and/or business cases and can form a collection of 5G network functions for each specific use case and/or business model, supporting different requirements from all vertical industries. 5G slicing can span across all domains of the network: 5G device (UE), Access/RAN, Transport network/SDN, Core, and Data network, which could be Multi-access Edge Computing (MEC).

A key slicing concept is to be an isolated and self-contained entity that has all the functions and capabilities chained together to best meet all needs, services, and use cases.

Network Slice Instance (NSI) realized the network slicing concept by comprising a group of network functions, resources, and connection relationships and typically covers multiple 5G use cases. Three fundamental network slice types have been identified for 5G: enhanced Mobile Broadband (eMBB), massive Machine Type Communications (mMTC), and Ultra-Reliable Low Latency Communications (URLLC).

Examples of use cases that can take advantage of Network Slicing are:

- Virtual experience and media use cases, such as 4/8K real-time delivery, smart home/office/city, AR/VR and broadcast services
- Massive connectivity use cases, such as smart metering for gas and electricity, smart environment, smart farms and fishery, smart grid and sensor networks, personal and wearable devices
- Mission-critical use cases, such as autonomous driving, robotic applications, remote controlled machines, Industry 4.0, tactile internet, public safety networks, eHealth and remote surgeries

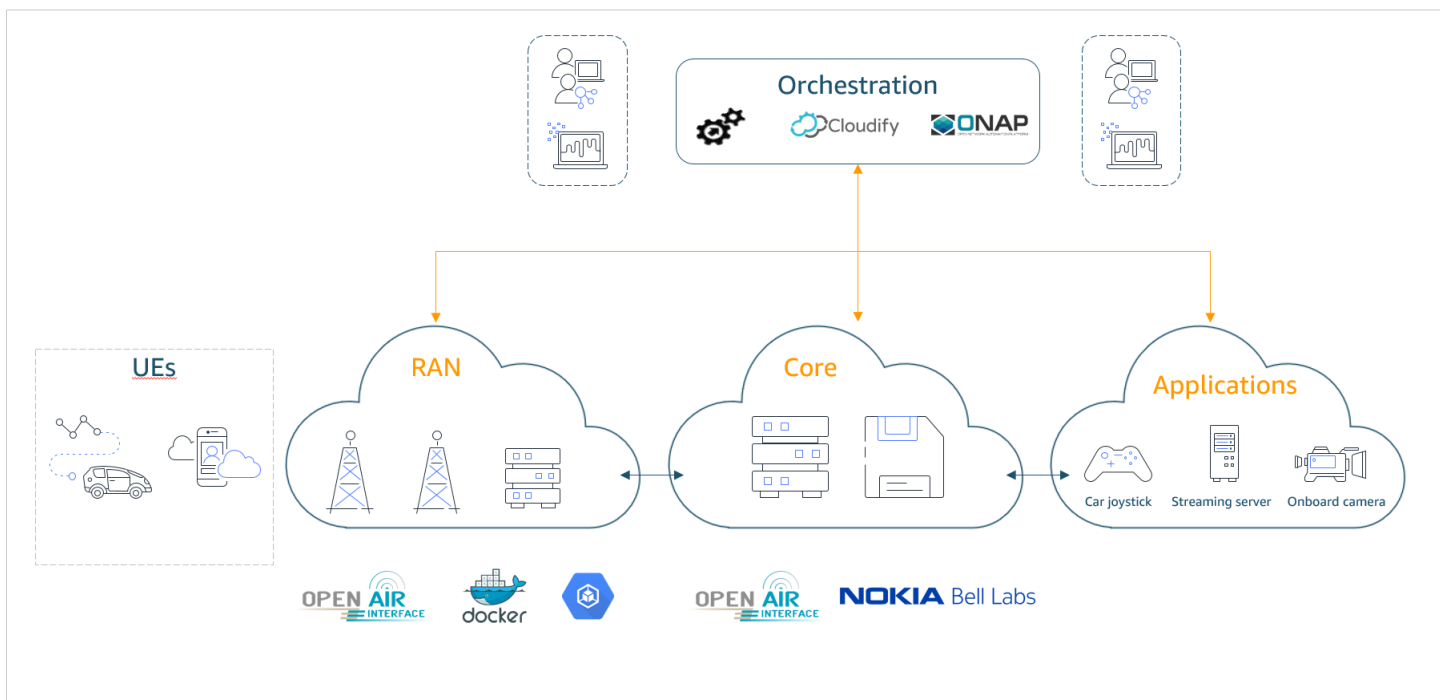
In this post, we describe in detail how [Cloudify](#)'s management and orchestration solution runs on AWS to create and manage the lifecycle of 5G network slices.

## The Cloudify Solution

Cloudify is an open-source, multi-cloud, and edge orchestration platform that allows organizations an effortless transition to the AWS Cloud and to cloud-native architecture by enabling them to automate their existing infrastructure alongside cloud native and distributed edge resources.

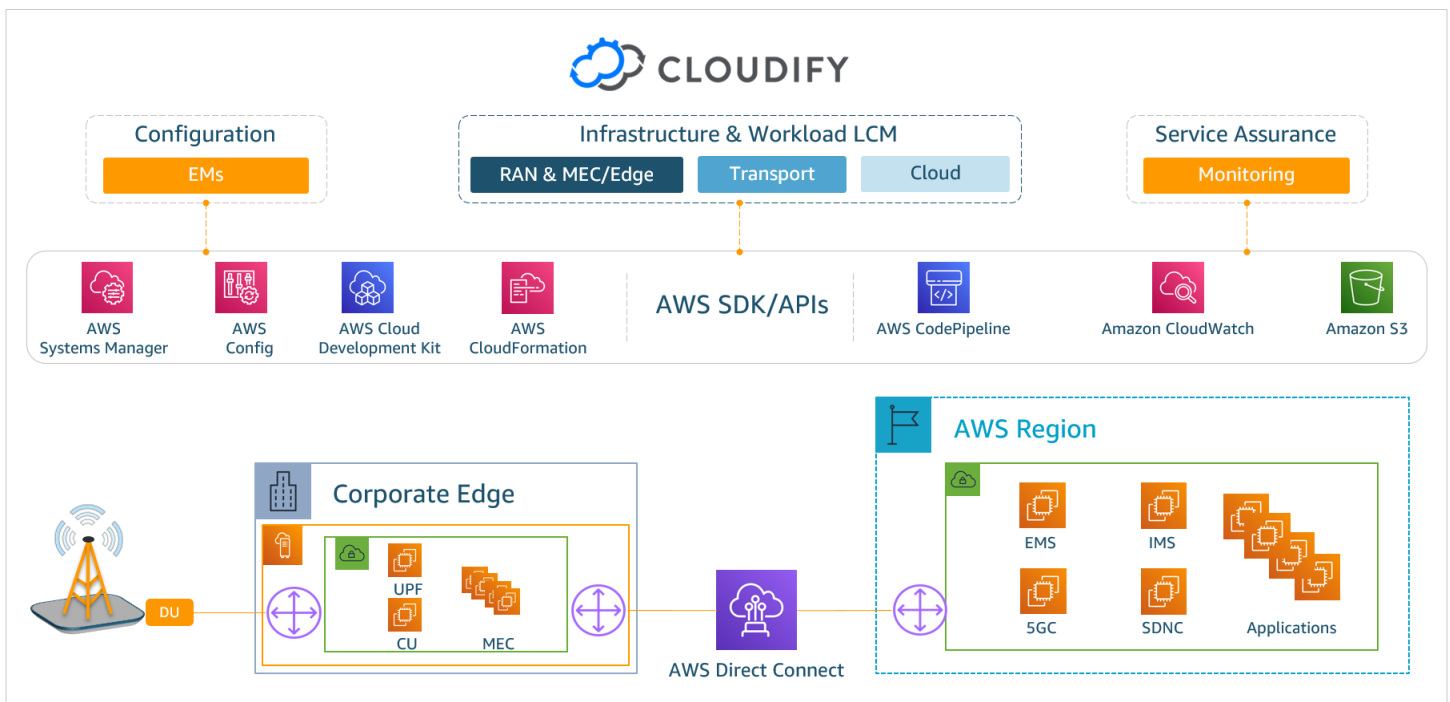
Cloudify is used by many tier-1 carriers as a way to manage their network services and cloud infrastructure. Cloudify is also being used by VNF vendors as a VNF Manager (VNFM), to manage their specific network domains.

Cloudify's flexibility to work with both existing on-premises infrastructure, managed network services such as SD-WAN, EPC, IMS, in addition to modern cloud-native services has been key to its ability to deliver successful NFV transformation projects to production. Cloudify is used by leading open-source community projects such as ONAP to demonstrate [a fully open 5G network slicing](#) as shown in the following **Figure 1**.



**Figure 1:** Use of Cloudify with ONAP to deliver full 5G network slicing.

Cloudify comes with native support for many AWS services including Amazon EKS, Amazon RDS, AWS Lambda, AWS CodePipeline in addition to AWS CloudFormation. The combination of native integration with AWS resources, in addition to support for AWS CloudFormation allows Cloudify to work with many AWS services without depending on specific resource mapping and integration as shown in the following **Figure 2**.



**Figure 2:** Cloudify architecture on AWS.

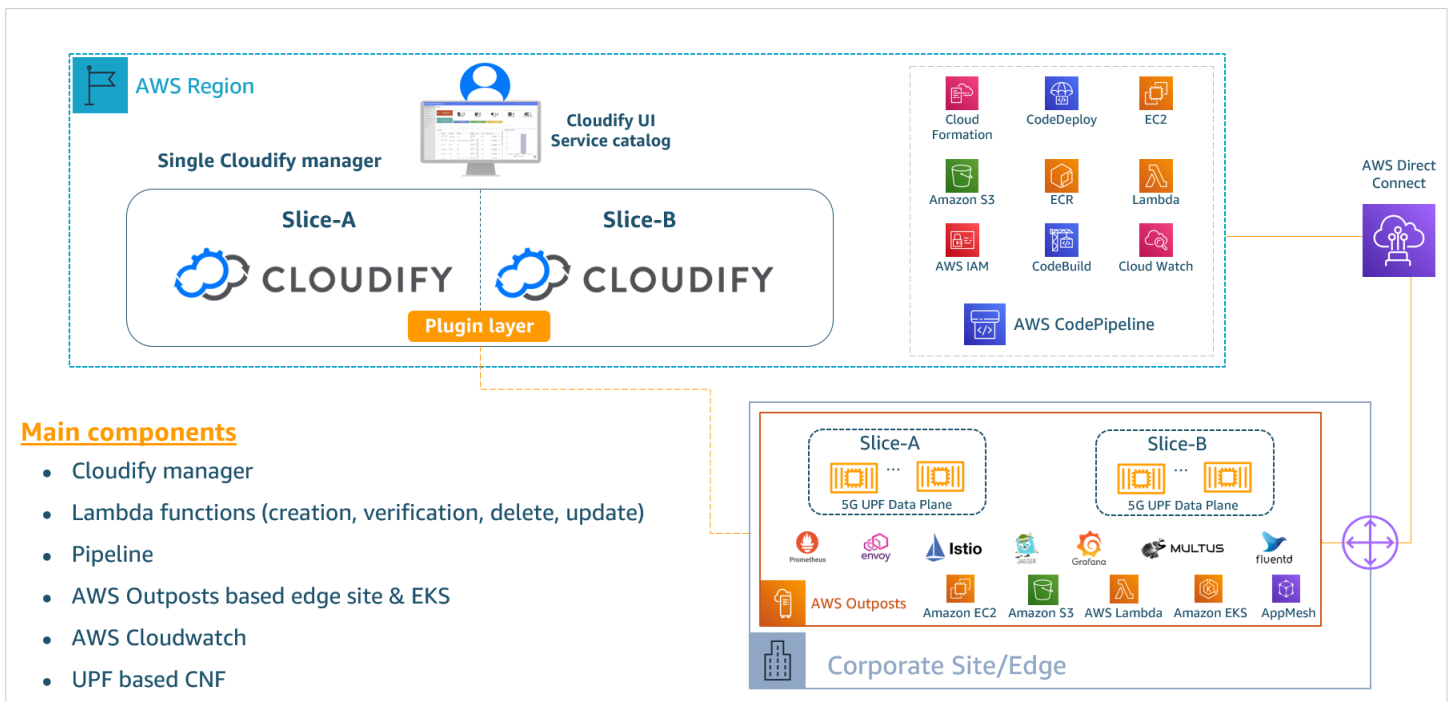
## Cloudify and AWS for Network Slicing

In order for a MANO solution to deliver 5G network slicing capabilities, it has to be able to create multiple network slices and manage lifecycles. This includes: slice template creation (design phase), slice instantiation (deployment phase), slice scaling, configuration management, and termination (operation phase).

AWS CodePipeline is used to fully manage the creation, deletion, management, and configuration of 5G core instances on Edge sites (AWS Outposts). Cloudify integrates to AWS CodePipeline via a set of plugins using an AWS Lambda set that acts as the northbound interface of AWS CodePipeline toward the Cloudify system. This enables the setup and configuration of the network slice from the operator view to the underlying core network as illustrated in the following **Figure 3**.

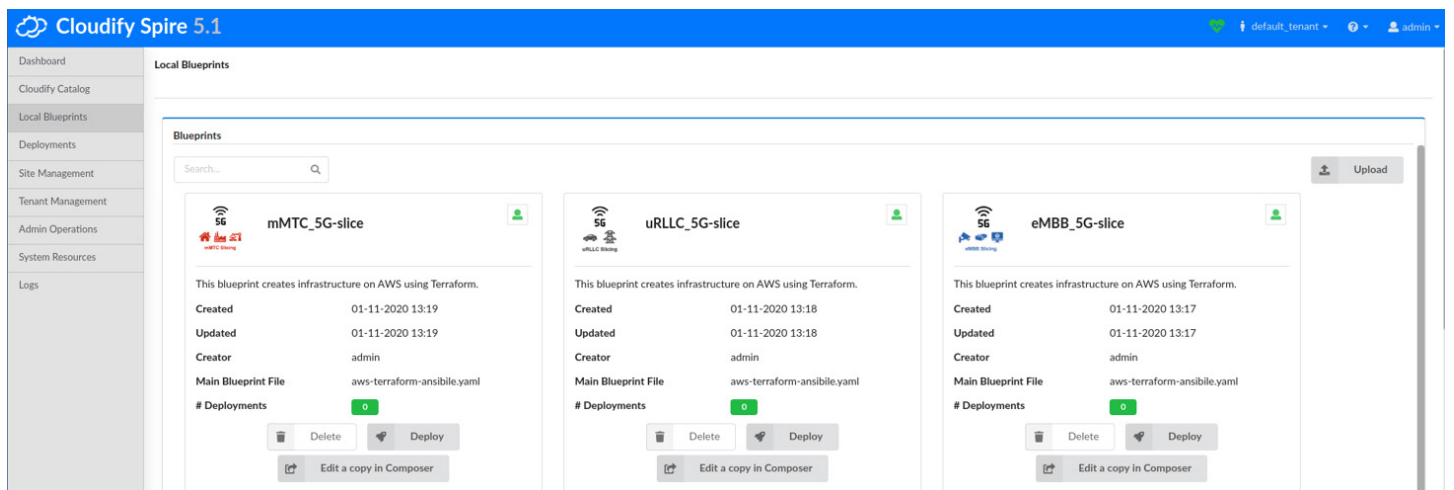
Each Lambda function is mapped to a specific AWS CodePipeline service and each AWS CodePipeline service is mapped to the lifecycle operation (create/delete/modify/verification) of a network slice.

A 5G mobile core slice domain (user plane function (UPF)) is created by instantiating a CNF on different Amazon EKS edge sites (Running on AWS Outposts). A process that is repeated for each network slice.



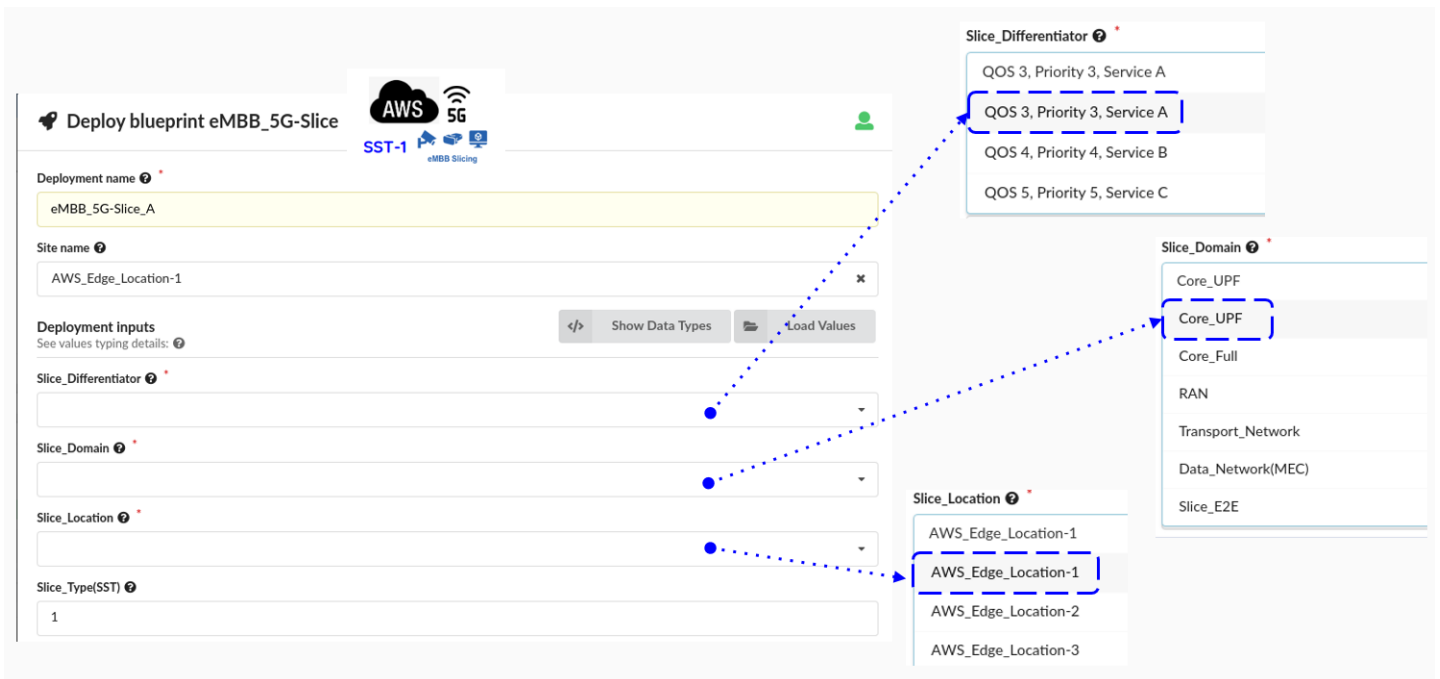
**Figure 3:** Cloudify on AWS for network slicing.

One of the key concepts behind this layered architecture demonstrated previously is to abstract the detailed orchestration workflow and all the network configuration associated with it from the operator. In this context, the Cloudify catalog service is used to allow a simple interface to run network slicing as illustrated in the following screenshot.



**Figure 4:** Providing a simple operator interface using Cloudify catalog service.

As part of slice creation, the user needs to fill several inputs which traverse into values on the fulfillment phase such as slice edge location, slice differentiator, slice domain (Core\_UPF, Core\_Full, RAN, MEC, Transport, Slice E2E). Slice execution includes Amazon EKS cluster deployment, UPF instantiation/licence activation/baseline config/day two config.



**Figure 5:** Configuring a network slice.

Slice modification is implemented on an existing slice, which includes Core UPF within an existing Amazon EKS cluster. The user deploys an additional pipeline using Cloudify UI within the slice, and slice parameters are updated as part of the workflow. For this purpose, we have utilized SD (Slice Differentiator) values which represents a flavor of the existing slice. Updating slice flavor using day two config (SD) will affect QOS/Priority/services within the slice. Two options were evaluated: direct slice modification within UPF day2 config and centralized modification via PCF.

When deleting a network slice as part of lifecycle management, the user could remove all slice resources (Amazon EKS cluster, UPF...) utilizing the uninstall workflow.

The Cloudify Manager uses consistent workflows, with each slice execution task followed by a graphical view based on a Cloudify fulfillment execution graph and detailed logs. Each step of the creation process is followed by a verification task against the real state of the resources. After the orchestration, the user has full visibility using Cloudify UI for all slice resource states (runtime properties). The following **Figure 6** shows the management portal.

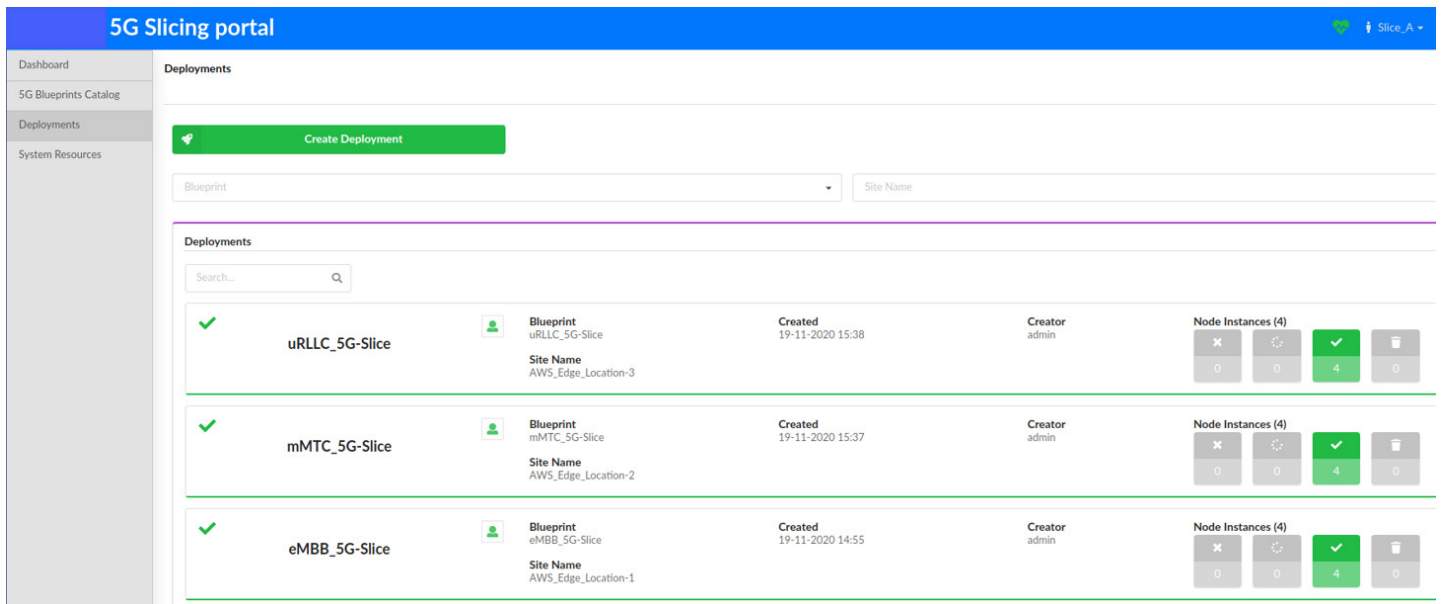


Figure 6: Manage network slices.

## Conclusion

5G network slicing is becoming a critical feature supported by 5G in order to allow communication service providers to support a wide range of applications with different Service Level Requirements on the same network while keeping them isolated and secure.

Cloudify's management and orchestration solution with its full integration to AWS services such as AWS CloudFormation, AWS CodePipeline, Amazon CloudWatch, Amazon EKS – alongside its support to provision workloads and clusters on-premises such as [AWS Outposts](#), or on Amazon EKS anywhere – make it a great solution to support 5G network slicing use cases.

[Read blog online](#)



# Whitepaper:

## 5G Network Evolution with AWS

Design scalable, secure, reliable, and cost-efficient cloud-native core and edge network on AWS

July 2020

### Abstract

5G is transforming the connectivity landscape, allowing lower latency and higher bandwidth across a larger scale of devices. Setting up and managing 5G mobile network functions on AWS allows for global scalability, cost reduction, elasticity, and hundreds of augmenting features (for example, AI, Analytics, IoT, DevOps). This paper highlights the best practices for designing and deploying 5G mobile network functions on AWS.

### Introduction

The advent of 5G brings dramatic improvements to both the radio access network (RAN) and its core network. Besides with the most important and significant change in the radio technology, such as using millimeter wave spectrum (mmWave) for better throughput and less latency in data transmission, network architectures of Core and RAN become to have an easier transition to the modern innovations of cloud-native software technologies such as microservices, containerized, service-based, and stateless architecture.

For the Core Network aspect, 3GPP defines 5G standalone (SA) Core Network to have decomposed architecture with the introduction of a service-based interface (SBI) using HTTP/2 as a baseline communication protocol, and control plane and user plane separation (CUPS). This function decomposition, SBI, and CUPS of 5G network function software strongly favor cloud-native container-based implementation. Even though 5G network function can be built based on the legacy generation of architecture such as virtual machine (VM) based monolithic architecture, the real benefits in terms of agility,

fast innovation, hyper scalability, elasticity, and simplified operations and lifecycle management (LCM) can be realized only when the 5G network functions (NFs) are designed and implemented on a cloud-native container-based microservices architecture.

In the case of RAN, decomposition of Central Unit (CU) and Distributed Unit (DU), the use of enhanced Common Public Radio Interface (eCPRI), and Open RAN (O-RAN) concept make traditional Baseband Unit (BBU) to be more easily transformed to a virtualized network function or containerized network function. This can bring the benefit of scalability and cost-efficiency in terms of both, operating expenditures and capital expenditure.

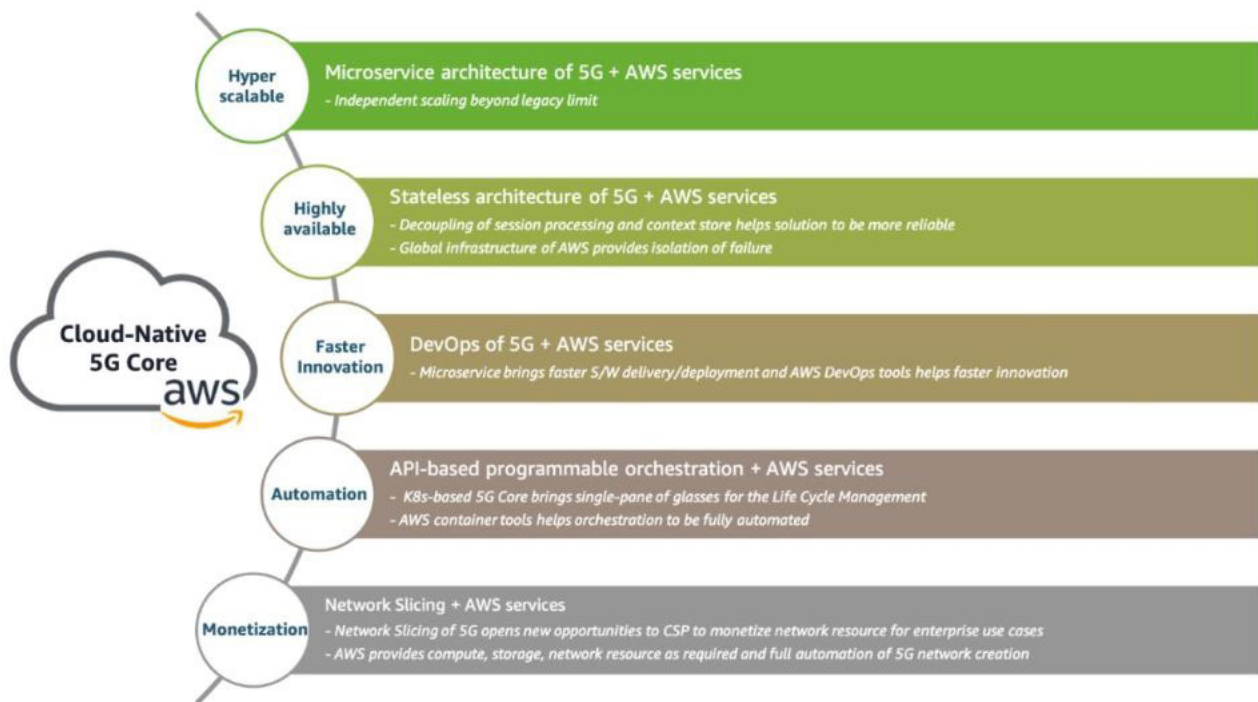
### AWS and 5G

The transformation of both Core and RAN in the 5G era makes AWS an ideal platform for hosting them. This is due to the breadth and depth of AWS services and an API-driven approach to designing modern, cloud-native applications. Additionally, the promise of network slicing is expected to bloom in the 5G era, creating a private and enterprise-oriented network. As such, the 5G network would be best built on the hyper-scalable platform. Therefore, AWS becomes a natural choice for providing 5G network creation, not only for its breadth and depth of services but also because of the strongest and widest partner ecosystem across network equipment providers (NEP) and communication service provider (CSP) in the telecom industry. Openness has been a long-standing requirement of CSPs and AWS Cloud provides exactly that – strong foundation for all NEPs to build and innovate on.

This whitepaper describes a reference architecture of 5G on AWS to help CSPs and NEPs build a carrier-grade 5G production network, in the wide spectrum of views such as design, deployment model, use cases, and AWS tools. It focuses on achieving operational excellence and agility needed by CSPs leveraging AWS products and services for networking, compute, storage, DevOps, and CI/CD pipeline. It also discusses how CSPs can monetize its network by building services using other AWS services like that for data lakes, IoT, and AI/ML.

The paper lays out the typical deployment journey of 5G evolution from using non-standalone (NSA) core, to a new standalone (SA) core network and RAN. Specifically, how relevant AWS services are used for helping cloud-native microservices and stateless architecture-based NFs will be introduced with providing a reference architecture for 5G SA core network that brings the benefit of hyper-scalability, reliability, unified programmable orchestration, and faster speed of innovation for a new service. Lastly, AWS has stretched its service coverage to the Edge cloud by using [AWS Outposts](#), Local Zones, and Wavelength. This whitepaper highlights how these AWS services can contribute to building a seamless 5G Network with providing the best value, the most cost-saving, and the better monetization strategy to the service provider.

[Read full whitepaper online](#)



Cloud-native 5G Core and AWS services

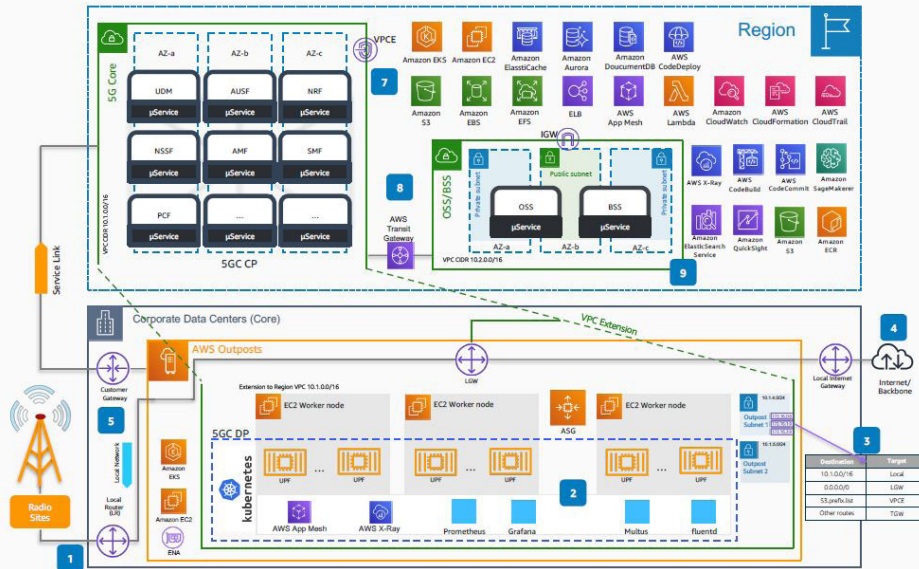
# Reference Architecture: Deploying 5G Core on AWS

This reference architecture explains how 5G Core can be distributed between on-premises data centers and AWS Regions.

## Deploying 5G Core on AWS

### Distribute your 5G Core to on-premises data centers

This reference architecture explains how 5G Core can be distributed between on-premises data centers and AWS Regions.



- 1 Subscriber's traffic from radio access network (RAN) ingress/egress AWS Outposts running 5G User Plane Function (UPF) via Outposts Local Gateway (LGW).
- 2 UPF instances are containers running on Amazon Elastic Kubernetes Service (Amazon EKS) with access to multiple network interfaces via AWS Multi-homing support and Multus.
- 3 Outposts has two subnets (for ingress and egress) with routing tables that contains paths to service end points and Transit Gateway to other virtual private clouds (VPCs).
- 4 Internet access to mobile subscribers is achieved via local gateway as a default route in the subnet route tables.
- 5 Service Link traffic in outpost get separated from local traffic via virtual LANS (VLANs), providing connectivity both locally and to AWS Regions.
- 6 AWS Direct Connect can be used to provide high throughput connection to VPC on an AWS Region (via public virtual interface).
- 7 The service end-point provides direct access to AWS regional services such as Amazon Simple Storage Service (Amazon S3) without traversing via internet.
- 8 AWS Transit Gateway is used to provide connectivity to other VPCs that performs 5G management and control services.
- 9 Orchestration, operational, and business support systems can run on AWS Regions with direct connectivity to on-premises data centers.



© 2020, Amazon Web Services, Inc. or its affiliates. All rights reserved.

AWS Reference Architecture

[View Reference Architecture online](#)

## AWS Architecture Blog

Cloud architecture guidance and best practices

READ & SHARE

## This is My Architecture

A technical video series that showcases unique or innovative cloud architectures

## Case Study:

# Ericsson Delivers Low-Latency Communication with 5G at the Edge Using AWS Outposts

2021

[Ericsson](#) is one of the leading providers of information and communication technology to service providers. As 5G network capabilities entered the market and low-latency communication emerged, Ericsson needed to take those benefits to the edge to deliver high-value use cases for its enterprise customers, and it deployed Amazon Web Services (AWS) solutions to do so. According to Rishi Bhaskar, head of hyperscale cloud business growth at Ericsson, "We're able to take advantage of an innovation that AWS has brought to the market in the form of [AWS Outposts](#)." AWS Outposts is a fully managed service that offers the same AWS infrastructure, AWS services, APIs, and tools to virtually any datacenter, co-location space, or on-premises facility for a truly consistent hybrid experience. Ericsson can leverage AWS Outposts with its 5G Core network to bring low-latency capabilities directly to its customers' enterprise networks.

Learn more in the video below and watch [Ericsson's re:Invent 2020 presentation](#) with Telefónica on unlocking industry transformation using [AWS for the Edge](#) services for critical network workloads.

[View case study online](#)



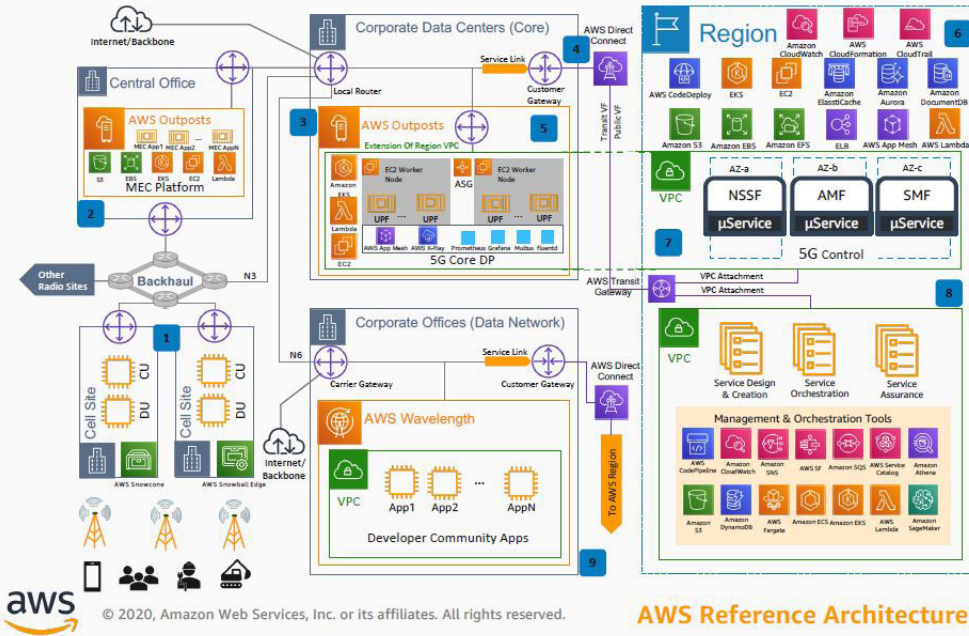
# Reference Architecture: Deploying E2E 5G Network with AWS

This reference architecture explains how different AWS Services can be used together to deliver an end-to-end 5G network.

## Deploying E2E 5G Network with AWS

### 5G RAN, Edge, Core, and Data Network

This reference architecture explains how different AWS Services can be used together to deliver an end-to-end 5G network.



- 1 Based on throughput requirement, AWS Snowcone (up to 100Mbps) or AWS Snowball Edge (up to 10Gbps) can be used for OpenRAN (Distributed Unit (DU) and Centralized Unit (CU)).
- 2 Multi-Access Edge Computing (MEC) capabilities are built using AWS Outposts with Services such as Amazon Elastic Compute Cloud (Amazon EC2), Amazon Elastic Container Service (Amazon ECS), Amazon Elastic Kubernetes Service (Amazon EKS), and Amazon Simple Storage Service (Amazon S3).
- 3 5G Core User Plane Function (UPF) is deployed on AWS Outposts on-premises to provide high throughput.
- 4 Use AWS Direct Connect to connect on-premises 5G Core components to an AWS Region for control and management.
- 5 5G Core user plane function (UPF) is implemented as micro-services on Amazon EKS taking advantage of Single-root input/output virtualization (SR-IOV), Data Plane Development Kit (DPDK), and dual homing capabilities.
- 6 The Control Plane runs on the AWS Region on the same virtual private cloud (VPC) as on-premises. Control plane functions are implemented on Amazon ECS or EKS.
- 7 VPC expansion to on-premises allows UPF instances to expand to AWS regions via Network Load Balancer (NLB) if needed.
- 8 Other VPCs can be interconnected via AWS Transit Gateway and host management and orchestration services.
- 9 AWS Wavelength is used to allow the developer community access to the communication service provider (CSP) environment, and to provide low latency apps to subscribers.

[View reference architecture online](#)

## How to Build This

A video series designed for builders of all skill levels to start building with AWS

## Back to Basics

A video series outlining basic architectural building blocks and best practices

# Blog:

## Open source mobile core network implementation on Amazon Elastic Kubernetes Service

by Christopher Adigun and Young Jung

As introduced in [Amazon Web Services \(AWS\)](#) whitepapers, [Carrier-grade Mobile Packet Core Network on AWS](#) and [5G Network Evolution with AWS](#), implementing 4G Evolved Packet Core (EPC) and 5G Core (5GC) on AWS can bring a significant value and benefit, such as scalability, flexibility, and programmable orchestration, as well as automation of the underlying infrastructure layer.

This blog post focuses on practical implementation steps for creating a 4G core network using the open source project [Open5gs](#).

In addition to showing the benefit of easy installation steps, we introduce how the following AWS services can help the mobile packet network operate efficiently in the cloud environment: [Amazon Elastic Kubernetes Service \(Amazon EKS\)](#), [Amazon Route 53](#) (DNS service), [Amazon DocumentDB](#), [Amazon Elastic Container Registry \(Amazon ECR\)](#), [AWS CloudFormation](#), [Amazon CloudWatch](#), and [AWS Lambda](#).

This generalized example of open source-based 4G core network implementation gives guidance for mobile network function developers and also can be a relevant tool for developers of orchestration, service assurance, and Operation Support System (OSS) solutions that require a general example of mobile packet core network running on AWS.

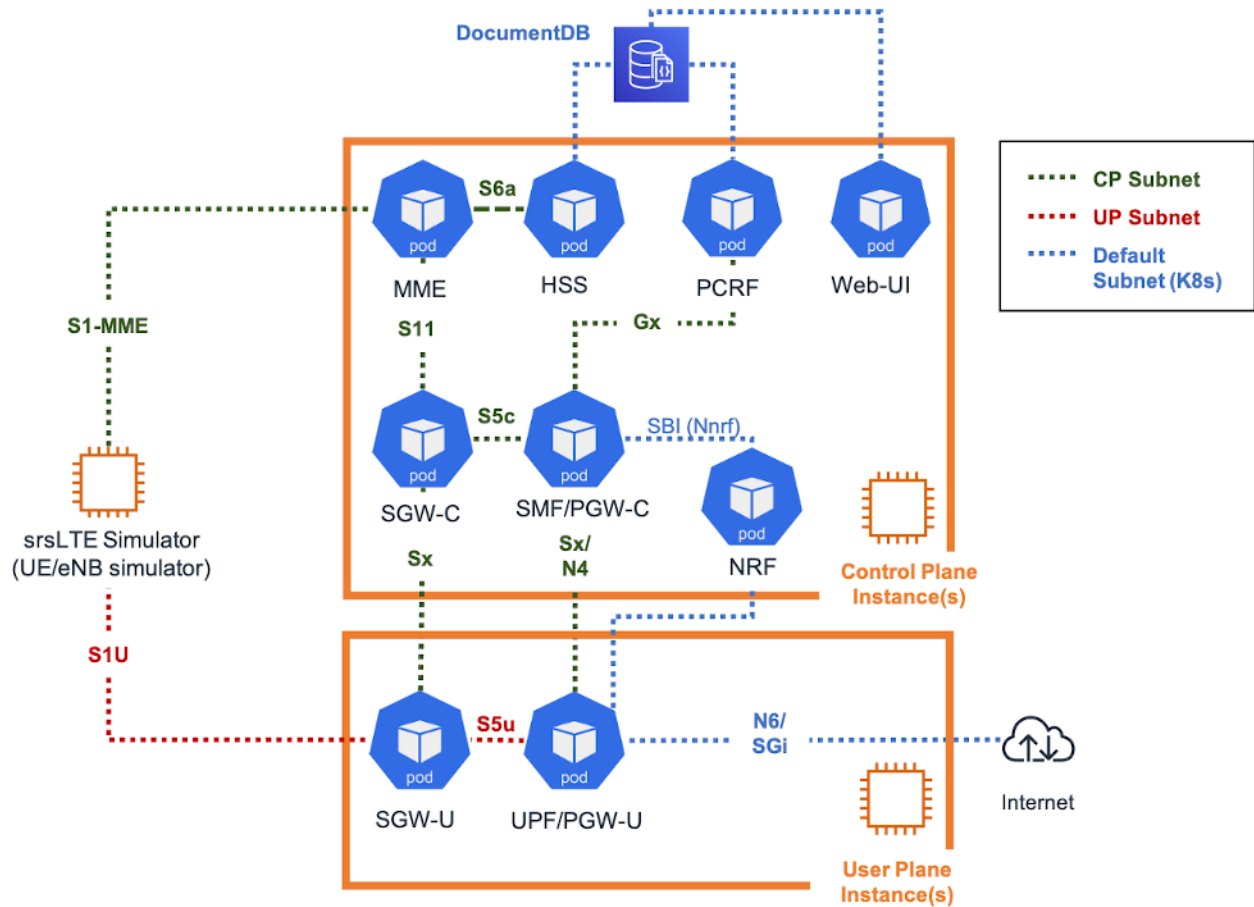
Time to read	About 10-15 minutes
Time to complete	About 45-60 minutes
Cost to complete (estimated)	\$489 (for a month, on-demand instance cost based)
Learning level	Advanced (300)
Services used	AWS CloudFormation, Amazon Elastic Kubernetes Service, Amazon DocumentDB, AWS Lambda, Amazon CloudWatch

## Solution overview

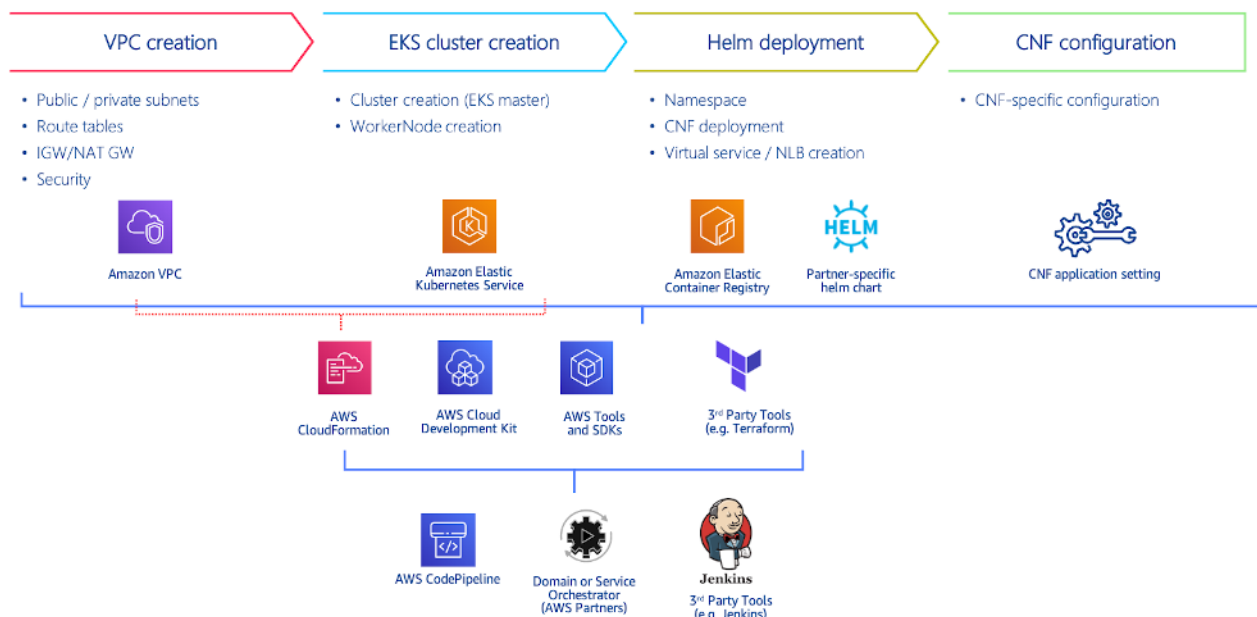
In this implementation, we have chosen Open5gs as a sample mobile packet core application. The Open5gs is an open source project that provides 4G and 5G mobile packet core network functionalities for building a private LTE/5G network under the GNU AGPL 3.0 license. Currently, it supports [3GPP](#) Release 16 with providing 5G Core (AMF, SMF+PGW-c, UPF+PGW-u, PCF, UDR, UDM, AUSF, NRF) network functions and Evolved Packet Core (MME, SGW-c, SGW-u, HSS, and PCRF) network functions.

Among the components in Open5gs, the network function applications in the following table are only used for a 4G EPC network demonstration, with having a 3GPP logical interface in the diagram. Note that even the Network Repository Function (NRF) is a 5G-only network function. It is introduced to use SMF and UPF, which play a role as PGW-c and PGW-u in the Open5gs project.

Network Function	Role
MME	Mobility Management Entity
HSS	Home Subscriber Server
PCRF	Policy and Charging Rules Function
SGW-c	Serving Gateway Control Plane
SGW-u	Serving Gateway User Plane
SMF+PGW-c	Session Management Function + PDN Gateway Control Plane
UPF+PGW-u	User Plane Function + PDN Gateway User Plane
NRF	Network Repository Function (it is only for NF registration of 5G functions)
Web-UI	GUI to configure subscriber and its profile for HSS/PCRF

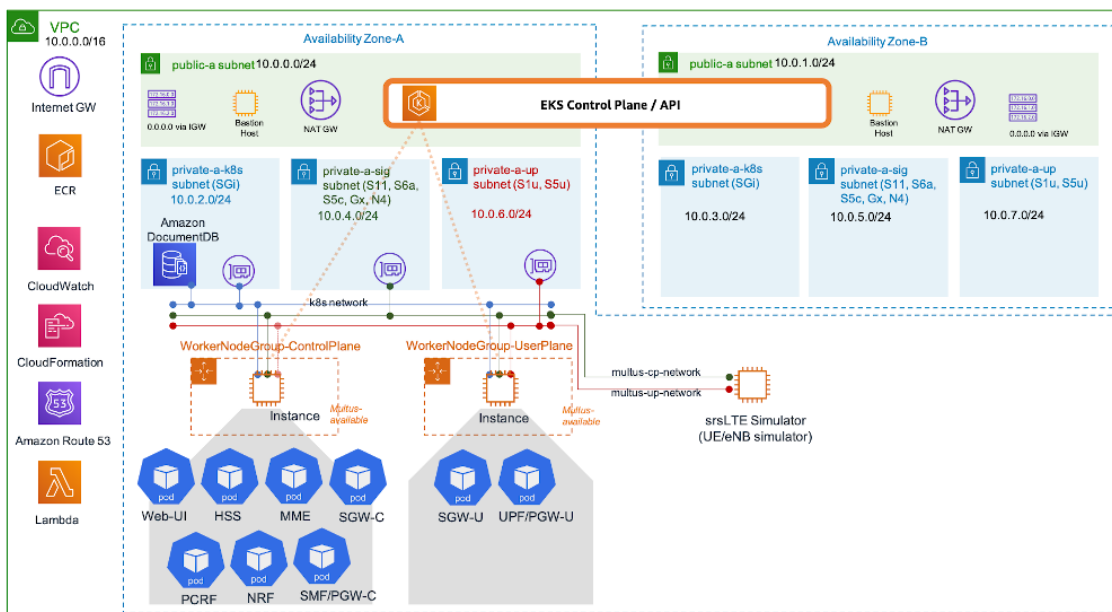


If we use container-based network functions on [Kubernetes](#) (K8s), we can generally standardize a deployment process of these network functions in the flow of VPC creation → EKS Cluster and worker node creation → [Helm](#) deployment → CNF configuration, as in the following diagram, which can be automated with various automation tools and scenario.



In this example, we use a CloudFormation to create an [Amazon Virtual Private Cloud \(VPC\)](#), an Amazon EKS cluster, and two worker node groups (one for the 3GPP control plane, the other for the 3GPP user plane). Importantly, when we deploy these types of open source EPC/5GC on EKS, because they are mostly using multiple network interfaces to serve all different protocols at each interface with having network separation, we have to leverage the [Multus CNI Plugin](#). As guided in [AWS GitHub](#), we can automate this process through AWS Lambda function and Amazon CloudWatch Event Rule. The bottom line is that two AWS CloudFormation templates create the following resources:

- Infrastructure creation template
  - *EpcVpc*: A VPC that will be used for the deployment.
  - *PublicSubnet1/2*: These subnets will host the bastion host for kubectl command run with having public internet access. Also, this will host the NAT-GW to provide internet access for the private subnets.
  - *PrivateSubnetAz1/2*: Subnets for the EKS control-plane in AZ1 and AZ2.
  - *MultusSubnet1Az1*: The first subnet that Multus will use to create secondary interfaces in the EPC control plane pods.
  - *MultusSubnet2Az1*: The second subnet that Multus will use to create secondary interfaces in the EPC user plane pods.
  - *EksCluster*: EKS cluster that will host network functions.
  - *DocumentDBCluster*: For profile store of subscribers, Open5gs originally used MongoDB for HSS and PCRF. In this implementation, Amazon DocumentDB is facilitated because DocumentDB has full compatibility with MongoDB.
  - *Route53 Private Hosted Zones*: For the discovery of service interfaces, such as S6a, Gx, S11, S5-c/u IP addresses, Amazon Route 53 is facilitated as one central DNS.
- EKS worker node group creation template
  - Worker node group for control plane network functions, such as MME, SGW-c, SMF, etc., with additional control plane subnet network.
  - Worker node group for user plane network functions, such as SGW-u and UPF, with additional control plane subnet and user plane subnet networks.
  - Lambda function for attaching additional Multus subnet networks to worker node groups.
  - CloudWatch Event Rule for monitoring instance scaling up and down to trigger Lambda hook to attach additional Multus networks to worker node groups.



Additionally, two more controllers have been developed and introduced for the further steps of automation.

- *DNS update controller:* While we use Amazon Route 53 to resolve the service IP given to the Multus interface, we also created a controller to register this service IP to respective Route 53 private hosted zone automatically. Each EPC service interface uses a separate DNS private hosted zone, created by the open5gs-infra CFN template.
- *Multus IP update controller:* The other controller is used to associate the Multus secondary IPs to the EC2 instance in which the pod is running. The controller listens for pods with designated annotations, and it searches for the secondary IPs and then calls Amazon EC2 API to associate the IP at the Multus interface of the pod to the respective ENI of the host instance. It also disassociates the IP from the host ENI when the POD gets deleted.

After a successful deployment of Open5gs, the functionality of the 4G core network can be tested with other tester or simulators. In this article, we have used [srsLTE simulator](#) as an example, but it can be chosen according to the user's preference.

[View full walkthrough online](#)

## Whitepaper: Continuous Integration and Continuous Delivery for 5G Networks on AWS

### Abstract

This whitepaper introduces continuous integration and continuous delivery (CI/CD) for 5G networks, and how Amazon Web Services (AWS) tools and services can be used to fully automate the deployment and upgrades of 5G network functions. The whitepaper provides detailed description of the different stages of CI/CD for 5G network functions, including network setup, infrastructure deployment, cloud-native network functions deployment, and continuous updates of network functions. It also provides details on integration with open-source and third-party tools for testing, observability, and orchestration.

This whitepaper is aimed at Communication Services Providers (CSPs) as well as Independent Software Vendors (ISVs).

Historically, development, lab and field integration testing, and production deployment of new network nodes or new features in a

cellular network took weeks or even months to ensure the stability of mission and business critical telecommunications (telecom) services. The long cycle of deployment was caused by the monolithic architecture of traditional network nodes, a multi-vendor environment, and many point-to-point interfaces among network entities in the 2G, 3G, and 4G mobile networks.

As introduced in the [5G Network Evolution with AWS](#) whitepaper, 5G mobile networks, as standardized by 3GPP, now support a cloud-native architecture enabled by virtualization and containerization. More specifically, 5G networks introduce and support a new paradigm of the microservice, stateless, and service-based architecture.

This 5G architecture means that different network functions can work as loosely coupled independent services that communicate with each other through well-defined interfaces and APIs. Most importantly, each network function can be updated independently. This architecture

shift in 5G enables CSPs to achieve more agility and operational efficiency by making it easier to roll out updates for network functions more frequently, while maintaining the testing, security requirements, and standards through automation.

Integration and deployment of new features for a CSP generally start when the network function vendor releases a new network function software package, such as a [Docker](#) image in a container-based network function, or a new configuration file, such as a [Helm](#) chart in the [Kubernetes](#) application case. (A Helm chart is a collection of files that describe a related set of Kubernetes resources).

The idea of using the paradigm of CI/CD for 5G network function deployment is gaining traction, but the practical realization of this idea has been a challenge in the telecom industry.

AWS has pioneered the development of new CI/CD tools for software delivery to help a broad spectrum of industries develop and roll out software changes rapidly, while maintaining systems stability and security. These tools include a set of Software Development and Operations (DevOps) services such as [AWS CodeStar](#), [CodeCommit](#), [CodePipeline](#), [CodeBuild](#), and [CodeDeploy](#). AWS also evangelizes the idea of Infrastructure as Code (IaC) using the [AWS Cloud Development Kit \(AWS CDK\)](#), [AWS CloudFormation](#), and API-based 3rd party tools such as [Terraform](#). Using these tools, AWS can store the deployment processes of network function within AWS as a source code, and maintain this IaC source code in the CI/CD pipeline to realize continuous delivery.

This whitepaper describes detailed processes for leveraging AWS IaC and CI/CD tools for the deployment and update of the 5G network function. Additionally, this whitepaper covers integration with third-party tools for testing, observability, and orchestration.

AWS CI/CD tools are not restricted to 5G network functions. They are also employed for automating the deployment of 4G networks, which enables CSPs to rapidly and efficiently deploy and update 4G network functions. Most 4G network functions are Virtual Network Function (VNF) based. AWS CI/CD toolsets like AWS CloudFormation can be used to automate the deployment of 4G VNFs, bringing scale and time efficiency for 4G network deployments.

[Read full whitepaper online](#)



# Videos:

## Reinventing Cloud Native 5G Networks with AWS and DISH

Learn how AWS and DISH are collaborating to build the US Telecom Industry's first fully cloud-native 5G Network. Dave Brown, (VP of Amazon EC2, AWS) and Marc Rouanne (Chief Network Office, Dish Wireless) discuss how DISH is leveraging AWS infrastructure, including AWS Regions, AWS Local Zones, and AWS Outposts to build a more flexible, agile, and cost-effective 5G network that redefines the end user experience.



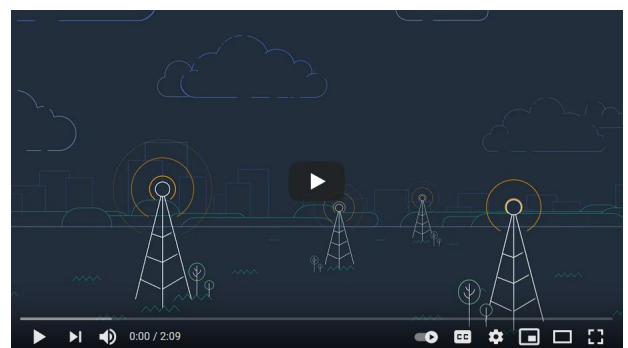
## Connecting the World of Tomorrow, Today, with 5G Edge Cloud: Verizon and AWS

5G is more than just mobile bandwidth for communication services providers (CSPs). By combining 5G with AWS edge cloud services, 5G improves application performance for end users in countless industries, offering up to 10 times the speed of 4G. 5G can power emerging use cases like live video event streaming, smart cars and cities, Internet of Things (IoT), and augmented reality. Learn from Srinivasa Kalpala, VP Technology Strategy, Verizon and Dave Brown, VP EC2 AWS, about how 5G cloud-native private networks are now even easier to deploy with AWS, enabling a complete Industrial 4.0 revolution.



## AWS Wavelength - Edge Computing for 5G Networks

AWS Wavelength brings AWS services to the edge of the 5G network. Wavelength Zones are AWS infrastructure deployments that embed AWS compute and storage services within communications service providers' (CSP) datacenters at the edge of the 5G network, so application traffic can reach application servers running in Wavelength Zones without leaving the telecommunications network. This avoids the latency that would result from application traffic having to traverse multiple hops across the Internet to reach their destination, enabling customers to take full advantage of modern 5G networks. Learn more about AWS Wavelength at <https://amzn.to/33t18oe>

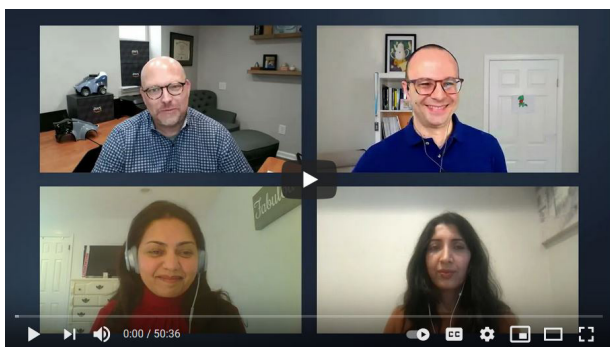


## Accelerating VR Adoption Using 5G Edge Computing - AWS Online Tech Talks



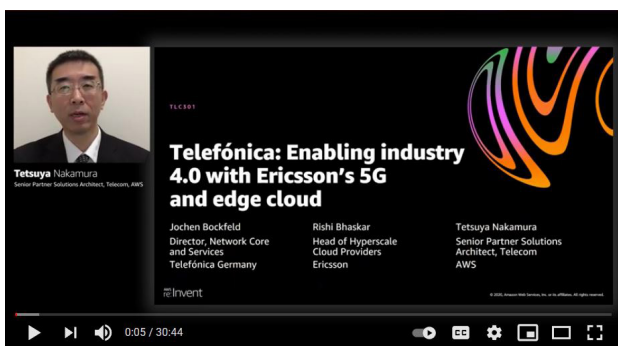
Virtual reality is changing the way people work, learn, collaborate, and enhance their productivity. VR, in its early days, depended heavily on custom bulky hardware mounted with high compute-intensive graphical processors, or tethered with high-powered GPU hardware. AWS Wavelength provides the GPU powered instances that can be located at the edge of the 5G mobile networks, enabling applications to achieve ultra-low latency and high-bandwidth connectivity between AR/VR devices and GPU-based compute on AWS. In this tech talk, we talk about how we can disaggregate the product design from the compute infrastructure using the power of 5G and Edge computing.

## All Things Automotive | S1 Ep2 | C-V2X and 5G MEC with Verizon



C-V2X (Cellular Vehicle to Everything) is helping to make the driving experience safer and more efficient. In this episode, Dean & Stefano will be joined by Verizon to unpack how 5G and Edge computing are playing a key role in driving this innovation.

## AWS re:Invent 2020: Telefónica: Enabling industry 4.0 with Ericsson's 5G and edge cloud



In this session, Telefónica will talk about leveraging the benefits of the AWS edge cloud for critical network workloads, namely the Telefónica 5G Core (5GC) network. Telefónica Germany is deploying Ericsson's cloud-native 5GC for private networks on AWS, and a major German automotive manufacturer is a pilot customer for a proof of concept (involving Telefónica, Ericsson, and AWS) on industrial use cases enabled by standalone 5G and powered by AWS Outposts.

Learn more about re:Invent at <http://bit.ly/3c4NSdY>