



AWS Whitepaper

Next-Generation Mobile Private Networks Powered by AWS



Next-Generation Mobile Private Networks Powered by 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	i
Abstract	1
Introduction	2
Mobile private network components	4
Mobile private network components	4
Network management systems	6
Operations and orchestration	7
AWS edge and IoT capabilities	7
Reference architecture on AWS	9
Enterprise campus network	11
Connected remote site network	11
Disconnected remote site network	12
Hybrid campus and remote site network	13
Full cloud-based core network	14
Network management systems	15
Operations and orchestration	16
Application and use cases	18
Real-time wide-area, high-definition video surveillance	18
Asset management and uptime assurance	19
Autonomous and remote-controlled robotics	20
Remote diagnostics and predictive maintenance	21
Quality sensing	22
Connected worker	23
Physical distancing	24
Connected campuses and remote learning	24
Wayfinding	25
Conclusion	27
Contributors	28
Further reading	29
Document revisions	30
Glossary	31
Notices	33

Next-Generation Mobile Private Networks Powered by AWS

Publication date: **January 27, 2021** ([Document revisions](#))

Abstract

This whitepaper introduces the relevant use cases, solutions, and best practices for designing and deploying mobile private networks powered by AWS. Cloud-enabled mobile private networks allow enterprises, governments, and professional organizations to autonomously deploy their own geo-dispersed, secure mobile private networks. These networks cover private facilities while meeting their performance, reliability, availability, security, and scalability requirements.

Introduction

Large enterprises, businesses, and organizations are increasingly in need of a dedicated reliable mobile private network to address their complex requirements. A mobile private network requires resilience, quality of service, mobility, security, and integration with business applications. The Internet of Things (IoT) and human communication across the facilities are targeted use cases for improving productivity, reducing costs, and improving customer experience.

Commercial mobile network service providers focus on public consumer needs. Enterprises require mobile private network with specific service-level performance. Enterprise-class Wi-Fi lacks seamless voice call mobility and can't provide satisfactory coverage in an environment with metallic structures, buildings, trees, masonry, and other physical structures.

These private networks address gaps with interference and aren't limited to clear and unobstructed transmission paths. Therefore, licensed and unlicensed spectrum mobile private networks are seen as a strong augmentation to the commercial mobile networks, Wi-Fi connectivity, and indoor distributed antenna systems.

AWS has worked with global telecommunications APN Partners to create a set of cost-effective hybrid mobile private network solutions with carrier-grade core Virtual Network Function (VNF), Cloud Native Network Function (CNF), Operations Support System/Business Support System (OSS/BSS), and business applications, hosted both in the AWS Cloud as well as on premises, that satisfy the telecommunications industry requirements. These solutions can be deployed and operated by businesses themselves, third-party network providers, and virtual or traditional mobile network operators.

AWS brings the following unique set of capabilities and services that can enable an end-to-end mobile private network solution:

- **Extending the cloud to the edge** – Private networks on AWS can be seen as an extension of the cloud, where you can make use of the scale and the AWS services on the cloud while still being able to host business and mission-critical applications and network functions on the edge. AWS also brings in the flexible pay-as-you-go charging models of the cloud into private networking, moving away from the traditional Capex intensive telco models.
- **Extending private networking to the cloud** – You can make use of the multiple Regions and Availability Zones offered by AWS. Take applications and network functions that don't need to be hosted at the edge, and move them to a reliable, highly available, and resilient private network extension in the cloud.

- **Security** – AWS is designed to help you build secure and high-performing infrastructure for their applications. AWS security services, such as data protection, AWS Identity and Access Management (IAM), infrastructure protection and threat detection, and continuous monitoring help ensure AWS infrastructure is secure in the cloud as well as at the edge.
- **Automation** – AWS brings the continuous integration and continuous deployment (CI/CD) frameworks of the cloud into private networks, which allows you to deploy applications and network functions in an automated way, as well as easily manage the patching and updating of applications.
- **Edge and IoT** – AWS brings in a wealth of edge and IoT capabilities and products to fit with many IoT and industrial use cases. AWS edge products range from Outposts, for heavy processing on premises, to Snowcone, which is a rugged, portable storage and compute device.
- **APN Partner ecosystem** – AWS actively works with many of the leading independent software vendors (ISVs) and infrastructure vendors to enable their solutions on AWS services. In addition, you have access to the AWS Marketplace where there are thousands of applications from partner ISVs that can be deployed on AWS Region and/or the edge.

In the following sections, we discuss AWS edge capabilities and services. We then explain how to design and deploy these private networks with solutions that meet security, availability, and performance requirements on AWS. Starting with detailing key private networking components, we cover reference architectures, customer applications, and use cases.

Mobile private network components

Conventional cellular mobile network systems, generally marketed as 4G and 5G wireless networks, are composed of the radio access network (RAN) and the Mobile Core networks. These networks have been primarily designed to seamlessly support internet protocol (IP) connectivity between the user equipment (UE) and the Packet Data Network (PDN), with high capacity and data bandwidth support for data and graphics-intensive applications.

Mobile private networks are designed to deliver reliable network coverage across business facilities and operational areas, providing the benefits and functions of a 4G and 5G mobile network in terms of quality of service, security, and reliability.

We start by exploring the AWS edge and IoT portfolio and then detail the mobile private network components and the different architecture options.

Mobile private network components

Unlike a network managed by mobile network operators, a mobile private network is a mobile network dedicated to an enterprise customer with dedicated components deployed on premises. The network utilizes dedicated RAN equipment to serve one or more enterprise campuses with voice and data functionality. This feature provides your business with greater control of the network performance metrics, such as quality of service (QoS), latency, and bandwidth management. This control enables application-aware mobile infrastructure to prioritize the traffic for the business-critical applications.

The following table summarizes key components of mobile private networks powered by AWS.

#	Components	Deployment	Partner/Provider	Required (R)
1	Licensed/unlicensed spectrum	On premises	Regulatory body or CSP	R
2	Spectrum Access Systems (SAS), Domain	AWS Cloud	APN SAS Partners	R (for CBRS)

#	Components	Deployment	Partner/Provider	Required (R)
	Proxy for CBRS Spectrum			
3	4G/5G access points	On premises	APN Partners	R
4	4G/5G Core Networks	AWS Cloud/AWS edge infrastructure	AWS ISV Partners	R
5	AWS edge infrastructure	On premises	AWS Outposts/AWS Snow Family	R
6	Customer premises equipment (CPE)/user equipment (UE)/CBSD	On premises	Third-party providers	R (CBSD for CBRS)
7	SIM cards	On premises	AWS ISV Partners	R
8	Network management system	On premises/AWS Cloud	AWS ISV Partners	R
9	Operations and orchestration	On premises/AWS Cloud	AWS ISV Partners	R
10	ISV or enterprise applications	AWS Cloud/AWS edge infrastructure	AWS ISV Partners or developers	depends on the use case

#	Components	Deployment	Partner/Provider	Required (R)
11	IoT platform	Hybrid	AWS IoT applications and solutions	depends on the use case

Table 1 – Key components of the mobile private network powered by AWS

Note

AWS ISV Partner products are available in [AWS Marketplace](#). For information about third-party provider equipment, see CBRS Alliance Certified Devices in [OnGo Certification Program](#).

Mobile private networks can utilize unlicensed spectrum technologies, such as Citizens Broadband Radio Service (CBRS) available in the United States. Spectrum technologies run wireless radio access points in unlicensed spectrum that are assigned to every citizen for non-exclusive usage subject to regulatory constraints. You can therefore independently deploy mobile private networks on unlicensed spectrum band without obtaining licenses from the telecommunications regulatory authority (for example, the FCC in the US, Ofcom in the UK).

You can also build mobile private networks using licensed spectrum bands, provided that your company owns the required spectrum issued by the telecommunications regulatory authority in your country. You can work with a communication services provider (CSP) to use your licensed spectrum to build the mobile private network.

Network management systems

Network management systems are part of mobile private networks. They handle the observability of the network. Telemetry systems covering the edge network logs and metrics, including RAN and Core in the AWS Cloud, are used to create performance and fault management reports.

Operations and orchestration

Automation and service assurance of mobile private networks are important components of a mobile private network deployment. Network management system outputs are used as inputs to trigger orchestration steps and further enforce service assurance targets. The network management system outputs are also used as inputs for operations like tuning the Radio Access Network (RAN).

AWS edge and IoT capabilities

AWS provides extensive services for you to manage, connect, and deploy your IoT devices and applications in a seamless and secure way. AWS also provides extensive edge capabilities that allow you to have access to compute and storage on premises as well as within a 5G network.

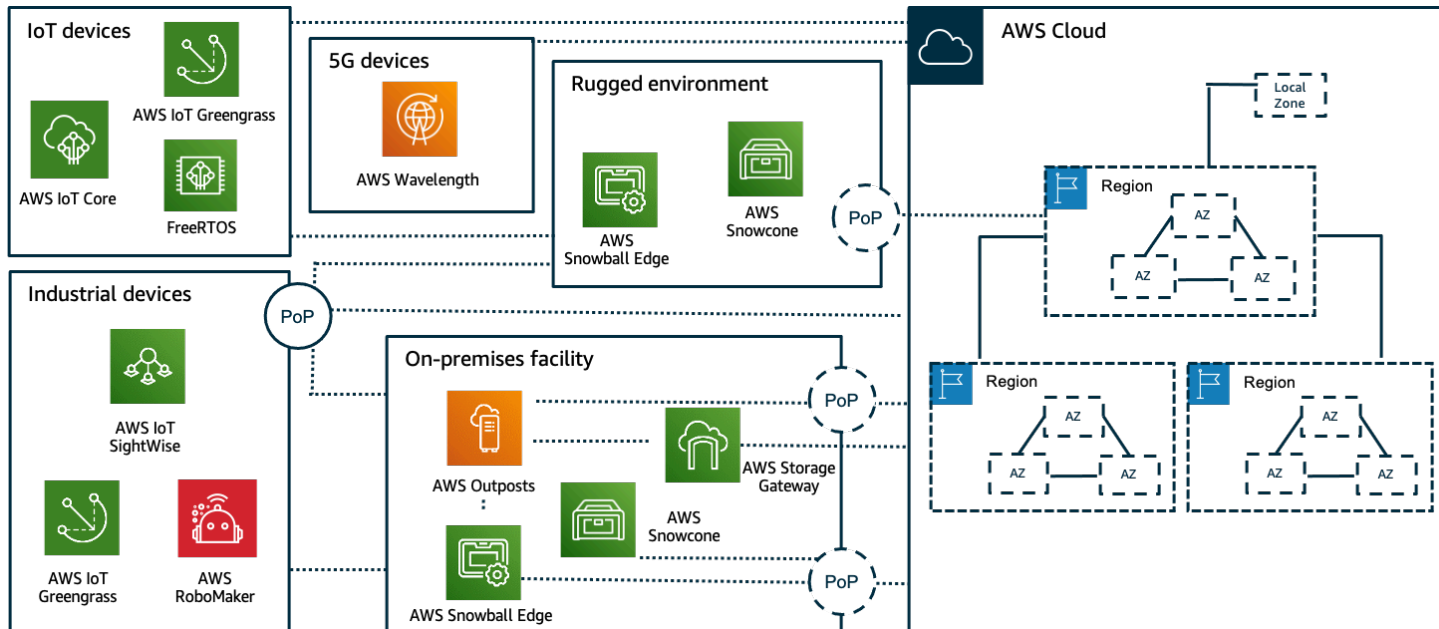


Figure 1 – AWS Edge and IoT capabilities

The following describes the different edge and IoT products and services.

- **IoT and industrial devices** – [AWS IoT Greengrass](#) and [Amazon FreeRTOS](#) help program, manage, and connect on-premises IoT devices in an easy and secure way. AWS IoT Core lets connected devices easily and securely interact with cloud applications. [AWS IoT SiteWise](#) makes it easy to collect, organize, and analyze industrial data at scale. [AWS RoboMaker](#) helps robotics developers simulate, test, and securely deploy robotics applications at scale.

- **AWS Wavelength** – [AWS Wavelength](#) combines the low latency and bandwidth of the 5G network with AWS Cloud services where application traffic can reach application servers running in Wavelength Zones without leaving the mobile network. This prevents the latency that would result from multiple hops to the internet and enables you to take full advantage of the advancements of 5G.
- **On-premises edge devices** – [AWS Outposts](#) provides the same AWS hardware infrastructure, services, APIs, and tools to build and run your applications on premises and in the cloud for a consistent hybrid experience. AWS compute, storage, database, and other services run locally on Outposts. AWS Snowball Edge and AWS Snowcone are portable, rugged edge resources that you can use for storage and compute on premises. These resources work offline as well as in connected mode. Finally, AWS Storage Gateway is a hybrid cloud storage service that gives you on-premises access to virtually unlimited cloud storage.

Reference architecture on AWS

Figure 1 shows the mobile private network infrastructure design with the key components described in [Table 1](#). The components are deployed using various AWS foundational, database, networking, and management services to deliver end-to-end cloud-based enterprise mobile network solutions.

In addition, the AWS IoT services, including [AWS IoT Core](#), [Amazon FreeRTOS](#) and [AWS IoT Greengrass](#), offer a complete suite for connecting, controlling, and managing IoT devices from the cloud or the Edge. This is further enriched by AWS IoT Analytics and machine learning services that help extract value from IoT data and drive decision making.

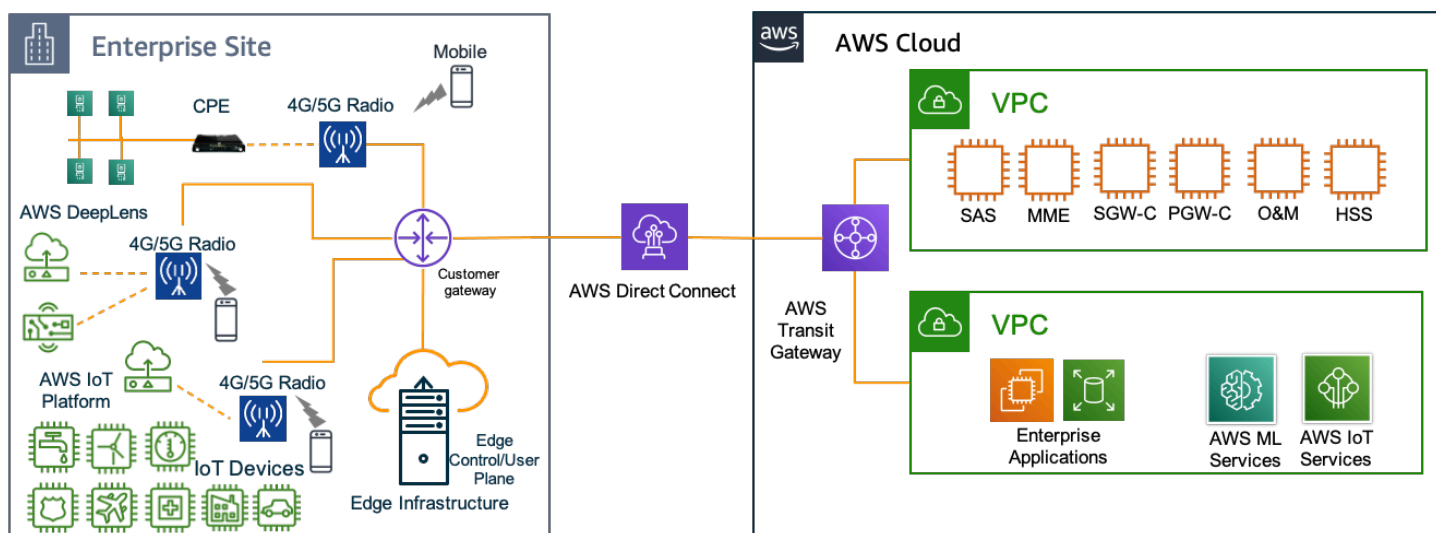


Figure 1 – High-level architecture for application-aware mobile private network

Enterprise sites vary based on business requirements. They can be stadiums, industrial plants, mining sites, university campuses, and other edge locations. Edge infrastructure is deployed on these edge locations to facilitate mobile core networks deployment.

AWS provides edge infrastructure and software that moves data processing and analysis as close as necessary to where data is created in order to deliver intelligent, real-time responsiveness and to streamline the amount of data transferred. This includes deploying AWS managed hardware and software to locations outside AWS Regions, such as [AWS Outposts](#) and the [AWS Snow Family](#).

Outposts is a fully managed service that offers the same AWS infrastructure, AWS services, APIs, and tools to virtually any data center, colocation space, or on-premises facility for a consistent hybrid experience. Outposts is ideal for workloads that require low latency access to on-premises

systems, local data processing, data residency, and migration of applications with local system interdependencies.

The Snow Family helps customers who need to run operations in austere, non-data center environments, and in locations where there is a lack of consistent network connectivity. The Snow Family, comprised of AWS Snowcone, AWS Snowball, and AWS Snowmobile, offers a number of physical devices and capacity points, most with built-in computing capabilities.

Based on edge infrastructure and requirements for connecting to the AWS Cloud, there are five options available to deploy mobile private networks:

- **Enterprise campus network** – In this option, you deploy mobile private networks with AWS Outposts. The networks are connected to the AWS Cloud. This option is ideal for campus networks.
- **Connected remote site network** – In this option, you deploy mobile private networks with AWS Snowcone. The networks are connected to the AWS Cloud. This option is ideal for sensor networks with lower throughput requirements on a lightweight hardware. You can also use AWS Snowball for large-scale connected remote site networks.
- **Disconnected remote site network** – In this option, you deploy mobile private networks with AWS Snowball. The networks are disconnected from the cloud. This option is ideal for sensor networks with lower throughput requirements within remote disconnected sites. You can also use AWS Snowcone for small-scale disconnected remote site network.
- **Hybrid campus and remote site network** – In this option, you deploy mobile private networks with AWS Outposts and AWS Snowcone. These networks are connected to the AWS Cloud. This option is ideal for a hybrid use case for industrial campuses and remote sensor networks with lower latency requirements on a lightweight hardware.
- **Full cloud-based core network** – In this option, you deploy mobile private networks with only Radio Access Network (RAN) installed on premises, with the core network components hosted in your local AWS Region. This option is ideal for applications with higher tolerance of network latency.

The following sections describe each option in more detail.

Enterprise campus network

This option requires connectivity from edge locations to the AWS Cloud. AWS Outposts is the underlying infrastructure for the mobile core user plane and control plane, as well as business applications, such as Machine Learning (ML), Industrial IoT, and communication applications.

The core control plane, subscriber provisioning, policy, and management functions are hosted in the AWS Region inside a virtual private cloud (VPC). In addition, the AWS Region facilitates the observability, control, and management of all AWS resources through different AWS services, such as Amazon CloudWatch and AWS CloudTrail.

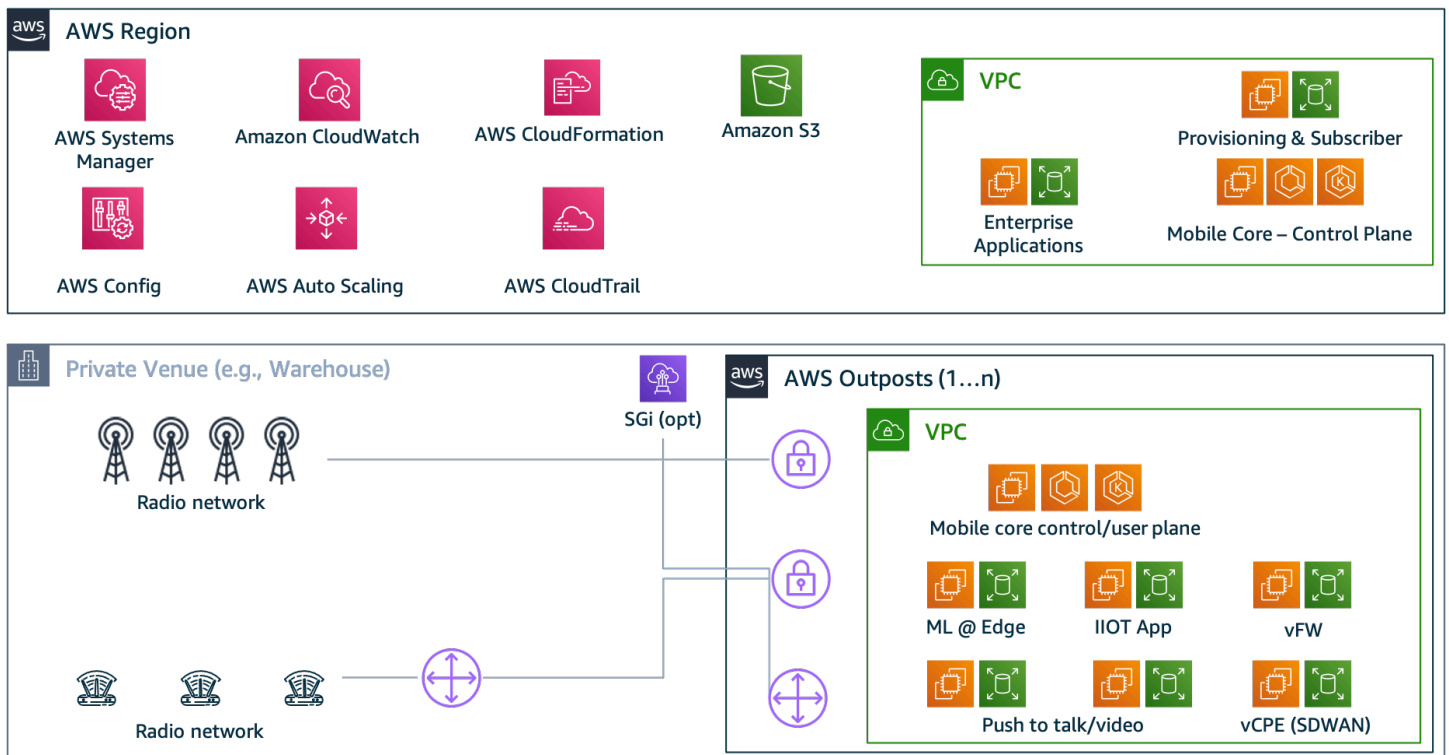


Figure 2 – Mobile private network with AWS Outposts

Connected remote site network

This option requires connectivity from edge locations to the AWS Cloud. AWS Snowcone is the underlying infrastructure for the mobile core user plane.

You can also use AWS Snowball for a large-scale connected remote site network. The AWS Region facilitates control, policy, and management functionality. Subscriber provisioning and mobile core control plane are also part of AWS Region deployment inside a VPC.

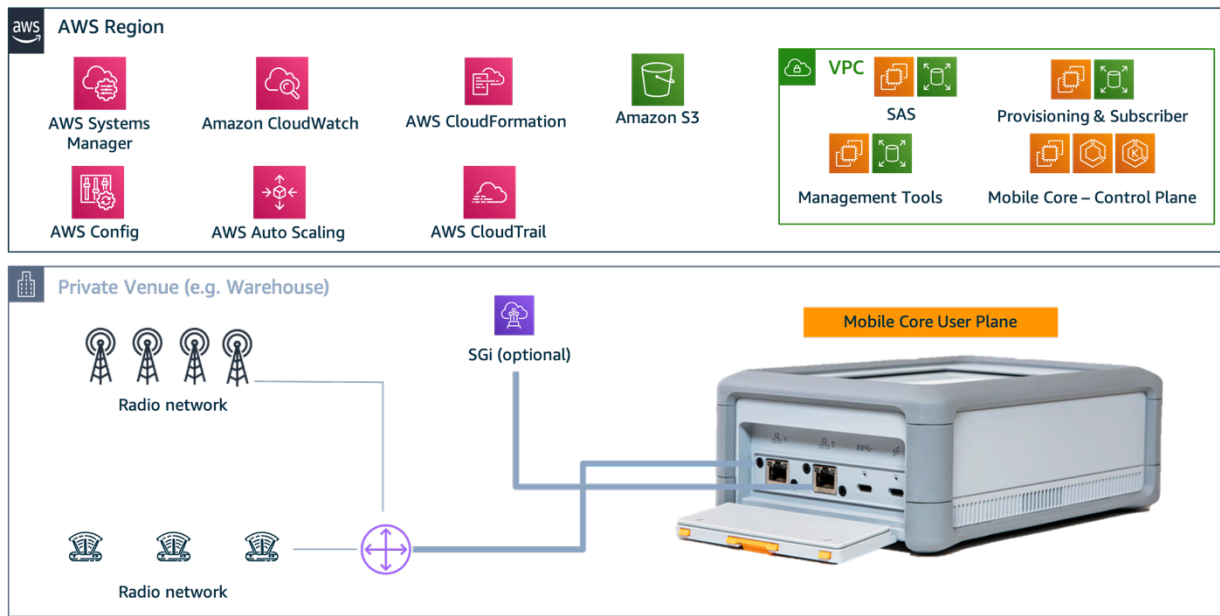


Figure 3 – Mobile private network with AWS Snowcone

Disconnected remote site network

This option does not require connectivity from edge locations to the AWS Cloud. AWS Snowball is the underlying infrastructure for the mobile core user plane and the control plane, as well as business applications, such as ML, Industrial IoT, and communication applications.

You can also use AWS Snowcone for a small-scale disconnected remote site network. The AWS Region can be the optional component to facilitate control, policy, and management functionality. Depending on local regulations, a RAN on certain spectrums, such as CBRS in the US, still require connectivity to the Spectrum Access Systems (SAS).

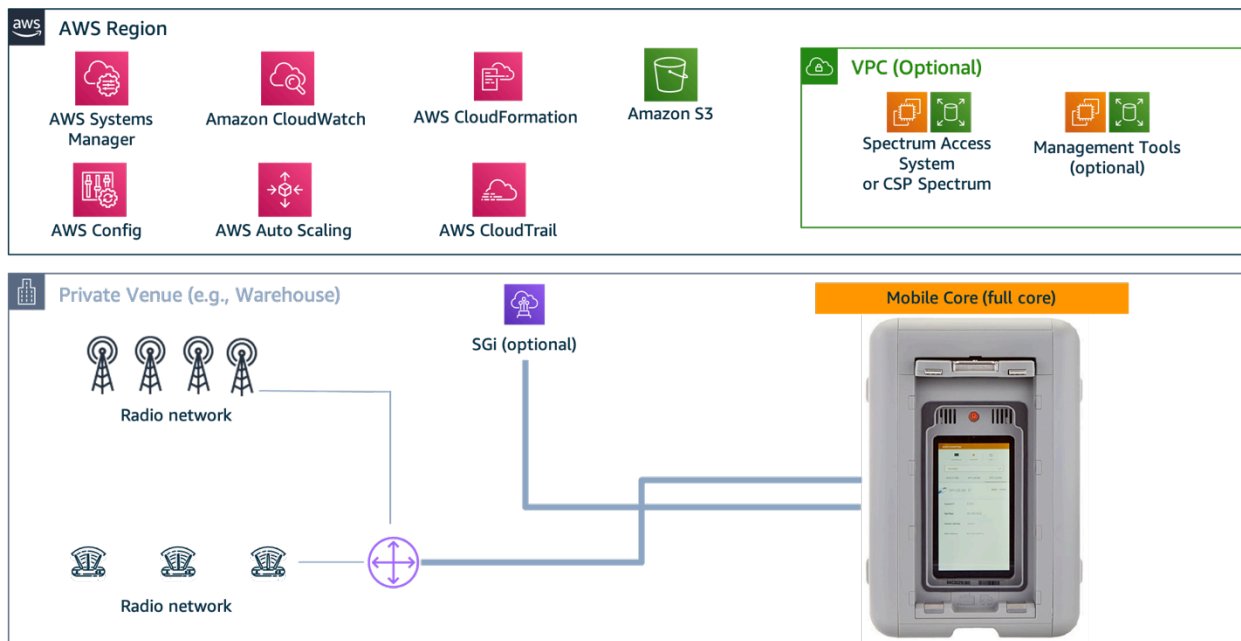


Figure 4 – Mobile private network with AWS Snowball

Hybrid campus and remote site network

This option requires connectivity from edge locations to the AWS Cloud for a combination of industrial campus and remote site coverage. AWS Outposts is the underlying infrastructure for the mobile core user plane and the control plane, as well as business applications located in industrial campus.

AWS Snowcone is the underlying infrastructure for the mobile core user plane, for a remote site requiring lightweight compute for low-latency applications.

The AWS Region facilitates control, policy, and management functionality. Subscriber provisioning and mobile core control plane are also part of the AWS Region deployment inside a VPC.

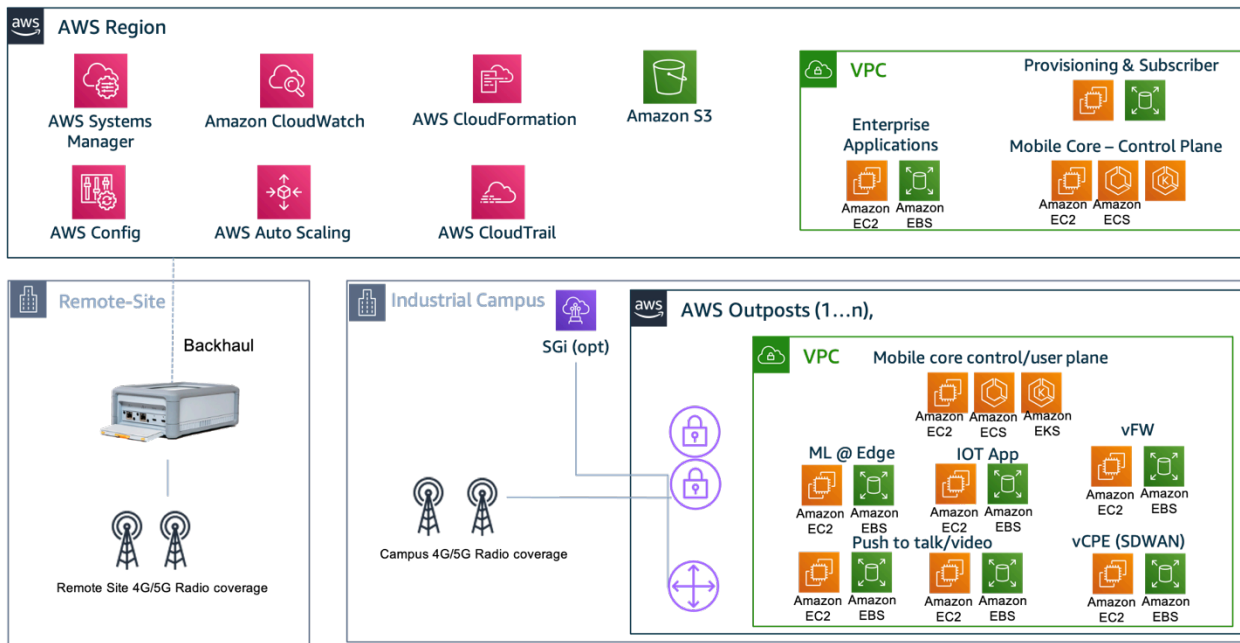


Figure 5 – Mobile private network with AWS Outposts and AWS Snowcone

Full cloud-based core network

This option requires connectivity from on-premises locations to the AWS Cloud. Only a RAN is deployed on premises. The AWS Region is the underlying infrastructure for the mobile core user plane and the control plane, as well as business applications such as ML, Industrial IoT, and communication applications.

The AWS Region facilitates control, policy, and management functionality. Subscriber provisioning and mobile core control plane are also part of AWS Region deployment inside a VPC.

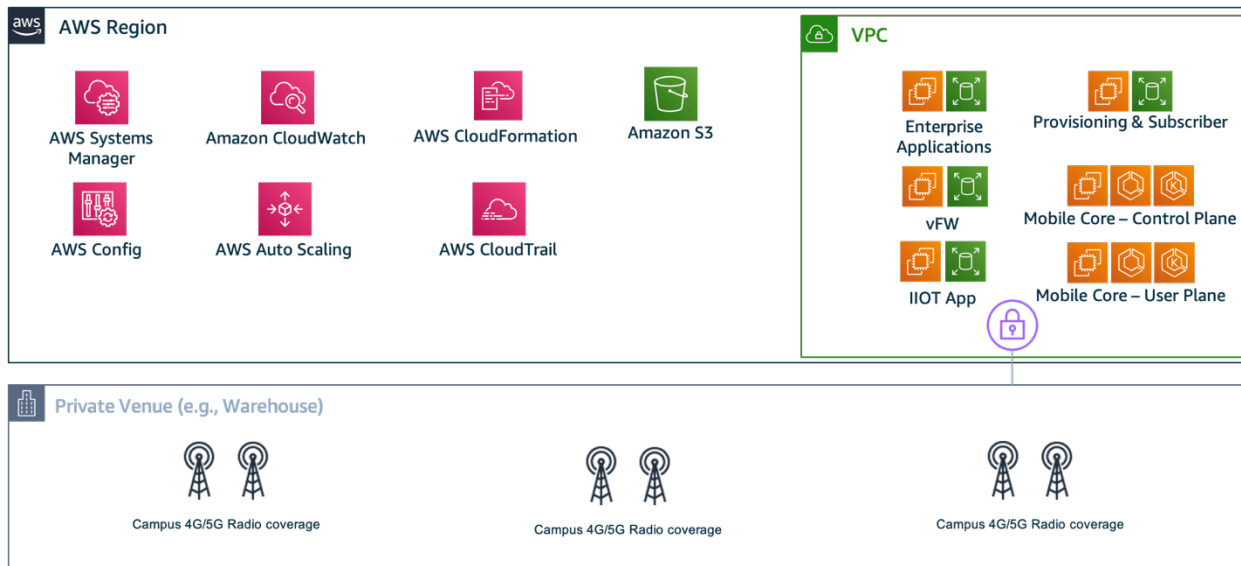


Figure 6 – Mobile private network with full core network in AWS Region

AWS supports the five options of deployment depending on business needs. In the next section, we explain in more detail the network management system (NMS), and the operation and orchestration aspects of the mobile private network architecture.

Network management systems

Like other cellular networks, private networks need an NMS. Typically, NMS includes RAN management, spectrum management, and application telemetry, among other network management functions.

- **RAN management** deals with the management of cell sites, including locking, unlocking, and restarts. It also collects telemetry data on the performance of the RAN. This data can be used to provide observability of the RAN performance and can help operators take corrective actions as needed.
- **Spectrum management** maintains the licenses of the CBRS spectrum, helps allocate frequency spectrums to cell sites, and helps in maintaining the spectrum licenses. Operators have the flexibility of turning off licenses based on usage and enforce efficient use of the spectrum.
- **Application telemetry** collects logs and metrics from the application itself. The data collection varies from application to application. The telemetry data collected from applications provides visibility into the performance of the end-user application. This data also helps operators take actions based on the application telemetry data.

In disconnected mode private networks, NMS is deployed on the Edge Cloud (AWS Snow Family). The NMS in disconnected mode private networks is kept minimal due to the resource form factor. The NMS mainly has RAN Management in disconnected mode private networks. Application telemetry and spectrum management have a mixed mode. In mixed mode, there are modules that collect data while disconnected and then sync up with the NMS Spectrum management and application telemetry hosted in the AWS Region, once the connectivity is established.

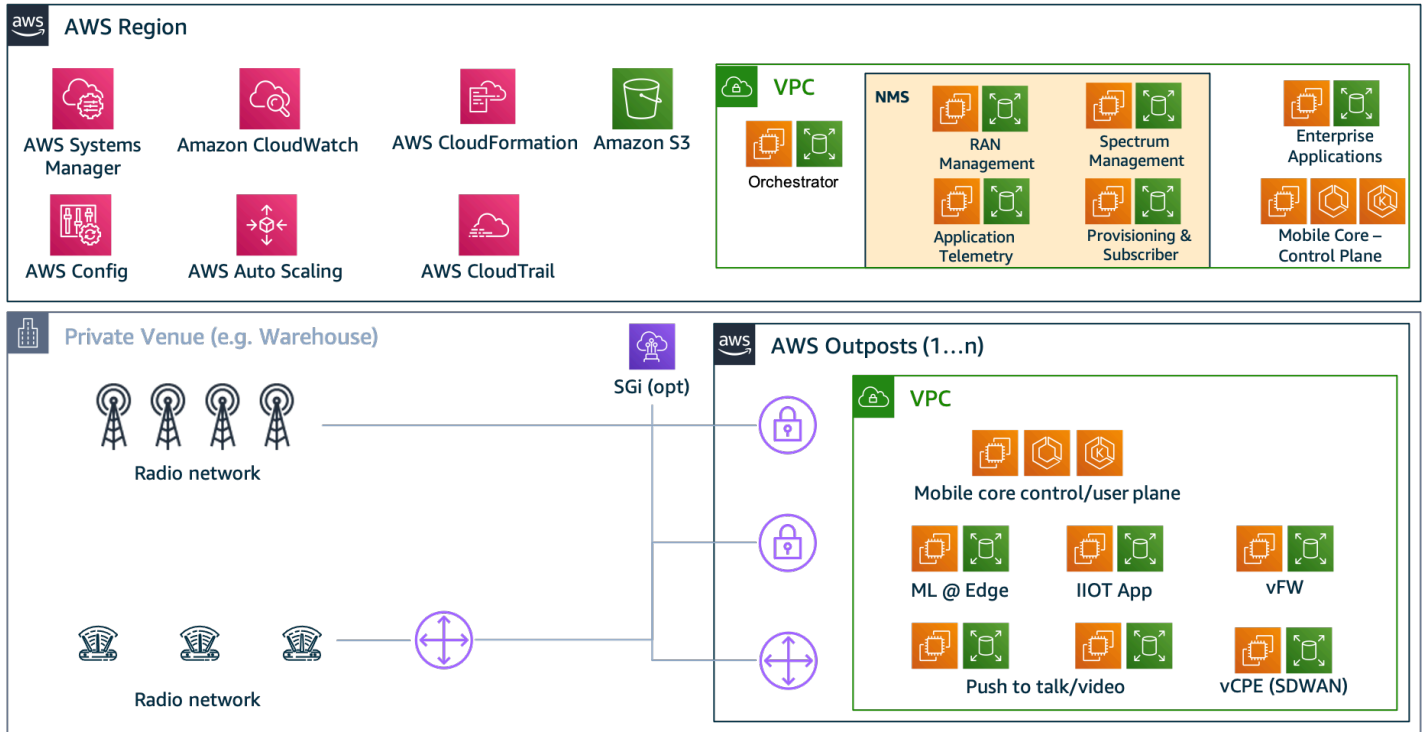


Figure 7 – Mobile private network NMS

Operations and orchestration

Operating a private network is essential for continuous monitoring and management of the lifecycle of physical and virtual resources in the end-to-end network in order to maintain the services delivered to end customers. Examples of basic operations are creating alarms, setting alerts, sending network status notifications, and automatically responding to changes by automated scaling, automated healing, or closed-loop automation.

Network Management Systems (NMS) are often deployed alongside Core and Radio Access Network (RAN) offering E2E management covering monitoring, troubleshooting, configuration, automation, and optimization of the network. You can use NMS to provision, delete, and update subscriber data.

Additionally, some NMS provide features for subscriber mass provisioning. This allows you to provision subscribers in bulk. Orchestration frameworks are often aligned with industry standards, such as ETSI MANO, to provide lifecycle management of services (such as creation, onboarding, scaling, and termination), policy-based automation, and service assurance. ONAP is an example of an open-source project of orchestration framework.

In a regular cellular network, orchestration solutions are often complex and architected to have multiple layers of abstraction to support a wide range of use cases and workloads. A private network, however, requires a much simpler and dedicated orchestration solution that combines various management elements into a few and often integrate with an existing NMS.

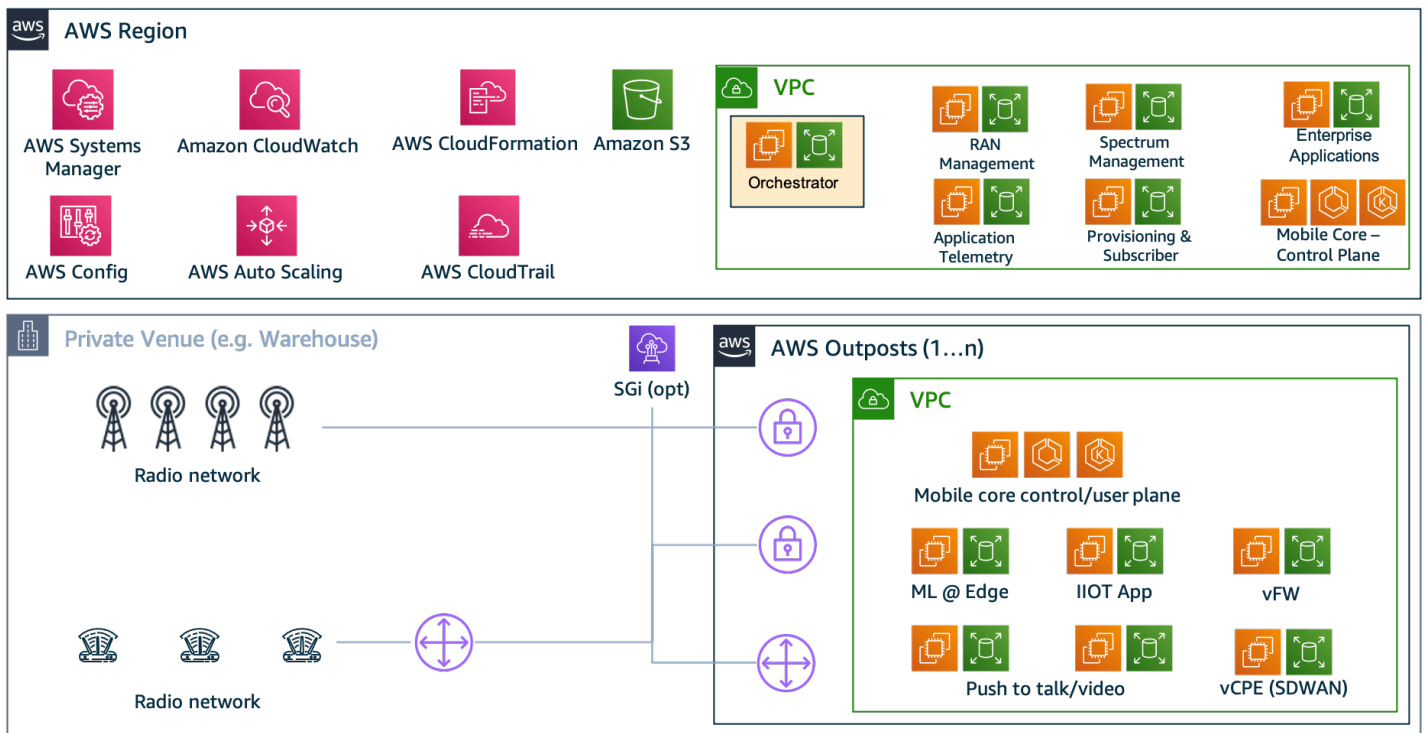


Figure 8 – Mobile private network operations and orchestration

Application and use cases

Mobile network operators and enterprises can implement mobile private networks using APN Partner products and services. The APN Partner capabilities are supported by AWS Cloud-native and AWS edge services, such as AWS Outposts and the AWS Snow Family, to deliver mobile private network use cases.

AWS powered mobile private networks enable your organization to address various business-critical challenges such as:

- Real-time wide-area high-definition video surveillance
- Asset management and uptime assurance
- Autonomous and remote-controlled robotics
- Remote diagnostics and predictive maintenance
- Quality sensing
- Connected worker
- Physical distancing
- Connected campuses and remote learning
- Wayfinding

These use cases evolve with customer needs and will be updated accordingly. The following section discusses each use case in more detail.

Real-time wide-area, high-definition video surveillance

Real-time high-definition (HD) remote video surveillance is typically used by businesses, large organizations, and public sector entities. The intent is to avert crime, vandalism, misuse of the property, and ensure the safety of employees, patrons, and citizens.

Although it can be deployed through fiber optics or a 4G-LTE network, deploying fiber-optic cable only for video surveillance in wider areas is not a cost-effective solution. 4G-LTE wireless communication offers a viable alternative. However, video surveillance requires a commercial 4G-LTE network. These networks usually do not provide the coverage and throughput that high-definition video surveillance in remote and wider areas need. This is where a private LTE comes into place.

In addition, if cameras and video recorders are based on an IP address, the video surveillance IP systems can always connect to the internet. This fact opens a lot of possibilities for video surveillance in wide and remote areas, such as cloud-based AI/ML and analytics services.

AWS DeepLens is the world's first deep-learning enabled video camera. It is optimized to run machine learning models and perform inference on the device. For example, the footage collected by AWS DeepLens cameras can be streamed using Amazon Kinesis Video Streams and sent to Amazon SageMaker to train a data model and deploy it to AWS DeepLens. In addition, Amazon Kinesis Video Streams can be sent to Amazon Rekognition for advance analytics, such as biometric facial recognition, or unusual behavior examination.

A private LTE network provides a robust backhaul network to ensure 24/7 connectivity of AWS DeepLens cameras to the AWS Cloud infrastructure for uninterrupted streaming of large amounts of HD video data and analytics. With AWS Outposts, network operators and enterprises can create a hybrid cloud infrastructure with private LTE-enhanced packet core (EPC) components hosted on Outposts placed on premises in a private data center. Or it can be placed in a control and user plane separated manner with the EPC user plane function placed on Outposts and control plane functions hosted in the AWS Cloud. Outposts extends AWS Cloud capabilities to an on-premises facility. This provides the same AWS native services, infrastructure, APIs, deployment tools, and control plane that customers use in AWS to support LTE virtualized network functions (VNFs) and other telecommunications and IT workloads that need to remain on premises to achieve low latency response times or process data locally.

Asset management and uptime assurance

In large factories, oil wells, hospitals, and airports with multiple devices and systems, system uptime is one of the key performance indicators used to reduce waste and minimize losses due to maintenance issues. Similarly, in the large corporations, tracking, tracing, and management of thousands of IT and non-IT assets is critical to prevent theft and maintain inventory for warranty, licenses, and auditing purposes.

AWS IoT services enable you to bring machines, cloud computing, analytics, and people together to improve the performance and productivity of industrial processes. AWS powered private LTE networks provide a low latency, high capacity, and reliable wireless backhaul owned and managed by businesses to interconnect millions of IoT sensors with on-premises edge computing services, such as AWS Outposts, in conjunction with the AWS Cloud Region to trace, track, manage, and operate the devices, systems, and other assets. Uptime assurance requires real-time decisions that need to be made at the device level, requiring more intelligence at the network edge. AWS edge

cloud services provide infrastructure and software that move data processing and analysis close to the endpoint to deliver intelligent, real-time responsiveness, and reduce the amount of data transferred.

Smart devices analyze information on site and send only relevant information back to the data center. This reduces network demands, not to mention the analyzing and storing of all that data. So, you gain all the benefits of AWS Cloud capabilities, such as Amazon S3 data lake solution for the events ingestion, AWS IoT Analytics services, and IoT device management and control services, together with AI and ML capabilities.

Private LTE also provides built-in controls for the IoT applications that public LTE networks can't offer.

Autonomous and remote-controlled robotics

Today's robotics technology has evolved towards autonomous capability. This capability helps with critical use cases, such as remote mining site exploration, remote area search and rescue operations, and public safety projects.

In most cases, the robot has to provide a communication channel within remotely disconnected areas that have no internet connectivity. The communication flow happens between team members during remote area search and rescue operations, while the robot walks together with the personnel providing mobile private network (MPN) connectivity. Another communication flow possibility is control communication between the robot operations staff with remote exploration robots inside mining tunnel.

AWS Snowcone can serve as an edge computing infrastructure by providing core MPN functionality. AWS Snowcone weighs 4.5 lbs. (2.1 kg). It is 8.94 inches long, 5.85 inches wide, and 3.25 inches tall (227 mm x 148.6 mm x 82.65 mm). The weight and size make it suitable for a robotics payload during remote disconnected site exploration.

A robotics use case in a connected metro area can be combined with an AWS Wavelength zone to provide a low-latency video analytics solution. The robots can be equipped with IP cameras that deliver video stream to a video analytics application. AWS Wavelength can be used to deploy a real-time video analytics application to deliver analytics results to the robot's operator.

APN Partners, together with the AWS Professional Services team, provide software, hardware, and services to implement MPN for robotics use cases.

Remote diagnostics and predictive maintenance

The next industrial revolution is underway with a shift to selling products as services, predictive and preventive maintenance, remote diagnostics in large and wide-area industrial complexes, increased automation, and factories reconfigured with full mobility.

With wireless connected Industrial IoT (IIoT) sensors, LTE wireless mobile technology provides a robust dedicated wireless backhaul network for device connectivity and data transmission. As noted earlier, commercial LTE options do not fulfill IIoT service level requirements and industries have no control on the network. But private LTE networks are owned and managed by the industries and therefore can be tailored for industrial applications with greater control of the network performance, such as quality of service (QoS), latency and bandwidth management, cellular-based security, industry grade reliability, and ultra-low latency.

A private LTE network provides capacity for a large number of wireless connected IoT sensors with superior coverage, both indoors and outdoors, and seamless mobility with service continuity for the mobile industrial equipment. Also, the traffic between the machines and sensors stays local.

A key requirement for predictive maintenance is a collection of large amounts of data for each component of the industrial equipment. Prediction, also referred to as inference, requires ingesting data, building and training data models before those models can be used for failure predictions and preventive maintenance.

AWS Internet of Things (IoT) service provides broad and deep IoT and Industrial IoT services from the edge to the cloud, including cloud computing, machine learning (ML) and analytics capabilities. AWS IoT provides services required for data collection from industrial equipment through IoT sensors in near-real-time (sub-second), hourly, or daily, depending on your business requirements, connectivity, and budget.

Sensor data is collected and transmitted from the equipment to the AWS Cloud for real-time monitoring and analytics, as well as to build, train, and evaluate ML models that are used for maintenance predictions. AWS powered private LTE networks provide ultra-low latency, high bandwidth, and reliable connectivity between the IIoT sensors and AWS Outposts edge cloud, where AWS IoT Greengrass core and virtualized enhanced mobile packet core (vEPC) instances are hosted.

Quality sensing

As the world moves at a rapid pace, industrial output and production are being continuously upgraded to deliver in mass and at shorter time cycles. Due to time restrictions, industrial output can escape quality. There's also a shorter time for manual error detection.

This issue can be solved with a quality sensing solution on a private network. This is a proactive solution that can help reduce quality errors and can enhance a traditional reactive quality monitoring.

We use AI/ML models to learn from machines and collect quality data. Using this data, we can potentially predict quality error before it occurs. This prediction can reduce cost and time of production. The predictive nature of the solution helps prevent quality issues that pass from the production floor, thereby reducing the production line downtime and cost.

With a private wireless network at the edge, the quality-sensing application hosted on a private wireless network brings in latency and computing resources to the edge. LTE wireless mobile technology provides a robust, dedicated wireless backhaul network for the device connectivity and data transmission. Commercial LTE options do not fulfill service level requirements and industries have no control on the network. Private LTE networks are operated and managed by the industries. This allows for industrial applications with greater control on the network performance, such as quality of service (QoS), latency and bandwidth management, cellular-based security, industry grade reliability, and ultra-low latency.

The quality sensing application is used for predictive maintenance, which helps industries and manufacturing units take proactive actions. As part of the solution, industrial sensors are deployed in factory, and the data is monitored at each individual component level. The sensors track the individual machine performance for different environment metrics, such as temperature, pressure, speed, and power. In addition, operational data is collected from various assets on the factory floor. This data is ingested and algorithms are used to detect differences between ideal production specifications and actual production data. The analysis is used to create current and historical notifications for process quality in real time. Secured API operations provide front-end applications or qualified third-party systems access to current, historical, raw, or processed data for additional analysis and design updates.

The following is a summary of the quality sensing application:

- Monitors production setpoints, and analyzes variance data to help predict conditions that might result in quality issues, and alert floor technicians

- Feeds advanced AI/ML models to help improve machine performance predictions for quality
- Creates powerful data visualizations for live and historical quality factors and performance
- Enables reporting functionalities, pulling from high-performance data architecture

AWS Internet of Things (IoT) provides broad and deep IoT and IIoT services from the edge. Sensor data is collected and transmitted from the equipment to the AWS edge cloud for real-time monitoring and analytics, as well as to build, train, and evaluate ML models that are used for quality sensing. AWS powered private LTE network provides ultra-low latency, high bandwidth, and reliable connectivity between the industrial IoT sensors and AWS Outposts multi-access edge computing (MEC) where the quality sensor application and virtualized LTE packet core (vEPC) are hosted.

Connected worker

The health and safety of industrial workers is a priority for enterprises. This is especially so in vulnerable or hazardous working environments, including ports, mines, oil and gas fields, processing plants, construction sites, and manufacturing floors and warehouses. Within these environments, workers are in close proximity to heavy machinery, moving objects, and working at height. Some workers are exposed to extreme conditions, such as heat, high noise, and dangerous substances. The safety gear that workers wear, such as high-visibility vests and helmets, are vital to keeping the workers safe. However, today's technology enables us to provide the workers with capabilities to keeping them connected to further ensure their safety.

Private LTE connectivity ensures that workers are always connected, even in harsh, remote, isolated environments. Private LTE solutions provide the necessary coverage, reliability, latency, and quality of service to enable services such as push-to-talk (PTT) and push-to-video (PTV). These solutions allow workers to communicate with each other and with their command office in a reliable way. Also, body-worn cameras and sensors that keep track of workers biometrics, including heart rate and fatigue level as well as noise level sensors, help give an all-around picture of the worker's condition while on site. These cameras and sensors also help workers react quickly, prevent worker-down situations, and warn workers of hazards. High accuracy geo-location services help trace people and form no-go zones and ring-fencing areas where workers are not supposed to be or areas that are too dangerous for them.

Finally, providing workers with wireless devices connected to machines and equipment in the field for field maintenance and using VR/AR services in field maintenance jobs improves workers efficiency.

AWS edge services, including Snowcone and Snowball Edge, can enable you to deploy an isolated private network in a harsh environment. This ability allows for the high computation capacity and low latency needed for the connected worker use cases. Also, AWS IoT services enable you to connect the IoT sensors and devices worn by workers where the data can be kept on site, shared between users, and transmitted to the cloud for processing and analytics.

Physical distancing

Social distancing has become an important aspect of our day-to-day life, where we need to adhere to a certain level of physical distancing from other people. Keeping track of whether people are adhering to physical distancing rules is a tricky task, especially in public venues such as shopping malls, transportation hubs, busy squares, as well as classrooms in universities. This requires sophisticated real-time video analytics capabilities, which are computation intensive, requiring both low latency communications and a high throughput.

Computation capabilities have evolved from CPU-based to GPU-based and lately into field-programmable gate array (FPGA)-based to support advanced real-time analytics and streaming requirements. FPGA-based accelerators can deliver the required performance but pose challenges related to hardware integration, device programming, management, and application integration. Recent advancements have addressed some of these barriers, enabling FPGA acceleration of real-time analytics for higher processing efficiency and lower latency.

AWS is working with its partners to provide a comprehensive vertical stack to abstract the complexity of FPGA programming and management and develop a physical distancing solution on [Amazon EC2 F1](#) instances utilizing FPGAs. This, coupled with AWS edge capabilities, allows you to deploy a full private network on premises, leveraging the low latency and high processing capacity to enable a high computation-intensive application, such as a video analytics solution for physical distancing.

Connected campuses and remote learning

A fast, reliable, and secure wireless connectivity solution has become a crucial need for university campuses. Many campuses face increasing connectivity challenges to effectively serve their students, staff, and faculty. This is especially the case during the global COVID-19 pandemic, where more students must learn remotely.

Schools and universities must quickly adapt to e-learning and support technologies such as augmented reality (AR) and virtual reality (VR). They must provide collaboration tools for students

to work together seamlessly, whether on or off campus. These requirements have prompted university administrations to explore setting up high-speed campus-wide wireless services that can help students connect effectively on premises while also having access to all the university facilities when off premises.

An LTE-based private cellular network on premises (for example, using CBRS spectrum) offers capabilities such as seamless coverage across campus, high capacity, low latency, and security. The private network provides effective inter-building connectivity. This allows students access to learning resources, whether they are in the library, lab, café, or anywhere on campus. The low latency and high throughput network also enables better smart classrooms where students and teachers can use technologies for interactive learning experiences, such as AR, VR, HD video streaming, and smart boards.

The private network also helps provide effective access to students off campus by removing the connectivity bottleneck on campus. This ensures that all services can be provided over the public internet securely and effectively.

Finally, the private network enables more efficient facilities and building management by providing connectivity to staff (push-to-talk) and connecting the sensors and on-campus IT systems to the centralized network private core.

The different AWS edge solutions allow campuses to deploy on-premises private networks in different sizes, starting with Snowcone for a small network, all the way to Outposts for a private network at scale. Also, AWS is working with partners to provide innovative access solutions based on the CBRS spectrum as well as other licensed solutions to provide an end-to-end private network.

Wayfinding

Navigating large venues, such as shopping malls, exhibitions, transportation hubs, and stadiums can be confusing. Such navigation can be an exhausting experience, especially when you have a limited time. However, new technologies, such as augmented reality and mixed reality, have enabled venues to enrich their customers' experience. Customers can use their phones to view their surrounding in an augmented fashion, such as identifying points of interest, navigating to points of interest, signposting different locations with helpful information, in addition to showing virtual avatars of assistants that can help customers have a better experience. Gamification can be introduced, whereby users can be engaged in a game where they try to find certain artifacts that would further improve their experience.

All these use cases require a communication platform with high capacity, low latency, and reliability for them to work properly and at scale. A private network with edge compute capability on premises satisfies these requirements because all computation is performed on site, allowing for a consistent experience for users.

Conclusion

Private 4G and 5G mobile networks enable businesses to utilize next-generation mobile network technologies for a range of customer use cases. Mobile private networks are designed to deliver wireless network services across enterprise campuses and remote sites with high throughput and low latency connectivity while allowing businesses to autonomously deploy and manage their own wireless networks.

The APN Partner ecosystem provides key components required to build mobile private networks. AWS Regions and edge services, such as AWS Outposts and AWS Snow Family, support carrier-grade core virtual network functions, cloud-native network functions, operations support systems/business support systems, and business applications for carrier-grade enterprise mobile private network deployment.

Contributors

Contributors to this document include:

- **Shonil Kulkarni**, Manager, GSI Solutions Architecture, Amazon Web Services
- **Sigit Priyanggoro**, Sr Partner SA, Dedicated Edge, Amazon Web Services
- **Hisham Elshaer**, Sr Consultant, AWS Telecom, Amazon Web Services
- **Rabi Abdel**, Principal Consultant, AWS Telecom, Amazon Web Services
- **Vara Prasad Talari**, Principal Consultant, AWS Telecom, Amazon Web Services
- **Robin Harwani**, Head of Global Telecom Partner Technology, Amazon Web Services
- **Tipu Qureshi**, Principal Engineer, AWS Premium Support, Amazon Web Services

Further reading

For additional information, see:

- [Amazon EC2 Overview and Networking Introduction for Telecom Companies](#)
- [Carrier-Grade Mobile Packet Core Network on AWS](#)
- [5G Network Evolution with AWS](#)

Document revisions

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

Change	Description	Date
Minor Updates	Updated sections for technical accuracy	January 27, 2021
Initial publication	Whitepaper first published	December 15, 2020

Glossary

- **AF** – Application Function
- **AMF** – Access & Mobility Management Function
- **AUSF** – Authentication Server Function
- **CBRS** – Citizen Broadband Radio Service
- **CBSD** – Citizen Broadband Radio Service Device
- **CHF** – Charging Function
- **CNF** – Cloud-native or Containerized Network Function
- **CSP** – Communication Service Provider
- **CU** – RAN Central Unit
- **CU-CP** – CU Control Plane
- **CU-UP** – CU User Plane
- **CUPS** – Control and User Plane Separation
- **DN** – Data Network
- **DP** – Domain Proxy
- **DU** – RAN Distributed Unit
- **EPC** – Evolved Packet Core
- **FCC** – Federal Communications Commission
- **ISV** – Independent Software Vendor
- **MANO** – Management and Orchestration
- **MEC** – Multi-Access Edge Computing
- **MPN** – Private Mobile Network
- **NEF** – Network Exposure Function
- **NFV** – Network Function Virtualization
- **NFVI** – Network Function Virtualization Infrastructure
- **NFVO** – Network Function Virtualization Orchestrator
- **NRF** – Network Repository Function
- **NSA** – Non-Standalone 5G
- **NSSF** – Network Slice Selection Function

- **Ofcom** – Office of Communications, UK
- **PCF** – Policy Control Function
- **PoP** – Point of Presence
- **RAN** – Radio Access Network
- **RU** – RAN Radio Unit
- **SA** – Standalone 5G
- **SAS** – Spectrum Access Systems
- **SBI** – Service-Based Interface
- **SCTP** – Stream Control Transport Protocol
- **SMF** – Session Management Function
- **UDM** – Unified Data Management
- **UE** – User Equipment
- **UPF** – User Plane Function
- **vEPC** – Virtual Evolved Packet Core
- **VIM** – Virtualized Infrastructure Manager
- **VNF** – Virtual Network Function
- **VNFM** – Virtual Network Function Manager

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.