

Overview of the Samsung Push to Talk (PTT) Solution on AWS

October 2017



© 2017, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Notices

This document is provided for informational purposes only. It represents AWS's current product offerings and practices as of the date of issue of this document, which are subject to change without notice. Customers are responsible for making their own independent assessment of the information in this document and any use of AWS's products or services, each of which is provided "as is" without warranty of any kind, whether express or implied. This document does not create any warranties, representations, contractual commitments, conditions or assurances from AWS, its affiliates, suppliers or licensors. 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.

Contents

Introduction	1
AWS Overview	1
AWS Infrastructure and Services for Samsung PTT Solution	2
Regions and Availability Zones	2
Amazon Elastic Cloud Compute	2
Elastic Load Balancing	3
Amazon Elastic Block Store	3
Amazon Machine Image	3
Amazon Simple Storage Service	3
Amazon Virtual Private Cloud	4
AWS Security and Compliance	4
AWS Features Enabling Virtualization of Samsung PTT Solution	4
Samsung PTT Solution on AWS	6
Samsung PTT Solution Components	6
Samsung PTT Architecture on AWS	7
Benefits of Samsung PTT Solution on AWS	9
Samsung PTT on AWS Use Cases	11
Conclusion	11
Contributors	11
Document Revisions	11

Abstract

The Samsung Push to Talk (PTT) solution is a popular suite of integrated components that enables mobile workforce communication. This whitepaper provides an architectural overview for running the Samsung PTT solution suite on AWS.

Introduction

All major enterprises, public safety, and communications service organizations with mobile workforces can benefit from a Push to Talk (PTT) solution. The PTT solution is a two-way radio-type service that enables customers to push a button and instantly communicate with large audiences over a variety of devices and networks. Sectors such as construction, hospitality, security, oil and gas, utilities, manufacturing, field services, education, and transportation already rely on previous-generation technologies to perform this function. However, cloud adoption among enterprises is growing rapidly, with many adopting a cloud-first strategy for new projects and migrating their existing systems from on-premises to Amazon Web Services (AWS). Enterprises can deploy the Samsung PTT solution on AWS.

This whitepaper provides an overview of the Samsung PTT solution and a reference architecture for deploying Samsung PTT on AWS. We also discuss the benefits of running the Samsung PTT solution on AWS and various use cases.

AWS Overview

AWS provides on-demand computing resources and services in the cloud, with pay-as-you-go pricing. As of this publication, AWS serves over a million active customers in more than 190 countries and is available in 16 AWS Regions worldwide. You can access servers on AWS and log in, configure, secure, and operate them just as you would operate servers in your own data center.

When you use AWS resources for your compute needs, it's like purchasing electricity from a power company instead of running your own generator, and it provides many of the same benefits, including:

- The capacity you get exactly matches your needs.
- You pay only for what you use.
- Economies of scale result in lower costs.
- The service is provided by a vendor who is experienced in running large-scale compute and network systems.

AWS Infrastructure and Services for Samsung PTT Solution

This section describes the AWS infrastructure and services that are part of the reference architecture that you need to use to run the Samsung PTT solution on AWS.

Regions and Availability Zones

Each AWS Region is a separate geographic area that is isolated from the other Regions. Regions provide you the ability to place resources, such as Amazon Elastic Compute Cloud (Amazon EC2) instances, and data in multiple locations. Resources aren't replicated across Regions unless you do so specifically.

An AWS account provides multiple Regions so that you can launch your applications in locations that meet your requirements. For example, you might want to launch your applications in Europe to be closer to your European customers or to meet legal requirements.

Each Region has multiple, isolated locations known as Availability Zones. Each Availability Zone runs on its own physically distinct, independent infrastructure, and is engineered to be highly reliable. Common points of failure, such as generators and cooling equipment, aren't shared across Availability Zones. Each Availability Zone is isolated, but the Availability Zones in a Region are connected through low-latency links.

For more information about Regions and Availability Zones, see [Regions and Availability Zones](#) in the *Amazon EC2 User Guide for Linux Instances*.¹ For the most current list of Regions and Availability Zones, see [AWS Global Infrastructure](#).²

Amazon Elastic Cloud Compute

Amazon Elastic Compute Cloud (Amazon EC2) is a web service that provides resizable compute capacity in the cloud that is billed by the hour. You can run virtual machines (EC2 instances) ranging in size from 1 vCPU and 1 GB memory to 128 vCPU and 2 TB memory. You have a choice of operating systems

including Windows Server 2008/2012, Oracle Linux, Red Hat Enterprise Linux, and SUSE Linux.

Elastic Load Balancing

Elastic Load Balancing automatically distributes incoming application traffic across multiple Amazon EC2 instances in the cloud. It enables you to achieve greater levels of fault tolerance in your applications, seamlessly providing the required amount of load balancing capacity needed to distribute application traffic. Elastic Load Balancing can be used for load balancing web server traffic.

Amazon Elastic Block Store

Amazon Elastic Block Store (Amazon EBS) provides persistent block-level storage volumes for use with EC2 instances in the AWS Cloud. Each Amazon EBS volume is automatically replicated within its Availability Zone to protect you from component failure, offering high availability and durability. EBS volumes offer the consistent and low-latency performance needed to run your workloads.

Amazon Machine Image

An Amazon Machine Image (AMI) is simply a packaged-up environment that includes all the necessary bits to set up and boot your EC2 instance. Your AMIs are your unit of deployment. Amazon EC2 uses Amazon EBS and Amazon S3 to provide reliable, scalable storage of your AMIs so that we can boot them when you ask us to do so.

Amazon Simple Storage Service

Amazon Simple Storage Service (Amazon S3) provides developers and IT teams with secure, durable, highly-scalable object storage. Amazon S3 is easy to use. It provides a simple web services interface you can use to store and retrieve any amount of data from anywhere on the web. With Amazon S3, you pay only for the storage you actually use. There is no minimum fee and no setup cost.

Amazon Virtual Private Cloud

Amazon Virtual Private Cloud (Amazon VPC) lets you provision a logically isolated section of the AWS Cloud in which you can launch AWS resources in a virtual network that you define. You have complete control over your virtual networking environment, including selection of your own private IP address range, creation of subnets, and configuration of route tables and network gateways. You can leverage multiple layers of security, including security groups and network access control lists, to help control access to EC2 instances in each subnet. Additionally, you can create a hardware Virtual Private Network (VPN) connection between your corporate data center and your VPC, and then you can leverage the AWS Cloud as an extension of your corporate data center.

AWS Security and Compliance

The AWS Cloud security infrastructure has been architected to be one of the most flexible and secure cloud computing environments available today. Security on AWS is similar to security in your on-premises data center, but without the costs and complexities involved in protecting facilities and hardware. AWS provides a secure global infrastructure, plus a range of features that you can use to help secure your systems and data in the cloud. To learn more about AWS Security, see the [AWS Security Center](#).³

AWS Compliance enables customers to understand the robust controls in place at AWS to maintain security and data protection in the cloud. AWS engages with external certifying bodies and independent auditors to provide customers with extensive information regarding the policies, processes, and controls established and operated by AWS. To learn more about AWS Compliance, see the [AWS Compliance Center](#).⁴

AWS Features Enabling Virtualization of Samsung PTT Solution

The features used to support the function virtualization of Push to Talk solution from Samsung on AWS Cloud include following:

- **Elastic Network Adapter (ENA)** – ENA is the next generation network interface and accompanying drivers that provide enhanced

networking on EC2 instances. ENA is a custom AWS network interface optimized to deliver high throughput and packet per second (PPS) performance, and consistently low latencies on EC2 instances. Using ENA, customers can use up to 20 Gbps of network bandwidth on specific EC2 instance types. Open source licensed ENA drivers are currently available for Linux and Intel Data Plane Development Kit (Intel DPDK). The latest Amazon Linux AMI includes the ENA Linux driver support by default. ENA Linux driver source code is also available on [GitHub](#) for developers to integrate in their AMIs.⁵ There is no additional fee to use ENA. For more information, see the [Enhanced Networking on Linux](#) in the *Amazon EC2 User Guide for Linux Instances*.⁶

- **Support for single root I/O virtualization (SR-IOV)** – The single root I/O virtualization (SR-IOV) interface is an extension to the PCI Express (PCIe) specification. SR-IOV allows a device, such as a network adapter, to separate access to its resources among various PCIe hardware functions.
- **Support for data plane development kit (DPDK)** – The DPDK is a set of data plane libraries and network interface controller drivers for fast packet processing. The DPDK provides a programming framework and enables faster development of high-speed data packet networking applications.
- **Support for non-uniform memory access (NUMA)** – NUMA is a design where a cluster of microprocessors in a multiprocessing system are configured so that they can share memory locally. This design improves performance and enables expansion of the system. NUMA is used in a symmetric multiprocessing (SMP) system.
- **Support for huge pages** – Huge pages is a mechanism that allows the Linux kernel to use the multiple page size capabilities of modern hardware architectures. Linux uses pages as the basic unit of memory, where physical memory is partitioned and accessed using the basic page unit.
- **Support for static IP addresses** – Amazon EC2 instances can use static IP addresses (survives reboot), and these addresses can be associated with or dissociated from a different EC2 instance in any Availability Zone within a Region.

Samsung PTT Solution on AWS

This section covers the major components of the Samsung PTT solution and its architecture on AWS that you can use to deploy and configure it on AWS.

Samsung PTT Solution Components

The Samsung PTT solution offers advanced 3GPP Rel.13 MCPTT (Mission Critical Push-to-Talk) features, centralized online address book management, and security—all delivered over 4G LTE, 3G WCDMA/HSPA, and Wi-Fi networks. With PTT, users can carry a single device to conveniently access instant broadband data, voice service, workforce management, and mobile productivity applications. The Samsung PTT solution leverages eMBMS broadcast technology to transmit data to up to several thousand users within range of a given LTE base station. This method allows an extremely rapid flow of information during crisis situations, without slowing down traffic on the network.

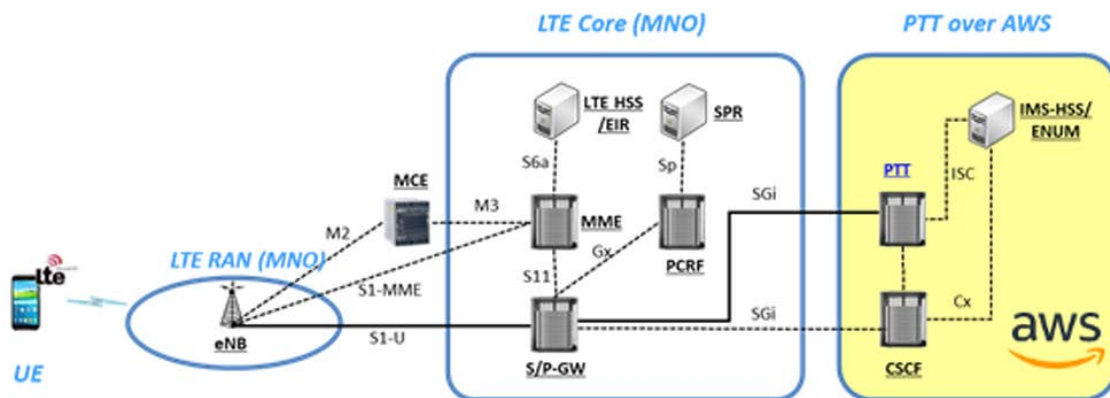


Figure 1: Push to Talk (PTT) network architecture

The solution consists of three main components:

- Samsung PTT server solution sends multimedia, such as video or high-quality images, to thousands of devices simultaneously using a single transmission channel. Each device seamlessly receives the incoming data at the same time, allowing real-time video communication among thousands of users. In contrast, when relying upon traditional unicast methods, in order to send multimedia to different devices, a single channel for each device is needed, consuming unnecessary air link

capacity, significantly degrading the quality of video and potentially causing video buffering or stuttering issues.

- **Samsung Call Session Control Function (CSCF)** is a collection of functional capabilities that play an essential role in the IP Multimedia Core Network Subsystem (IMS). The CSCF is responsible for the signaling that controls the communication of IMS User Equipment (UE) with IMS-enhanced services across different network access technologies and domains.
- **Samsung Home Subscriber Server (HSS)** is the main IMS database that also acts as a database in Evolved Packet Core (EPC). The HSS is a super home location register (HLR) that combines legacy HLR and authentication center (AuC) functions together for circuit-switched (CS) and packet-switched (PS) domains.

This component architecture integrates with Long-Term Evolution (LTE) handsets, eNodeB, and EPC components. We integrated this component architecture with AWS services via the public internet to create a test network. The next sections describe how we set it up.

Samsung PTT Architecture on AWS

The Samsung PTT solution setup included setting up a VPC with a public subnet that has a bastion host and three private subnets for CSCF, HSSs and PTT servers. The bearer packet processing acceleration was powered by the AWS ENA with DPDK applications and SR-IOV network port capabilities.

The EC2 instances within each of the private subnets reside in their respective placement groups, as shown in the following diagram.

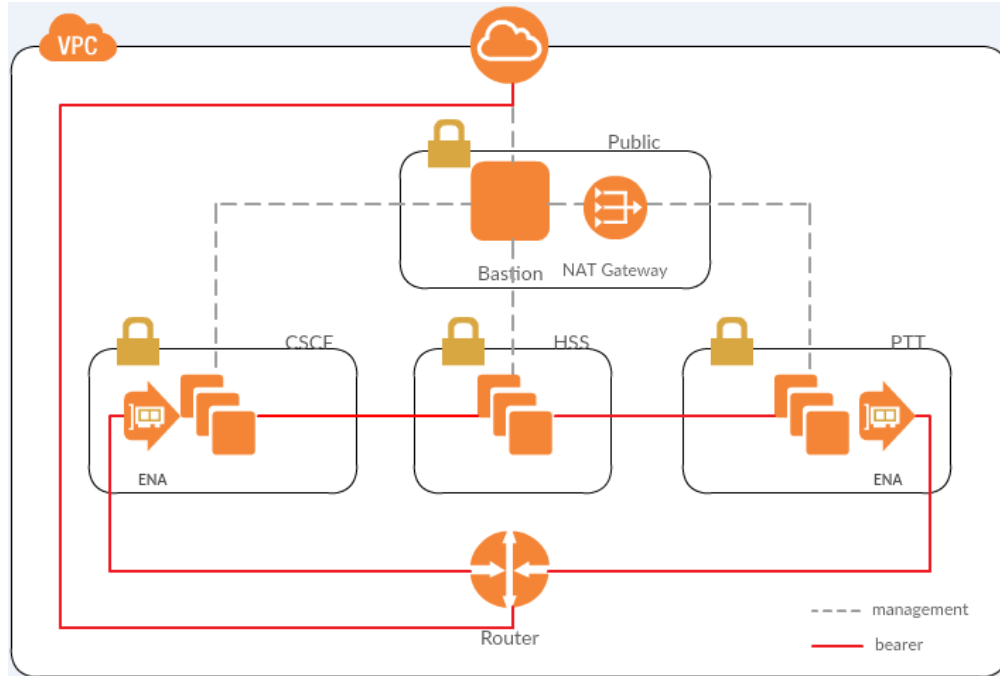


Figure 2: Push to Talk (PTT) deployment architecture on AWS

Effective and accurate dimensioning of the solution is critical for the virtual PTT solution. It’s always advisable to contact your Samsung team and get their input before implementing a solution for your organization. The configuration used for validation of the PTT solution on AWS is outlined in the following table, which lists each function, plane, number of instances, instance type, and feature that is enabled.

Table 1: EC2 Configuration used for Samsung PTT Solution Validation

Function	Plane	Number of Instances	Instance Type	Features Enabled
CSCF	Control plane	1	c4.4xlarge	DPDK, SR-IOV
CSCF	User plane	1	m4.2xlarge	DPDK, SR-IOV
PTT	Control plane	1	c4.4xlarge	DPDK, SR-IOV
PTT	User plane	1	m4.2xlarge	DPDK, SR-IOV
HSS	Control	1	m4.2xlarge	

Function	Plane	Number of Instances	Instance Type	Features Enabled
OSS	Operations & maintenance	2	m4.xlarge	Not applicable
Bastion	Management	1	T2.micro	Not applicable

Contact the Samsung team for accurate dimensioning of the solution for your organization.

Benefits of Samsung PTT Solution on AWS

The following sections outline the benefits of using Samsung PTT on AWS.

Cost Savings for Non-Production Environments

You can shut down your non-production environments when you aren't using them and save costs. For example, if you are using a development environment for only 40 hours a week (8 hours a day, 5 days a week), you can shut down the environment when it's not in use. You pay only for 40 hours of Amazon EC2 compute charges instead of 168 hours (24 hours a day, 7 days a week) for an on-premises environment running all the time. This can result in a saving of 75% for Amazon EC2 compute charges.

Unlimited On-demand Environments

In an on-premises environment, you usually have a limited set of environments to work with. Provisioning additional environments takes a long time or might not be possible at all. You don't face these restrictions when using AWS. You can create virtually any number of new environments in minutes as necessary.

You can have a different environment for each major project so that each team can work independently with the resources they need without interfering with other teams. Then, the teams can converge at a common integration environment when they are ready. You can terminate these environments when the project finishes and stop paying for them.

Lower Total Cost of Ownership

In an on-premises environment, you typically pay hardware support costs, virtualization licensing and support, data center costs, etc. You can eliminate or reduce all of these costs by moving to AWS. You benefit from the economies of scale and efficiencies provided by AWS and pay only for the compute, storage, and other resources that you use.

Right Size Anytime

Often customers oversize environments for initial phases, and then they're not able to cope with growth in later phases. With AWS, you can scale your organization's usage up or down at any time. You only pay for the computing capacity you use, for the duration that you use it. Instance sizes can be changed in minutes through the AWS Management Console, the AWS application programming interface (API), or the AWS Command Line Interface (AWS CLI).

Replace CapEx with OpEx

You can start a Samsung PTT solution implementation or project on AWS without any upfront cost or commitment for compute, storage, or network infrastructure.

No Hardware Costs

In an on-premises environment, you end up owning hardware that is depreciating in value every year. You are locked into the price and capacity of the hardware once it is acquired, plus you have ongoing hardware support costs.

With AWS, you can switch your underlying instances to faster, more powerful next-generation AWS instance types as they become available.

Low-Cost Disaster Recovery

You can build low-cost standby disaster recovery environments for your existing deployments and incur costs only for the duration of the outage.

No End of Life for Hardware or Platform

All hardware platforms have end-of-life dates, at which point the hardware is no longer supported and you are forced to buy new hardware again. In the AWS Cloud, you can simply upgrade the platform instances to new AWS instance types in a single click at no cost for the upgrade.

Samsung PTT on AWS Use Cases

Samsung PTT partners and customers are using AWS for a variety of use cases including the following:

- Implementing new Samsung PTT production environments
- Implementing disaster recovery environments
- Running Samsung PTT development, test, demonstration, proof of concept (POC), and training environments
- Scaling existing Samsung PTT production environments for incremental traffic
- Setting up temporary environments for migrations and testing upgrades
- Setting up temporary environments for performance testing

Conclusion

AWS can be an extremely cost-effective, secure, scalable, high-performing, and flexible option for deploying the Samsung PTT solution. This whitepaper outlines some of the benefits and use cases for deploying the Samsung PTT solution on AWS.

Contributors

The following individuals and organizations contributed to this document:

- Jeong Shang Ohn, Principal Engineer, Samsung Network Division
- Robin Harwani, Global Strategic Partner Solution Lead for Telecommunications, Amazon Web Services
- Andy Kim, Solution Architect, Amazon Web Services

Document Revisions

Date	Description
October 2017	First publication

Notes

- 1 <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-regions-availability-zones.html>
- 2 <https://aws.amazon.com/about-aws/global-infrastructure/>
- 3 <https://aws.amazon.com/security/>
- 4 <https://aws.amazon.com/compliance/>
- 5 <https://github.com/amzn/amzn-drivers>
- 6 <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/enhanced-networking.html>