AWS Whitepaper

Applying Security Practices to a Network Workload on AWS for Communications Service Providers

Applying Security Practices to a Network Workload on AWS for Communications Service Providers: AWS Whitepaper

Table of Contents

Applying Security Practices to a Network Workload on AWS for Communications Service Providers

Publication date: **June 30, 2023** (*Document revisions*)

Abstract

This whitepaper provides recommendations to Communication Service Providers (CSPs) on securing their telecommunications (telco) network workload on Amazon Web Services (AWS). These recommendations are based on the Security Pillar of the AWS Well-Architected Framework, and focus on AWS infrastructure and services. The Security Pillar provides guidance to help customers apply best practices in the design, delivery, and maintenance of an AWS workload. The information in this whitepaper informs how customers can introduce security controls into their workloads. By implementing these recommendations, CSPs can improve the security of their telco workload on AWS and help achieve their security goals and requirements.

Are you Well-Architected?

The AWS Well-Architected Framework helps you understand the pros and cons of the decisions you make when building systems in the cloud. The seven pillars of the Framework allow you to learn architectural best practices for designing and operating reliable, secure, efficient, cost-effective, and sustainable systems. Using the AWS Well-Architected Tool, available at no charge in the AWS Management Console, you can review your workloads against these best practices by answering a set of questions for each pillar.

Introduction

Running telecommunications (telco) network workloads on the public cloud enables CSPs to use the benefits of cloud computing for cost savings, elasticity, pay-as-you-go pricing, and supporting a global footprint. In addition, the underlying infrastructure of the public cloud is independently certified against many international frameworks. This provides a proven secure foundation on which to host workloads, lowering the overall security burden compared to on-premises deployments where CSPs are typically responsible for the security of the entire stack.

CSPs are looking for actionable guidance to design and manage the security of their workloads environments where they don't own the infrastructure. Another consideration for CSPs

contemplating the public cloud is compliance. CSPs generate increasing amounts of data containing personally identifiable information (PII), or subscriber data subject to regional and global regulations. Regulations in telco have a strong focus on security and require CSPs to implement state-of-the-art security measures to run and operate telecommunications and data processing systems. This whitepaper discusses domain security, data protection, and data privacy to help protect the data of telco network workloads on AWS. It provides guidance on how to manage, govern, and operate network workloads in AWS by recommending design principles, architectural concepts, and security controls that helps CSPs align with regulatory and compliance requirements.

The Security Pillar of the AWS Well-Architected Framework

The Security Pillar describes how to use cloud technologies to help protect data, systems, and assets in a way that can improve an AWS customer's security posture. It provides in-depth, best practice guidance for architecting secure workloads on AWS. The security pillar is made of seven design principles to help strengthen workload security:

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

The Shared Responsibility Model

Security and Compliance is a shared responsibility between AWS and the customer. AWS is responsible for protecting the infrastructure that runs all of the services offered in the AWS Cloud, known as *Security **of** the Cloud*. This infrastructure is composed of the hardware, software, networking, and facilities that run AWS Cloud services. The customer's responsibility is determined by the AWS Cloud services they select. This determines the amount of configuration work the customer must perform as part of their security responsibilities, known as *Security **in** the Cloud*. For example, for Amazon Elastic Compute Cloud (EC2) service, the customer will be responsible for the necessary security configurations and management from its networking, operating system, and application configuration including its patching and permissions. However, for abstracted services like Amazon Simple Storage Service (S3) where AWS operates the infrastructure, operating system and environment, the customer is provided access endpoints to use, store, and retrieve data. The customer will be responsible for managing the stored data to include applying encryption and appropriate access permissions. Applying this shared responsibility model to telco workloads means that, while AWS provides a secure infrastructure, CSPs and their Virtual Network Function/ Container Network Function (VNF/CNF) vendors should implement security measures to protect the workload. They can do this by adopting AWS security best practices and recommendations, and by following telco security standards as defined by multiple standard organizations such as 3GPP, ETSI, and IETF at the application level, to verify that the overall system is secured from each layer.

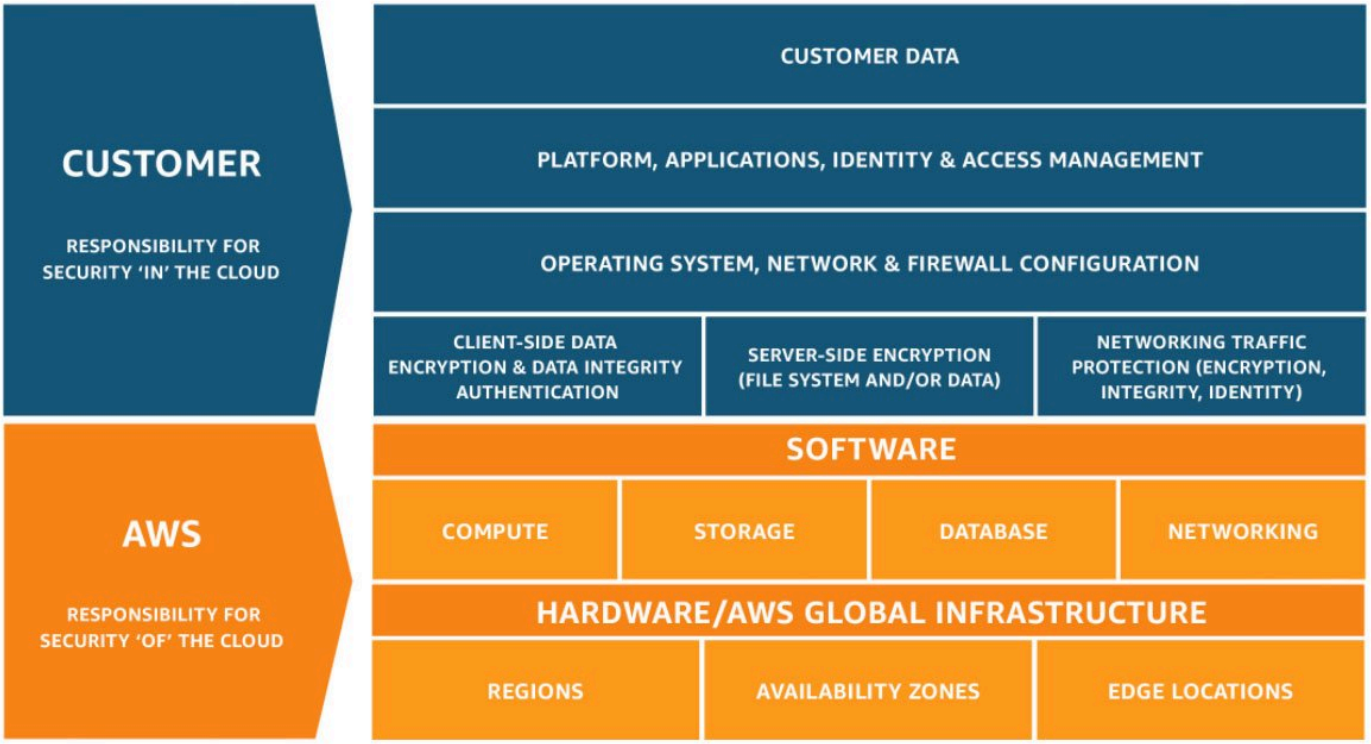

The Shared Responsibility Model in an AWS Region

Shared responsibility varies when using AWS services residing in a customer's data center; for example, when the Radio Access Network (RAN) functions such as Virtual Distributed Unit (vDU) are deployed on AWS Outposts. AWS Outposts is a family of fully-managed solutions delivering AWS infrastructure and services to virtually any on-premises or edge location. In AWS Outposts, the customer takes the responsibility of securing the physical infrastructure to host the AWS Outposts equipment in their own data centers. As a managed service, it inherits our well-tested security procedures, and includes built-in tampering and dedicated security components such as the Nitro Security card and key.

The preceding figure summarizes the shared responsibility model between AWS and the customer. AWS operates, manages, and controls the components from the host operating system and virtualization layer down to the physical security of the facilities owned by AWS. The customer assumes responsibility and management of the guest operating system and associated application or network functions as well as the configuration of the AWS services used.

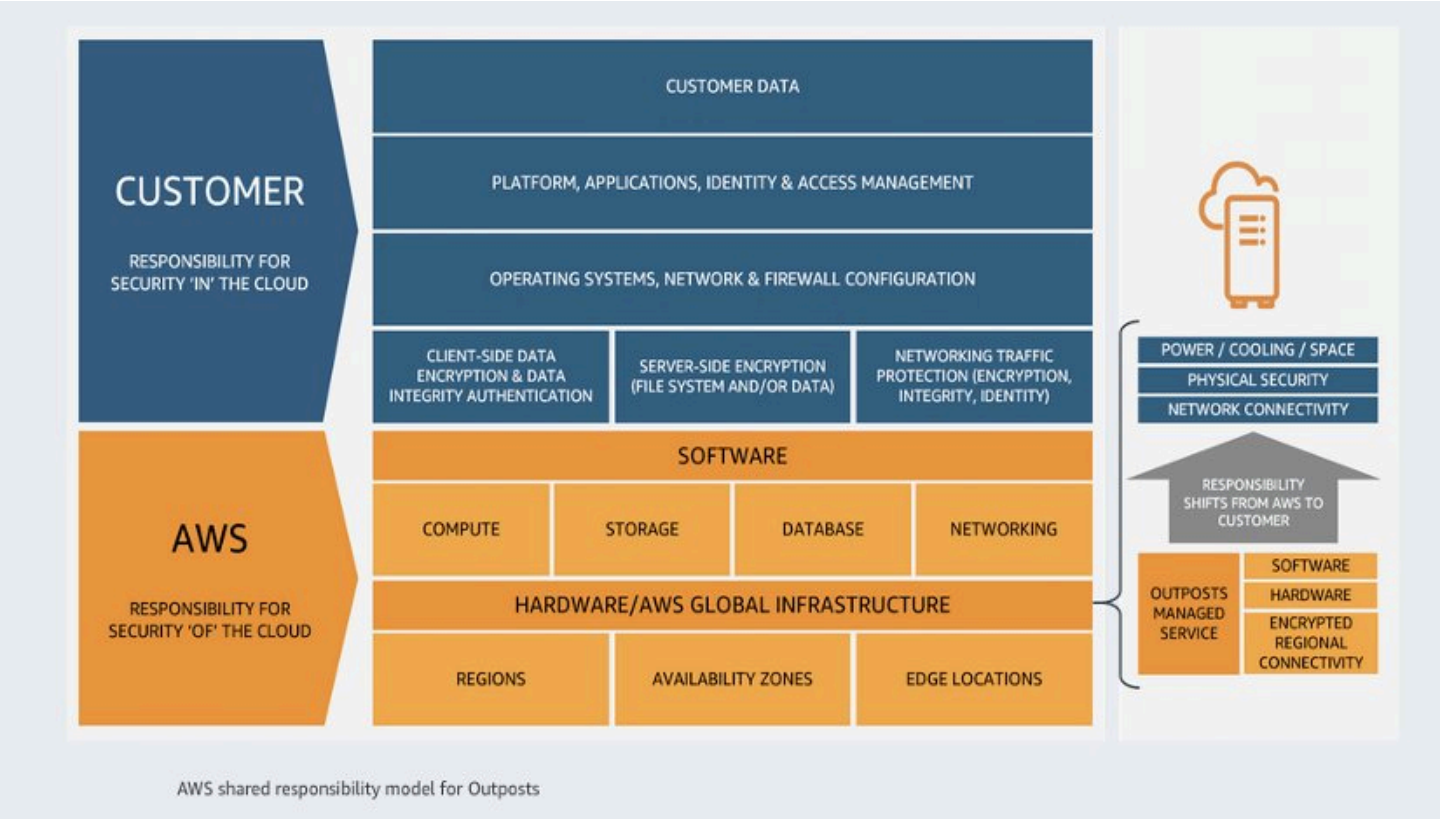

Shared Responsibility Model at the edge with AWS Outposts

The preceding figure shows an edge model with AWS Outposts, where the responsibility of the physical security, networking, cooling, and electricity for AWS Outposts is owned by the customer.

General design principles and controls

This section discusses the security design principles and controls that CSPs should consider when designing and running telco network workloads on AWS. It explains high-level security concepts, and what AWS services and service features can be used to support them. It also describes AWS infrastructures and how they help to secure telco workloads by design.

Security governance

Security governance is one of the foundational building blocks for defining and implementing a security strategy. Security governance involves defining how people, processes, and technology work together to support business objectives by defining policies and control objectives to help manage risks. This is a layered approach based on the AWS Shared Responsibility Model, and a typical starting point includes the separation of workloads across multiple environments (such as development, pre-production, production and so on) using AWS accounts as a logical boundary.

Landing zones are a mechanism to support the technical definition of security governance. A landing zone is a well-architected construct that uses multiple AWS accounts for workload isolation, billing, and to limit allocation. This is a starting point from which a CSP organization can quickly launch and deploy network workloads on AWS with confidence. Building a landing zone requires both technical and business decisions regarding account structure, networking, security, and access management, in line with business objectives.

A landing zone can be set up using AWS Control Tower and AWS Organizations. AWS Control Tower is a service that simplifies the orchestrating of multiple AWS services on a customer's behalf while helping to maintain the security and compliance needs of the organization. It provides automated mechanisms to set up a landing zone using AWS Organizations. AWS Organizations is an account management service that enables customers to consolidate multiple AWS accounts into an organization that customers create and centrally manage. It makes it easier for customers to manage and govern accounts, and follow security best practices based on AWS experience working with thousands of customers in their journey to the cloud. AWS Organizations includes account management and consolidated billing capabilities that enable you to better meet the budgetary, security, and compliance needs of your business.

Telco customers should consider the following when designing and implementing their governance model:

- **Roles and responsibilities boundaries** — For roles that have access to one or more accounts, attach role- and policy-based boundaries that limit the permission to the absolute minimum in line with the principle of least privilege.
- **Workload separation and/or isolation through accounts** — An AWS account is an isolated environment for a workload. Use account-level separation to isolate different network workloads or isolate account environments. This separation helps to limit the scope of impact should an event occur. Furthermore, separating workloads into different accounts prevents consuming resource quotas or potentially overprovisioning resources that can prevent other applications from working as intended. As an example, network workloads such as 5G and IP Multimedia Core Subsystem (IMS) should have their own separate accounts for each type of environment.

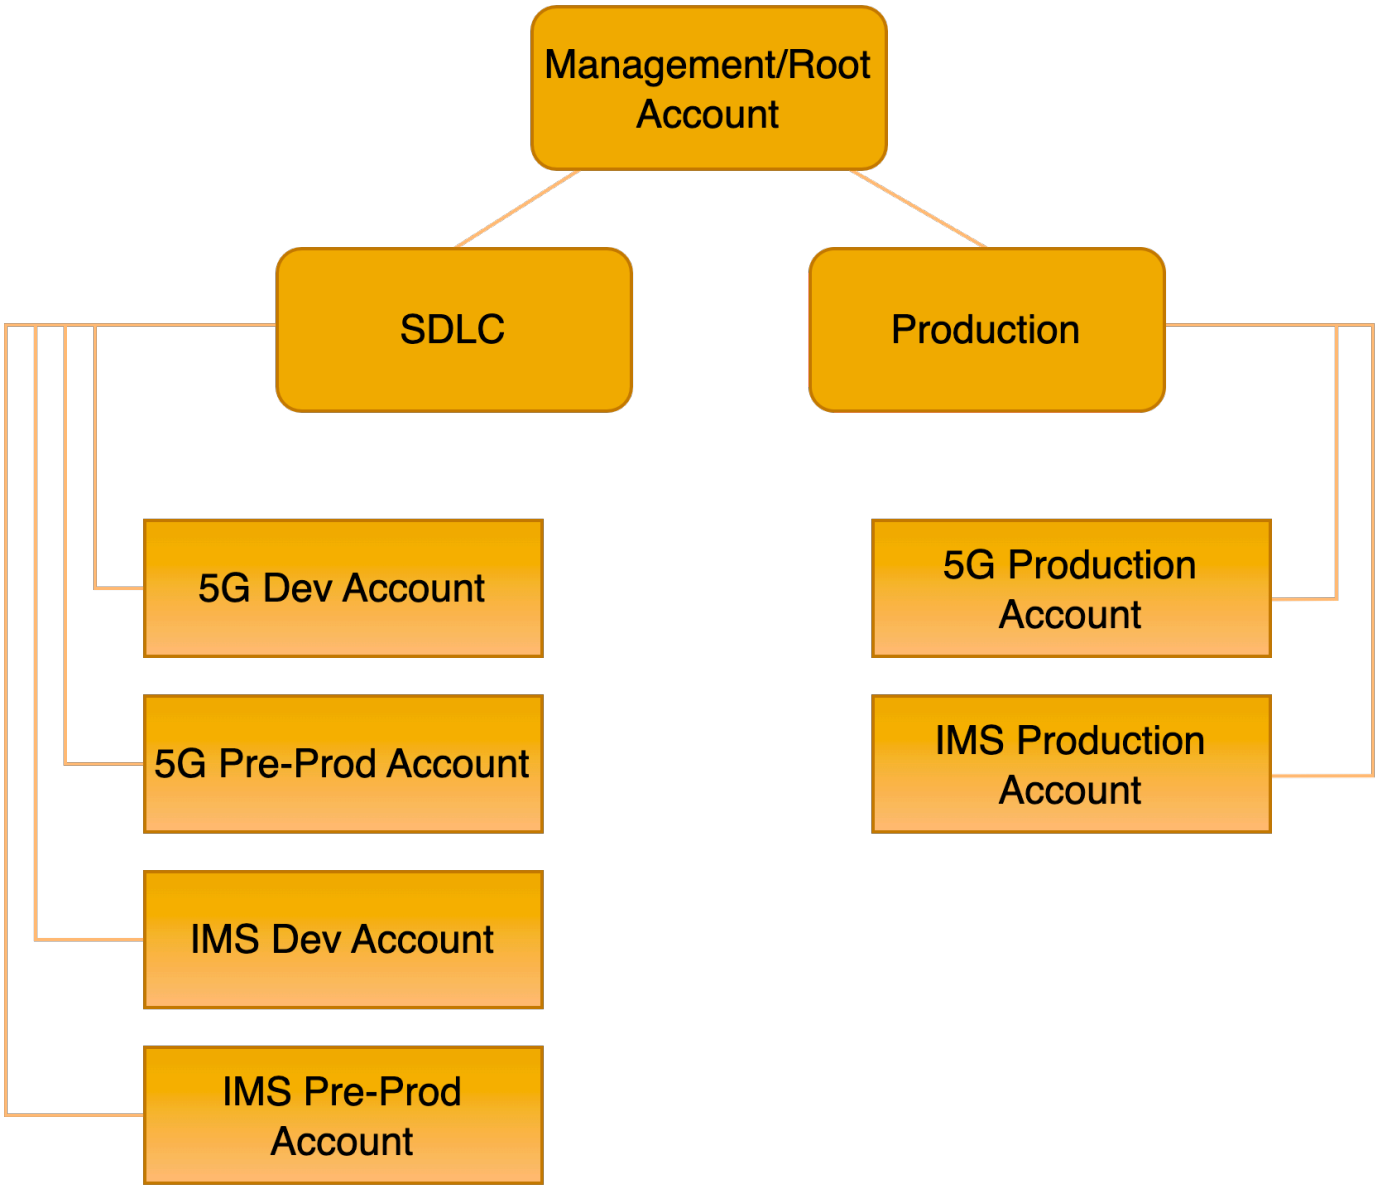

Example of a multi-account structure

- Compliance control identification — Different network workloads might have different risk profiles, requiring different control policies and mechanisms around them. Conformance packs, which are sets of AWS Config rules, and remediation actions can be used by customers to deploy in an account or across an organization in AWS Organizations. An AWS Config rule represents the desired configuration settings for specific AWS resources, or for an entire AWS account. Conformance packs provide customers with a general-purpose governance and compliance framework. They enable you to create security, operations, and cost-optimization governance.

  - **AWS Config** — Monitors and records your AWS resource configurations and allows you to automate the evaluation and remediation against desired configurations.

- **Central management** — Even though a landing zone contains multiple accounts, consistent management is key to maintain and administer a global policy framework to control account creation and management, including the available services and resources for each account as well as pre-defined and implemented security guardrails. CSPs should consider establishing a central function, which can act as both a management point and enablement function for the business and development teams. This could be in the form of, for example, a Cloud Center of Excellence (CCoE).
- **Data classification** — Classify data the workload will interact with. For example, control plane data, user plane data, and call data records are some data classifications that customers can use to determine if a chosen AWS service that processes or stores that data has the features to meet the security and compliance requirements specific to those data types.
- **Data isolation and anonymization** — Isolating data stores with the correct policies limits the number of people or applications that can access the data, and controls the exposure of traffic data or subscriber data. Call data records (CDRs) are one example where access should be limited only to people who need to access them. In the event that data is accessed, applying anonymization on sensitive details of CDR data adds another layer of exposure protection.
- **Limit allocation** — There are limits and quotas on AWS services per account. Separating workloads into different accounts prevents a single workload from consuming limits or potentially over-provisioning resources that may interrupt other workloads or incur unnecessary costs. For production workloads, make sure to check the limits so network functions are not prevented from scaling out the needed resources in case of sudden traffic peaks or disaster recovery (DR).
- **Security guardrails** — A *security guardrail* is a control that prevents deviations from expected or allowed behavior. Use security guardrails to implement preventive and detective controls across your AWS environment. You can implement these guardrails through the use of Service Control Policies (SCPs) enforced through AWS Organizations. SCPs are a type of organization policy that you can use to manage permissions in an organization.
  - **Preventive** – A *preventive guardrail* verifies that AWS accounts align with the policies you've set, because it disallows actions that lead to policy violations. AWS Control Tower comes with a default set of preventive guardrails based on best practices available to customers. Customers also have the option to set their own guardrails. For example, disallowing the creation of AWS resources that can compromise your network functions, such as the creation of a Network Address Translation Gateway (NATGW) inside your VPC that would provide a route from your

Amazon Elastic Compute Cloud (Amazon EC2) instance running network functions to connect to the internet.

- **Detective** – A *detective guardrail* detects noncompliance of resources within your accounts, such as policy violations, and can provide alerts through various mechanisms. For example, detecting a change in the security group of a network function compute instance to allow traffic from anywhere.

Secure workload operations

Operating network workloads in the cloud involves the whole lifecycle of a workload from design, build, run, and ongoing improvement. This includes applying DevSecOps principles. To achieve this, a recommended mechanism is to gather requirements and processes defined in the Operational Excellence Pillar of the Well-Architected Framework at an organizational and workload level, and apply them in all areas. The Operational Excellence Pillar discusses how an organization supports business objectives, the ability to run workloads effectively, gain insight into their operations, and to nearly continuously improve supporting processes and procedures to deliver business value.

Automation allows consistency and repeatability of processes. Customers should look to apply DevSecOps principles by aligning the security and development functions more closely: automating security processes, testing, and validating deployments help scale cloud operations. Adoption of AWS services can be supported across the development pipelines to apply security end-to-end across continuous integration/continuous deployment (CI/CD) pipelines, closed-loop workflows, and automated operations as the preferred methodology to deploy and manage the lifecycle of network workloads.

CSPs should also consider the following practices to support secure cloud operations:

- **Identify and prioritize threats and risks using a threat model** — Threat modeling provides a systematic approach to aid in finding and addressing security issues early in the design process. Earlier is better, because mitigations have a lower cost compared to later in the lifecycle. Use a threat model to identify and maintain an up-to-date registry of potential threats.
- **Identify and validate control objectives** — Based on your compliance requirements and the threats identified from the threat model, derive and validate the control objectives and controls to apply to the network workload. Ongoing validation of control objectives and controls helps measure the effectiveness of risk mitigation such as identifying your network workload's compliance requirements and identifying available AWS resources to assist you with your

compliance. More on AWS compliance resources can be found here and AWS security and compliance reports here.

- **Keep up-to-date with security recommendations** — Stay up-to-date with both AWS and industry security recommendations to evolve the security posture of the workload. AWS Security Bulletins contain important information about security and privacy notifications.
- **Evaluate and implement new security services and features regularly** — Evaluate and implement security services and features from AWS and AWS Partners that evolve the security posture of your workload. The AWS Security Blog highlights new AWS services and features, implementation guides, and general security guidance.
- **Automate testing and validation of security controls in pipelines** — Establish secure baselines and templates for security mechanisms that are tested and validated as part of CI/CD pipelines and processes. Use tools and automation to test and validate security controls nearly continuously in a DevOps fashion.
- **Enable logging** — Enable logging across components in conjunction with the Security Operations team to support visibility strategy and monitoring. Consider adopting a centralized logging approach for analysis and insights of security data by using Amazon Security Lake.
  - AWS CloudTrail — Provides event history of your AWS account activity, including actions taken through the AWS Management Console, AWS SDKs, command line tools, and other AWS services.
  - Amazon CloudWatch — A monitoring and observability service that collects monitoring and operational data in the form of logs, metrics, and events.
  - Amazon Security Lake — Automatically centralizes security data from AWS and third-party sources into a data lake stored in your AWS account. Amazon Security Lake gives you an understanding of the security posture across your entire organization.

Identity and access management (IAM)

This is a foundational element and requires a robust identity strategy that fits in with the overarching governance frameworks and supports business objectives.

Identity management

There are two types of identities which must be managed when approaching the operation of secure network workloads on AWS: *human identities* and *machine identities*.

- Implement identity management and permissions to verify that the right roles have access to the right resources under the right conditions. Apply the rule of least privilege, granting only the permissions required to complete a task.
- Define distinct IAM principals, differentiating between human (administrators, developers, operators, and consumers) and machine identities (network workload components, tools) with diverse IAM policies and permissions.
- Utilize a centralized and common Identity Provider (IdP), scoping multiple accounts in a common landing zone with services such as AWS IAM Identity Center and group their attributes appropriately.
- Use strong sign-in mechanisms for human principals, with specific password policies and multi-factor authentication (MFA) with software or hardware mechanisms.
- Require identities to dynamically acquire temporary credentials which have time-bound expiration, and engineer systems to require reauthentication once the session has expired.

Permissions management

Permissions control who can access what, and under what conditions. CSPs should assign permissions to specific human and machine identities to grant access to specific service actions on specific resources. Additionally, specify conditions that must be true for access to be granted.

- Define access requirements, supporting the principle of least privilege: The principle of least privilege determines that only the permissions needed to complete an activity should be granted. CSPs should look to clearly understand what is required to be performed across their network workload and grant only those permissions.
- Frequently review, refine, and reduce permissions: During early stages of design and test, permissions are often broad to allow for flexibility. Introduce a feedback cycle to review permissions and identify what has been used, and what permissions are not required.
- The use of IAM Access Analyzer helps you review and analyze the policies applied to the supported resources in your zone of trust. The organization or account you choose is known as the zone of trust for the analyzer.

Data protection

Applying security measures to protect data is an important consideration when running network workloads on AWS. The following measures to be discussed supplement the security features

and security mechanisms applied for LTE and 5G systems as explained in 3GPP TS 33.401 and TS 33.501 respectively. When deploying telco network workloads, country regulations and compliance frameworks mandate that data is protected at rest, in transit and, for some cases, in processing. In addition, control frameworks must be built around data handling to verify that mechanisms, tooling and processes are in place to prevent exposure of data.

- **Data identification** — The fundamental piece supporting data protection is knowing what data you need to protect. Data identification focuses on identifying and documenting the different points where data enters, is processed, and is stored throughout network workloads. It also includes the type of data, such as traffic information, subscriber information, or personally identifiable information (PII). As an example, specific network workload logs may contain identifiable subscriber information, subscriber activity and location, or network function descriptors that may be considered to be intellectual property from a given Independent Software Vendor (ISV).
- **Data classification** — The CSP should then build their own data classification framework which creates and defines classification *labels* (for example, sensitive, non-sensitive, and so on), with the support of examples and a data classification matrix. The data classification matrix is a guide on how to build the network workload architecture and which AWS services can be used.
- **Data tagging** — Aligned with the previous data classification, diverse policies for data classification and protection with encryption can be defined using resource tags. As an example, resources can be tagged for the associated network workload, the hosting environment, the existence of subscriber information, or other security considerations. This also enables advanced permissions management through the use of attribute-based access control (for users, services, and systems).
- **Data lifecycle management** — Based on the previous identification and sensitivity level, define data lifecycle policies, including the data retention duration, data destruction processes, data access management, data transformation, and data sharing. For example, it may be required that you hold sensitive data for a limited amount of time, after which it is automatically destroyed or anonymized and moved to an archive. In addition, customers can use AWS features that help protect against unintended or accidental data deletion during data lifecycle. Amazon Simple Storage Service (Amazon S3) objects help prevent objects or data from being deleted or overwritten for a fixed amount of time or indefinitely.
- **Identification and classification automation** — Automation supports the implementation of correct controls in a repetitive manner. For example, Amazon Macie recognizes stored sensitive data, such as PII, including names, phone numbers, Mobile Subscriber ISDN Numbers (MSISDN), International Mobile Subscriber Identity (IMSI), or International Mobile Equipment Identity (IMEI)

numbers using custom data identifiers. Amazon Macie provides dashboards and alerts that give visibility into how this data is being accessed or moved. Macie also enables the addition of resource tags to be added on objects based on their sensitivity.

Protect data at rest

Telecom network functions involve the processing of sensitive data such as subscriber data, access keys, location data, and PII that is protected by regulations such as the GDPR. For example, Call Detail Records, a formatted collection of a chargeable event used in billing and accounting, contains sensitive information that may identify specific callers and calls. It is critical to identify, classify and protect this data when saved on a persistent storage to help secure and comply with applicable regulations.

The following are recommendations about how data at rest can be protected on AWS:

- Encrypt all data at rest. Consider using a key management system such as AWS Key Management Service (AWS KMS) to generate keys for encryption and perform key lifecycle management. Optionally, customers can bring their own keys (BYOK) or connect AWS KMS to their own on-premises hardware security module (HSM) using AWS KMS External Key Store (KXS) for full control of the key material used for encryption.
- Use automation to validate and enforce data at rest controls nearly continuously. AWS Config rules can automate the validation when non-compliant settings have been applied. For example, an AWS Config rule that checks that Amazon Elastic Block Store (EBS) encryption is enabled by default. The rule is non-compliant if the encryption is not enabled.
  - AWS KMS — A fully managed key management service used to store and manage keys used to encrypt and decrypt data. Requests to use keys in AWS KMS are logged in AWS CloudTrail, so customers can understand who used which key, in what context, and when it was used. Event data logged to AWS CloudTrail cannot be altered. Also, AWS KMS is designed so that neither AWS (including AWS employees) nor third-party providers to AWS have the ability to retrieve, view, or disclose customers' primary keys in an unencrypted format.
  - AWS KMS External Key Store (XKS) — Customers who have a regulatory need to store and use their encryption keys on-premises or outside of the AWS Cloud can do so using this feature. This capability allows you to store AWS KMS customer managed keys on an HSM that they operate on-premises or at a location of their choice.

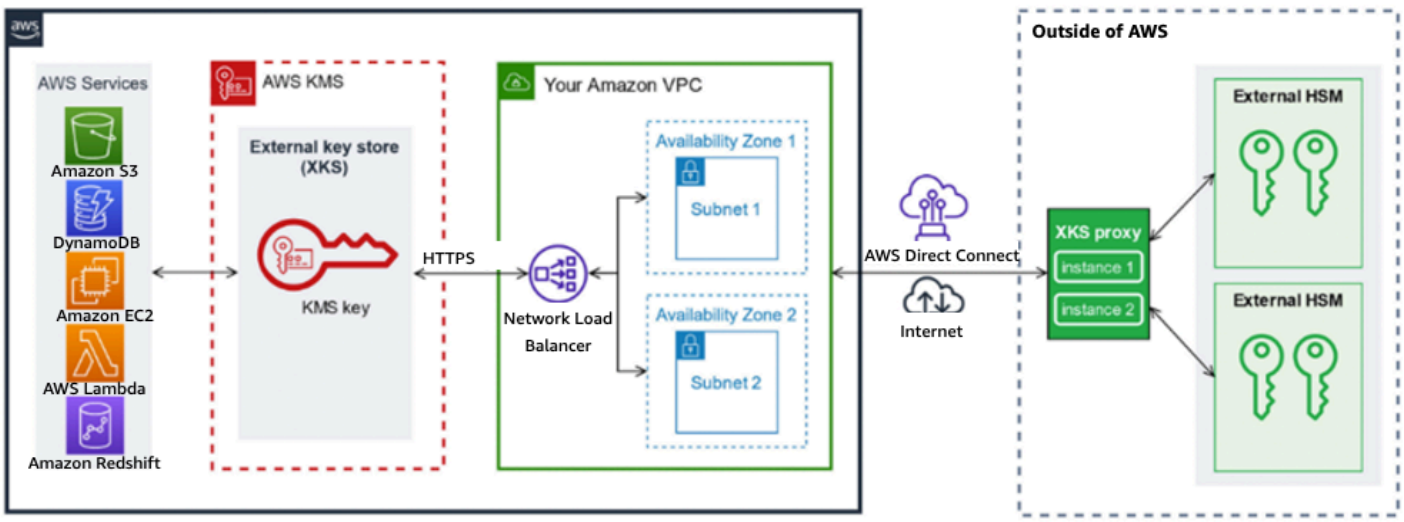

The HSMs that XKS communicates are on-premises

- AWS Config — A service that provides a detailed view of the configuration of AWS resources on an AWS account. It also provides information on how resources are related to one another, and how they were configured in the past.
- For workloads deployed using instance store volumes, data on Non-Volatile Memory Express (NVMe) instance store volumes are encrypted using an XTS-AES-256 cipher implemented on a hardware module on the instance itself. The keys are generated by, and only reside within, the hardware module, which is inaccessible to AWS personnel. For more information, refer to Data Protection in Amazon EC2.

Protect data in transit

3GPP TS 33.210 defines the need to implement security precautions to protect network domains using IPSEC and Transport Layer Security (TLS) encryption. These security precautions are supported in AWS, depending on the nature of the traffic, using different AWS services.

- Enforce encryption in transit. Data in transit is the data that is sent from one network function to another. This includes both entering an AWS environment from the on-premises network, or within a VPC or subnet in an AWS account.
- For Transmission Control Protocol (TCP) traffic, use TLS encryption. There are two AWS services for issuing and deploying X.509 certificates, whether they are public- or private-facing certificates, customized certificates, certificates you want to deploy into other AWS services, or automated certificate management and renewal.

  - **AWS Private CA**—Enterprise customers use this service for building a public key infrastructure (PKI) intended for private use within the organization inside the AWS Cloud. With AWS Private CA, you can create your own certificate authority (CA) hierarchy and issue certificates with it for authenticating internal users, computers, applications, services, servers, and other devices, and for signing computer code. Certificates issued by a private CA are trusted only within your organization, not on the internet.
  - **AWS Certificate Manager (ACM)** — This service manages certificates for customers who need a publicly trusted secure web presence using TLS. A managed service that handles the complexity of creating, storing, and renewing public and private SSL/TLS X.509 certificates and keys.
- Block unsecured ports such as HTTP using Amazon Virtual Private Cloud (Amazon VPC) security groups or network access control list (network ACL) rules. Use AWS Config to monitor unsecured security group rules. AWS Network Firewall and AWS DNS Firewall can also be used to enable the blocking of insecure methods of communication to and from the AWS environment.
- For non-TCP traffic such as UDP and SCTP, CSPs can apply transport encryption such as IPsec to help protect data in transit between network functions. For example, in transit encryption between on-premises and AWS, CSPs can use AWS Site-to-Site VPN with Virtual Private Gateway or a Transit Gateway. Alternatively, CSPs can use VPN virtual appliance on AWS to establish IPsec connection with on-premises network. Using an IPsec virtual appliance provides a better single tunnel bandwidth connection when compared to an AWS Site-to-Site VPN.
- Use dedicated private connection with AWS Direct Connect between your on-premises networks and AWS. Direct Connect links support L2 encryption in MACsec. CSPs should evaluate if MACSec encryption is sufficient for encryption in transit, because this would remove the need to use IPsec tunnels.
- If encryption in transit is required between EC2 instances running network functions and IPsec implementation is not possible for non-TLS supporting traffic, consider using Nitro VPC Encryption. Specific AWS instance types use the offload capabilities of the underlying AWS Nitro System hardware to automatically encrypt in-transit traffic between specific type of instances, using Authenticated Encryption with Associated Data (AEAD) algorithms with 256-bit encryption. Data Protection in Amazon EC2 talks more about this feature and the known considerations.

Protect data in process

Confidential computing is a term used to describe the ability to protect sensitive data in use by encrypting it while it is being processed by compute. Confidential computing provides an

additional layer of security for sensitive data, as it helps prevent unauthorized access to the data even if a bad actor is able to compromise the system where the data is being processed.

- Confidential computing is defined as the use of specialized hardware and associated firmware to protect customer data during processing from outside access. AWS Nitro System is a specialized hardware and is the underlying foundation for modern Amazon EC2 instances. The AWS Nitro system was designed to have no AWS operator access; there's no mechanism for a system or person to log in to EC2 servers (the underlying host infrastructure), read the memory of EC2 instances, or access data stored on instance storage and encrypted Amazon Elastic Block Store (Amazon EBS) volumes.
- AWS Nitro System is a combination of dedicated hardware and lightweight hypervisor which delivers practically all of the compute and memory resources of the host hardware to the instances for better overall performance and security. The AWS Nitro System enforces separation of duties and allows only the principals who have been specifically granted access to the data the ability to access it. For more information, refer to the Confidential computing: an AWS perspective blog post and *The Security Design of the AWS Nitro System* whitepaper.
- In addition to the AWS Nitro System, use instance types that have built-in memory encryption when required as an additional measure for protecting data in processing. AWS offers:
  - 3rd generation Intel Xeon scalable processors (Ice Lake) instances, with always-on memory encryption using Total Memory Encryption (TME): for example, M6i, C6i, R6i
  - 3rd generation AMD EPYC processors (Milan) instances, which includes support for AMD Transparent Single Key Memory Encryption (TSME), for example, M6a, C6a, R6a
  - Graviton2 and Graviton3 with always-on memory encryption, dedicated caches for every vCPU, and support for pointer authentication: for example, C6g, M6g, R6g

Data access control

When running network workloads on AWS, the compromise of subscriber data such as that found on the Home Subscriber Subsystems (HSS) or User Data Management (UDM) functions, or stored data such as CDRs, can be prevented by adopting data access control measures such as:

- Enforce controls for data access using the principle of least privilege and specific IAM principals that include prescriptive data access.
- Use resource policies, such as S3 bucket policies, to enforce data isolation boundaries between accounts or within the same account; for example, allowing only certain IAM principles access to specific S3 buckets which contain data of a certain type or classification.

- Consider enforcing specific rules such as preventing files or objects from being deleted or overwritten for a fixed amount of time, or indefinitely using Amazon S3 Object Lock.
- Consider automating the detection of unintended data access. With tools such as Amazon GuardDuty, CSPs can automatically detect suspicious activity or attempts to move data outside of defined boundaries.
  - Amazon GuardDuty — a nearly continuous security monitoring service that analyzes and processes data sources such as VPC Flow Logs, AWS CloudTrail management event logs, CloudTrail S3 data event logs, EKS audit logs, and DNS logs. It uses threat intelligence feeds such as lists of malicious IP addresses and domains, and machine learning (ML) to identify unexpected and potentially unauthorized and malicious activity within your AWS environment.

- Consider establishing an organization wide data perimeter. Use permission guardrails that restricts access outside of an AWS organization boundary. Use resource-based policies on AWS resources, service control policies to define permission on accessing AWS resources, and VPC endpoint policies to control access on VPC endpoints. More on data perimeters on AWS can be found here.

Infrastructure protection

Following are some recommended ways to protect your infrastructure on the network layer and on the computer layer.

Protect the network

Appropriate network protection controls that complement the business objective and provide the business teams confidence that their network workloads can be operated securely should be introduced into the design of the workload architecture. This section discusses how you can help protect your network domain on the infrastructure level.

Protecting the network – at the edge

- Analyze the connectivity requirements of the application and apply multiple controls with a defense-in-depth approach for both inbound and outbound traffic, including the use of security groups (stateful firewall), network Access Control Lists (ACLs) (stateless firewall), subnets, and route tables. Security groups and network ACLs are means to secure traffic that leaves and enters the VPC, which is considered to be your network domain.

- Analyze the option of using native edge protection services such as AWS WAF, AWS Shield, AWS Firewall Manager, and AWS Network Firewall to add additional layers of protection. These are native AWS services that customers can use without the heavy lifting of provisioning and maintaining the needed infrastructure and service. These services can be optionally applied in front of the Network Exposure Function (NEF). The NEF is a component of the 5G SBI architecture that is used to expose network information to external consumers. This information can include network resources, services, capabilities, and performance characteristics, and it can be accessed by other network functions or external entities through a set of APIs (Application Programming Interfaces).
  - AWS WAF — A web application firewall that helps protect your web applications or APIs against common web exploits and bots that may affect availability, compromise security, or consume excessive resources.
  - AWS Shield — An AWS service for the protection against Distributed Denial of Service (DDoS) attacks for applications running on AWS. There is also an enhanced version in AWS Shield Advanced. AWS Shield Advanced provides enhanced protections for applications and provides always-on, flow-based monitoring of network traffic and active application monitoring to provide near real-time notifications of suspected DDoS incidents.
  - AWS Firewall Manager — A service that simplifies your administration and maintenance tasks across multiple accounts and resources for a variety of protections, including AWS WAF, AWS Shield Advanced, Amazon VPC security groups, and AWS Network Firewall.
  - AWS Network Firewall — A stateful, managed, network firewall and intrusion detection and prevention service for VPC.

Protecting the network – within the customer boundaries

- Consider using VPC endpoints to privately connect to supported AWS services and VPC endpoint services powered by AWS PrivateLink. AWS PrivateLink provides private connectivity between VPCs, supported AWS services, and your on-premises networks without exposing your traffic to the public internet.
  - VPC Endpoint — Enables connections between a VPC and supported services, without requiring that you use an internet gateway, NAT device, VPN connection, or AWS Direct Connect connection. In addition, VPC endpoint access is controlled using an endpoint policy. An endpoint policy is a policy document that controls which AWS principal (e.g., user, roles, etc.) can use the VPC endpoint to access an endpoint service.

- For traffic between customer on-premises deployment (such as Outposts) and AWS Regions, there are multiple layers of encryption to be considered. At the transport layer, for Outposts, AWS encrypts in-transit data between your Outposts instance and its AWS Region. For more information, refer to Connectivity through service links.
- For on-premises commercial off-the-shelf (COTS) hardware, consider using AWS Direct Connect to connect your on-premises deployment to the AWS Region. You can use AWS Direct Connect connections that support MACsec to encrypt your data from your on-premises location to the AWS Direct Connect location. You can also combine AWS Site-to-Site VPN with Direct Connect to provide encryption at the transport layer (for example, IPsec VPN). At the application layer, use SSL/TLS to communicate with AWS resources, as mentioned in the Protect data in transit section.

Protect the compute

To protect the compute in your organization:

- **Use AWS Nitro-based instances** — AWS Nitro System is the underlying foundation for modern Amazon EC2 instances. There is no operator, administrator, or root access for administration. Access is strictly limited to a set of authenticated, authorized, and audited administrative APIs. None of the APIs have the capability to access customer data.
  - Nitro Trusted Platform Module (NitroTPM) — NitroTPM can be used for attestation, a process to demonstrate that an EC2 instance meets pre-defined criteria, allowing you to gain confidence in its integrity. It can be used to authenticate an instance requesting access to a resource (such as a service or a database) to be contingent on its health state (for example, patching level, presence of mandated agents, and so on). For example, a private key can be "sealed" to a list of measurements of specific programs allowed to "unseal."
- **Evaluate the use of pre-hardened compute** — Hardened images reduce exposure to unintended access by hardening operating systems and minimizing the components, libraries, and external services that are in use.
- **Patch management** — Rather than patching long lived instances, architect and engineer your workload to allow for newly patched images to take the place of outdated images.

  Consider scanning your container images with Amazon ECR or Amazon Inspector:

  - Amazon ECR — An AWS container image registry service. It can perform image scanning that can help identify software vulnerabilities in container images.

- Amazon Inspector — A vulnerability management service that nearly continuously scans AWS workloads for software vulnerabilities and unintended network exposure. Amazon Inspector automatically discovers and scans running Amazon EC2 instances, container images in Amazon Elastic Container Registry (Amazon ECR), and AWS Lambda functions for known software vulnerabilities and unintended network exposure.

Threat detection and incident response

Detection of threat events is underpinned by appropriate logging and monitoring of events and understanding the different threats to the workload. Threat modeling provides a systematic approach to identifying risks and threats specific to your workload, and informs what logging mechanisms need to be in place to support the detection and alerting of those threats.

Configure logging throughout the workload, including application logs, resource logs, and AWS service logs. For example, verify that AWS CloudTrail, Amazon CloudWatch Logs, Amazon GuardDuty, and AWS Security Hub are enabled for the accounts within your organization.

A foundational practice is to establish a set of detection mechanisms at the account level. This base set of mechanisms is aimed at recording and detecting a wide range of actions on the resources in your account. They allow you to build out a comprehensive detective capability with options that include automated remediation and partner integrations to add functionality.

In AWS, services that can implement this base set include:

- AWS CloudTrail — Provides event history of your AWS account activity, including actions taken through the AWS Management Console, AWS SDKs, command line tools, and other AWS services.
- AWS Config — Monitors and records your AWS resource configurations and allows you to automate the evaluation and remediation against desired configurations.
- Amazon GuardDuty — A threat detection service that nearly continuously monitors for malicious activity and unauthorized behaviour to help protect your AWS accounts and workloads. GuardDuty also provides threat detection for EKS clusters. We are seeing more Telco network workloads move to containerize their workloads for improved resilience, scaling, and operational management. Getting appropriate visibility and alerting from clusters is an important aspect of the overall security strategy.
- AWS Security Hub — Provides a single place that aggregates, organizes, and prioritizes your security alerts, or findings, from multiple AWS services and optional third-party products to

give you a comprehensive view of security alerts and helps you understand your overall security posture across all of your AWS accounts

Many core AWS services provide service-level logging features such as Amazon VPC. Amazon VPC Flow Logs enable you to capture information about the IP traffic going to and from network interfaces, subnet, or VPC. This information can provide valuable insight on connectivity history, and cue automated actions based on anomalous behaviour. This may be particularly important for network workloads which have a heavy traffic flow of subscriber data, and so on.

Visualizing high-volume data such as VPC Flow Logs is important to provide analytics and trend analysis. There are example solutions provided by AWS Specialists using other AWS services such as Amazon Athena and Amazon QuickSight to support efforts in storing, analyzing, and visualizing log data with more flexibility.

Responding to security events helps to reduce the impact should any occur. Therefore, it's imperative that, as a CSP, the organization is educated and well-prepared for different types of events. Levels of preparation strongly affect the ability of teams to cooperate effectively during an incident. Using automation to investigate and remediate events also reduces human effort and error, and enables CSPs to scale investigation capabilities. Regular reviews will help tune automation tools, and nearly continuously iterate.

CSPs should define incident response objectives and performance indicators, and define response playbooks or plans which help guide the responders on the correct series of steps for specific types of incidents. Overall, the CSP should maintain an appropriate strategy that covers education, preparation, simulation, and iteration for response activities. More details on these can be found on the Incident Response section of the Security Pillar or the NIST SP 800-61 Guide on Computer Security Incident Handling.

Example architectures and associated security considerations

To support telco customers and CSPs with additional context on the security considerations on network workloads, we have provided some example architectures of 5GC workloads along with a description of how AWS services have been applied to support security objectives.

Example architecture #1

An example architecture of a 5GC workload running in an AWS Region. The 5G Control Plane and User Plane are running in the Region and interconnected with an on-premises environment.

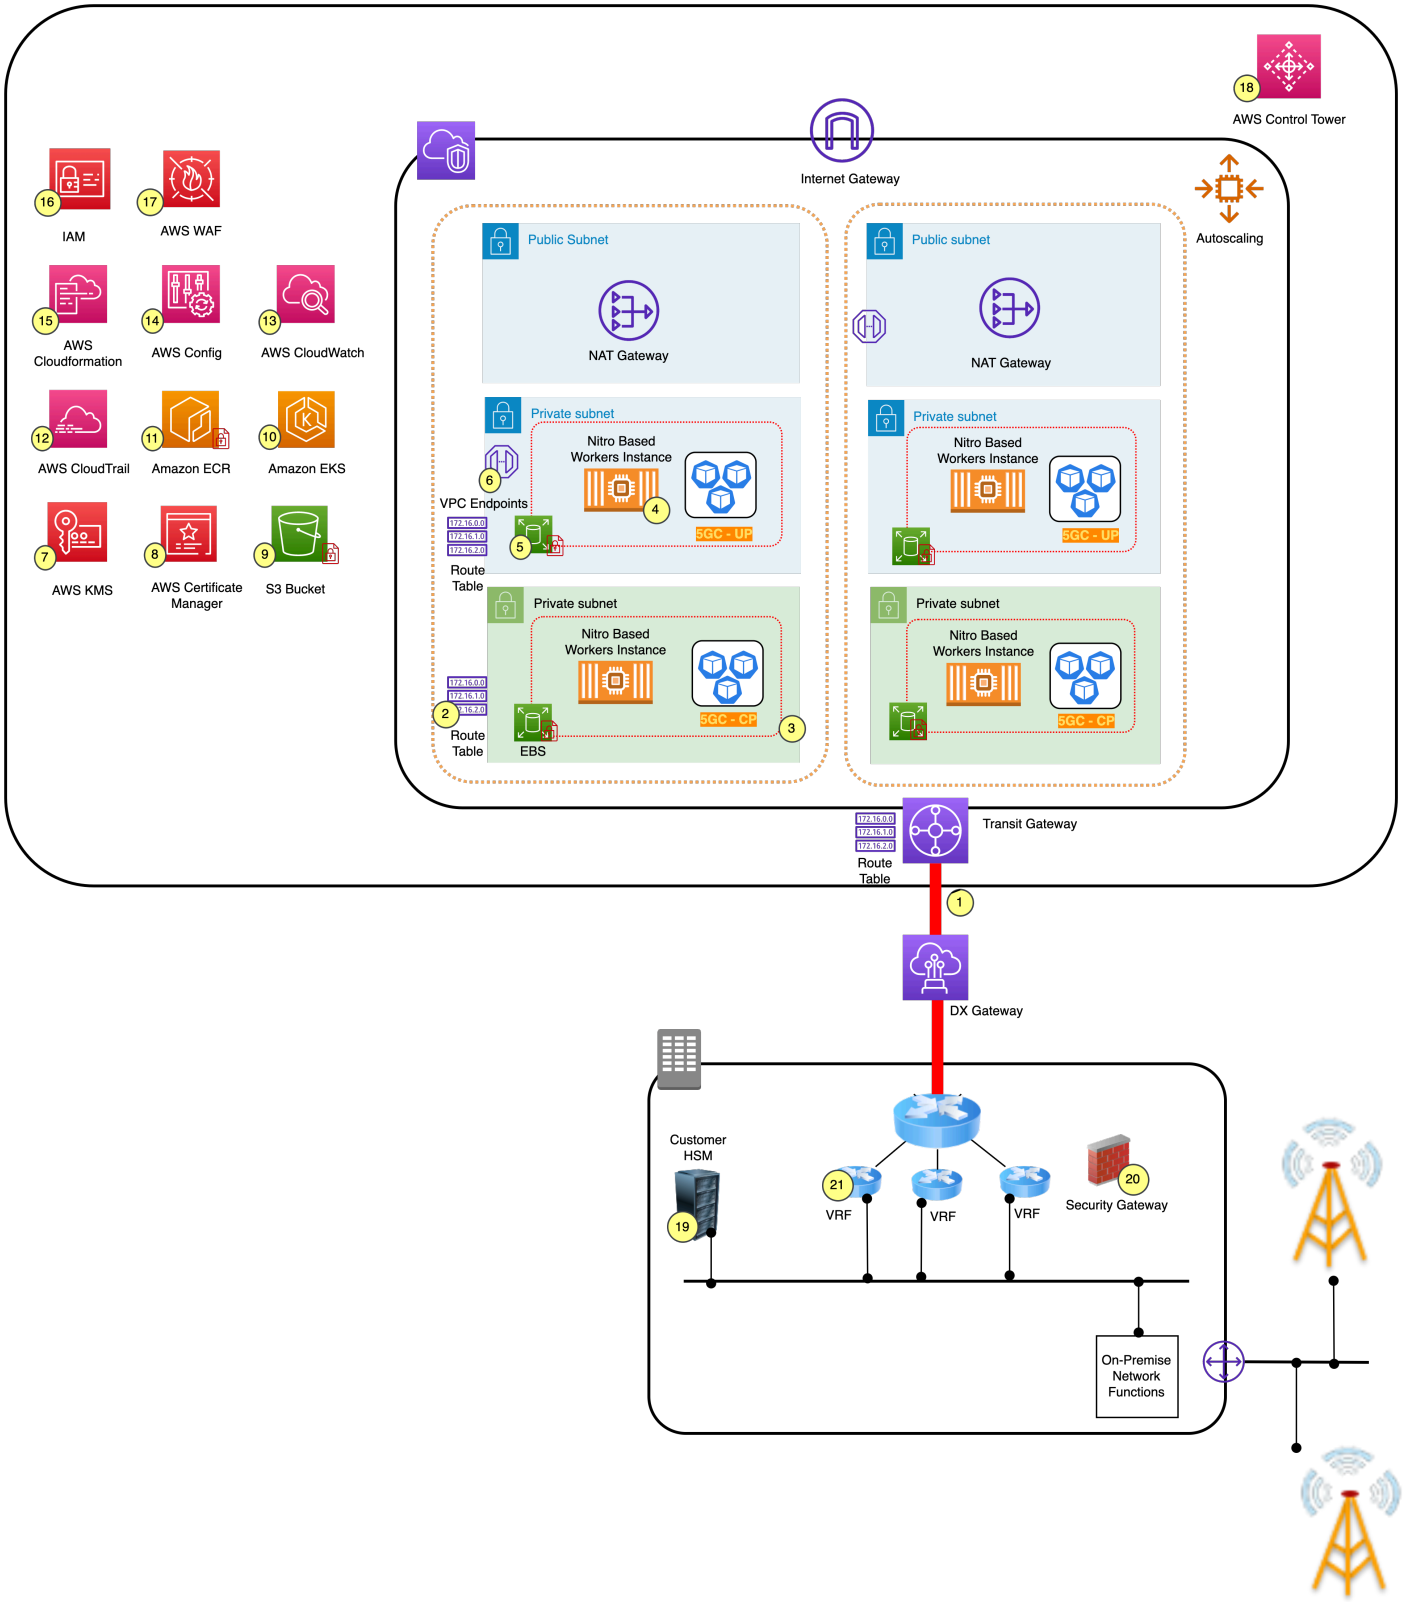

Architecture of 5G Core on AWS Region

Security description of the AWS services used in the example architecture of 5G core deployment on the Region:

1. AWS Direct Connect instances are physical links that connect the customer's on-premises network to AWS over a standard Ethernet fiber-optic cable. AWS Direct Connect connections support MACsec to encrypt data between on-premises facilities to the AWS Direct Connect location. MAC Security (MACsec) is an IEEE standard that provides data confidentiality, data integrity, and data origin authenticity. Network traffic is to be encrypted over an AWS Direct Connect line using IPSEC or TLS encryption, as indicated on 3GPP TS 33.210, to secure a network domain.
2. *VPC routing tables* are route tables used to control where network traffic is directed. Customers can remove direct routing towards the internet from the compute nodes.
3. Traffic going in and out of the instances are filtered using security groups. Security groups are stateful firewall rules that filter inbound traffic based on source address. In addition, there are network ACL rules that can filter traffic on a subnet level. Network ACLs are stateless firewall rules.
4. Nitro hardware-based instances are the compute or worker nodes where applications or network functions are deployed.
5. Persistent data at rest stored in EBS volumes are encrypted by a data key managed by AWS KMS. This verifies that data such as CDRs and subscriber profiles are protected.
6. Access to AWS services that do not reside inside the VPC is through VPC endpoints. These are endpoints that enable you to privately connect to an AWS service without traversing the internet.
7. Use AWS Key Management Service for management of encryption keys — AWS KMS supports BYOK. Alternatively, connect to on-premises HSM using AWS KMS XKS.
8. Use AWS Certificate Manager to manage imported SSL/TLS certificates. Certificates are for encrypting data in transit to support 3GPP TS 33.210.
9. Snapshots, Amazon Machine Images (AMIs) , manifest files, or backup data can be stored in Amazon S3 storage. Data at rest is encrypted using AWS KMS, and access to data is restricted with IAM policies.
10. Amazon EKS is used for Kubernetes-based container orchestration. Amazon EKS is a service used to run Kubernetes on AWS without needing to install, operate, and maintain your own Kubernetes control plane or nodes. For more information on security concepts using EKS, refer to Security in Amazon EKS.

11 Amazon ECR is used to store container images. Amazon ECR can also be used to scan images for security vulnerabilities. Images are encrypted using AWS KMS.

12 AWS CloudTrail helps enable governance and supports operational and risk auditing of an AWS account. Actions taken by a user, role, or an AWS service are recorded as events in CloudTrail. Events include actions taken in the AWS Management Console, AWS Command Line Interface, and AWS SDKs and APIs.

13 Amazon CloudWatch monitors AWS resources and applications that run on AWS in near real-time. Amazon CloudWatch can be used to collect and track metrics, which are variables that can be measured for resources and applications. Logs from Intrusion Detection Systems (IDS) and Intrusion Prevention Systems (IPS) firewall applications can be forwarded to CloudWatch and can be used for future investigations.

14 AWS AppConfig provides a detailed view of the configuration of AWS resources in an AWS account. This includes how the resources are related to one another, and how they were configured in the past so that customers can see how the configurations and relationships change over time. AWS Config rules are used to evaluate the configuration settings of AWS resources. When AWS Config detects that a resource violates the conditions in one of its rules, AWS Config flags the resource as non-compliant and sends a notification. AWS Config nearly continuously evaluates AWS resources as they are created, changed, or deleted. As an example, AWS Config can detect changes in the security groups and flag them as non-compliant.

15 Amazon CloudFront is an infrastructure as a code (IaaC) service that helps set up AWS resources automatically. Using IaaC templates helps eliminate human errors caused by manually configuring resources, security and networking rules.

16 AWS Identity and Access Management helps to securely control access to AWS resources. IAM is used to control who is authenticated (signed in) and authorized (has permissions) to use AWS resources.

17 AWS WAF is a firewall that helps protect application endpoints or APIs against common web exploits and bots that may affect availability, compromise security, or consume excessive resources.

18 AWS Control Tower provides the simplest way to set up and govern a secure, multi-account AWS environment.

19 A customer-owned on-premises HSM generates cryptographic keys for importing to AWS KMS, or use with AWS KMS XKS.

20 Customer security gateways (SEGs) are entities on the borders of the IP security domains used for securing native IP-based protocols.

21. Virtual routing and forwarding (VRF) devices, virtual router, and forwarding devices are used to segregate the VPN.

Example architecture #2

An example architecture of a 5GC workload with AWS Outposts. The 5G control plane and user plane are running on-premises.

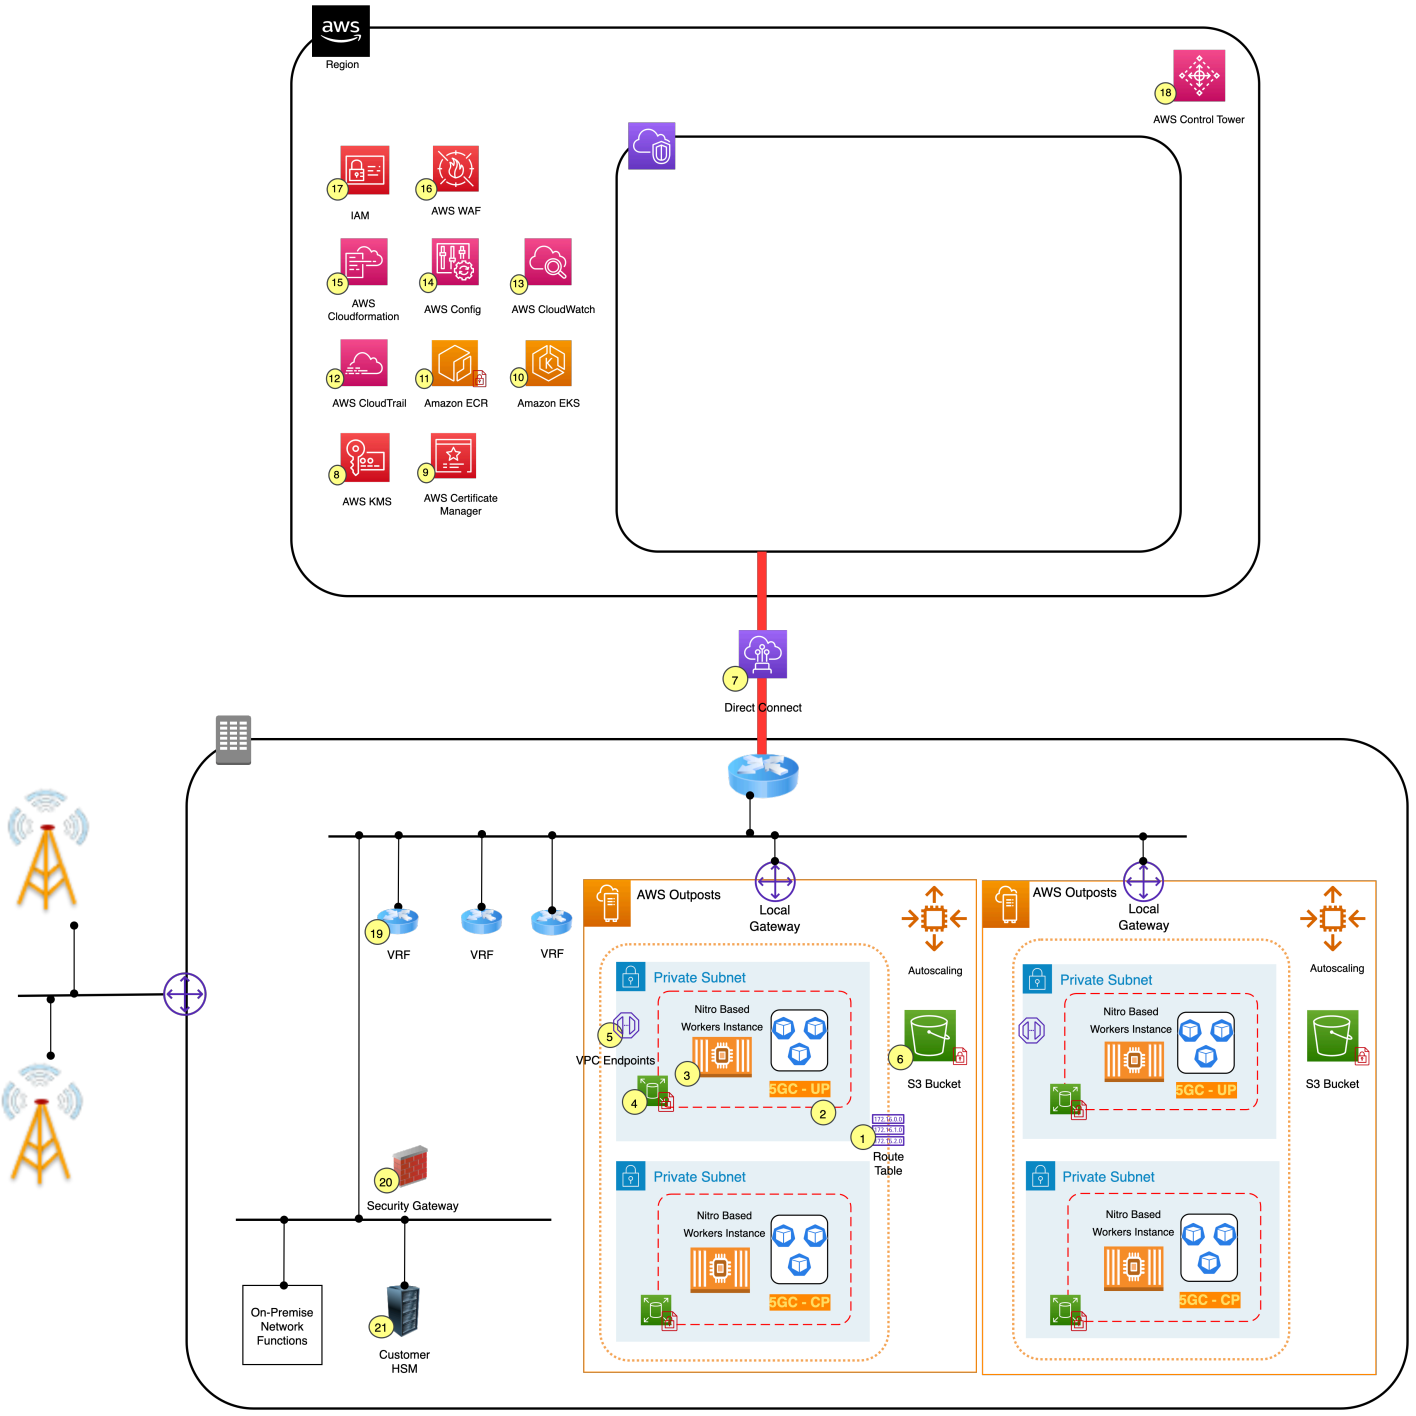

Architecture of 5G core on AWS Outposts

Security description of the example architecture of 5G core network function on AWS Outposts:

1. *VPC routing tables*. As an example, customers can direct the user plane or internet traffic to on-premises network using the AWS Outposts local gateway.

2. Traffic going in and out of the instances are filtered using security groups. In addition, there are network ACL rules that can filter traffic on a subnet level. Network ACLs are stateless firewall rules.
3. Nitro hardware-based instances.
4. Persistent data at rest stored in EBS volumes.
5. Access to AWS services that do not reside inside the VPC is through VPC endpoints.
6. Snapshots, AMIs, manifest files, or backup data can be stored in Amazon S3 storage. Data at rest is encrypted using AWS KMS, and access to data can be restricted with IAM policies.
7. AWS Direct Connect instances.
8. AWS KMS for management of encryption keys.
9. AWS Certificate Manager to manage imported SSL/TLS certificates.
10. Amazon ECR is used to store container images.
11. Amazon EKS service is used for Kubernetes-based container orchestration.
12. AWS CloudTrail helps enable governance, and supports operational and risk auditing of an AWS account.
13. Amazon CloudWatch monitors AWS resources and applications that run on AWS in near real-time.
14. AWS Config provides a detailed view of the configuration of AWS resources in an AWS account.
15. AWS CloudFormation helps set up AWS resources automatically.
16. AWS WAF helps protect application endpoints or APIs against common web exploits and bots.
17. AWS IAM helps to securely control access to AWS resources.
18. AWS Control Tower provides a simple way to set up and govern a secure, multi-account AWS environment.
19. VRF devices, virtual router, and forwarding devices are used to segregate the VPN.
20. Customer SEGs are entities on the borders of the IP security domains used for securing native IP based protocols.
21. Customer owned on-premises HSM to generate cryptographic keys for importing to AWS KMS or use with AWS KMS XKS.

Example architecture #3

An example architecture of a 5GC workload on the AWS Region, RAN CU on an AWS Local Zone, and RAN distributed unit (DU) on customer premise.

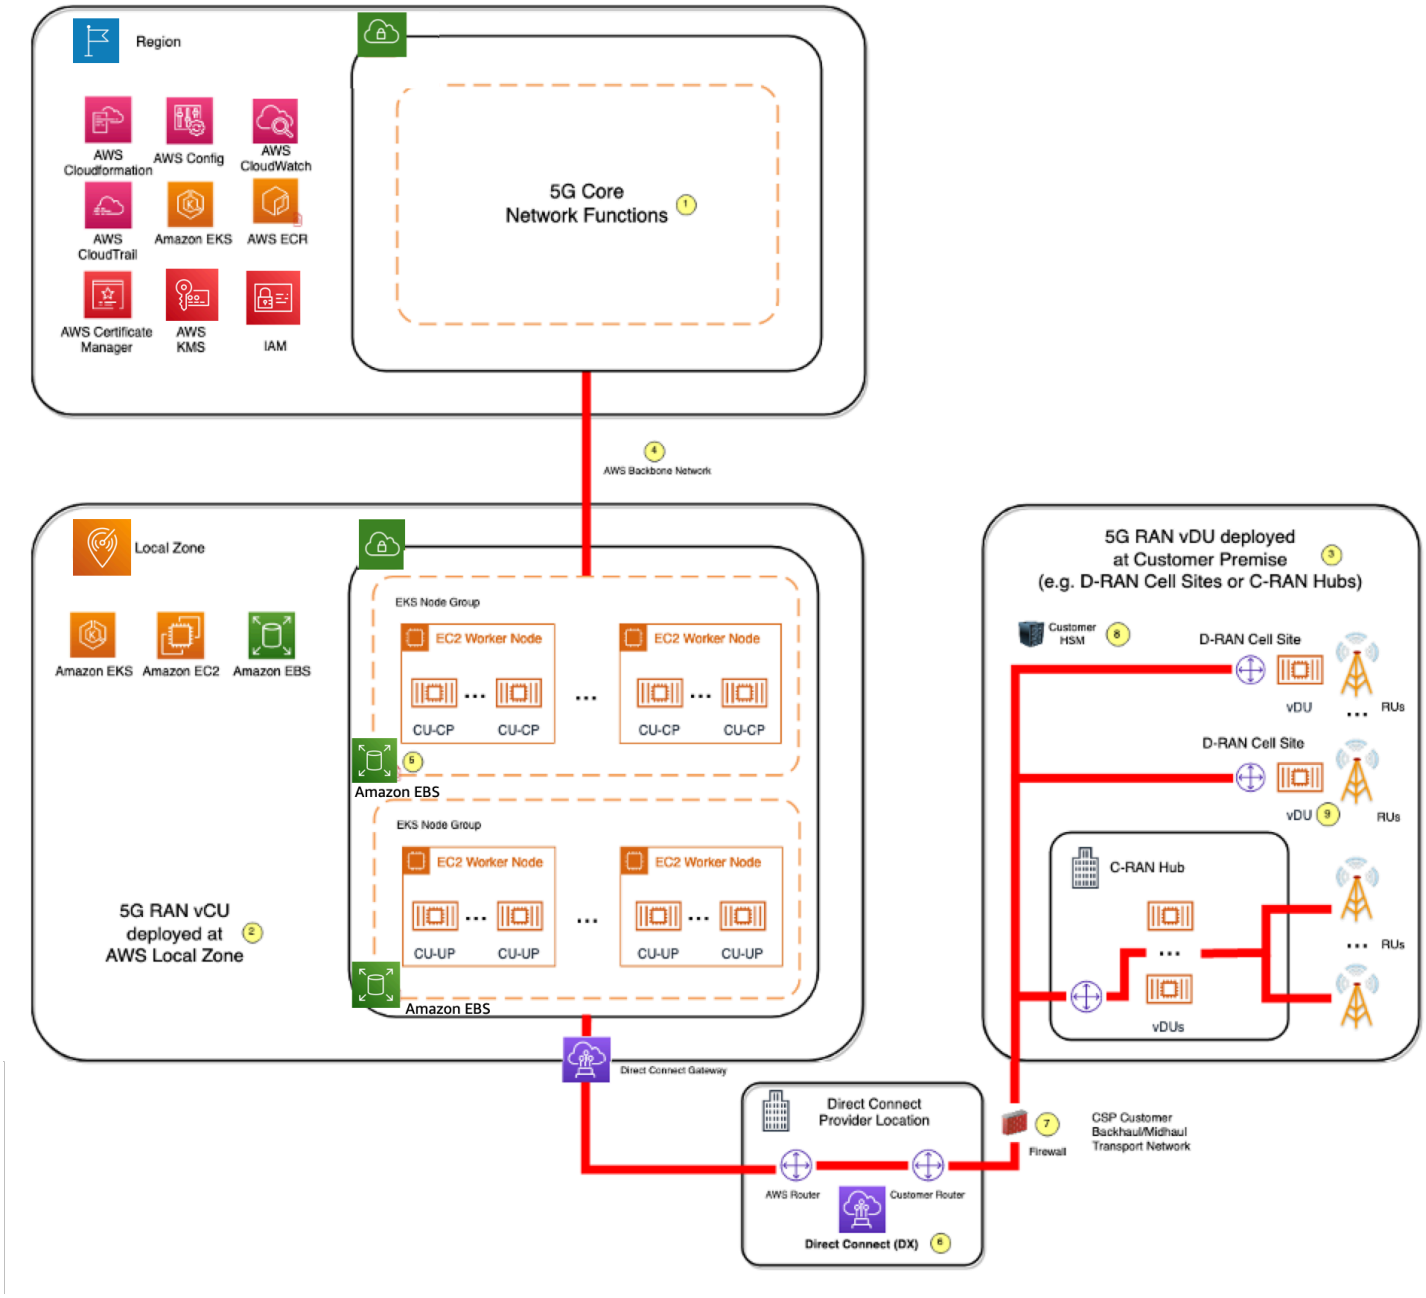

Architecture of 5G RAN on an AWS Local Zone and on-premises network

Security description of the example architecture of 5G RAN network function on AWS Local Zones and customer on-premises network:

1. Assuming 5G core network functions are deployed in the AWS Region, the security-related services and best practices described in the previous 5G Core deployment examples are also applied here.
2. 5G RAN vCU deployment at the AWS Local Zones. Due to the latency requirements of RAN network functions, the centralized unit (CU) functions of RAN need to be placed within tens of milliseconds from the end users. Therefore, AWS Local Zones are ideal edge locations to host CU function due to their low-latency access to the end users. AWS Local Zones are fully managed by AWS, and includes secure cloud infrastructure providing compute, storage, database and other select AWS services to customers.
3. 5G RAN vDU deployment at customer on-premises network: due to the ultra-low latency requirements (~100 microsecond) of the RAN DU network functions, they are typically deployed at customer on-premises locations, such as D-RAN cell sites, or C-RAN hubs.
4. The infrastructure-level connectivity between the AWS Regions and AWS Local Zones uses high-speed and secure AWS backbone networks. At the service level, you can extend a VPC from the parent Region into AWS Local Zones by creating a new subnet and assigning it to the AWS Local Zone.
5. EBS volumes are encrypted by default using Amazon EBS Encryption for data at rest and data in transition between the Local Zone and its parent Region. By default, Amazon EBS encryption uses AWS KMS and AWS-managed keys. However, customers can specify Customer Managed Keys as the default encryption key.
6. AWS Direct Connect is recommended to provide a dedicated private network connection between the customer on-premises network and AWS networks. While in transit, your network traffic remains on the AWS global network and never touches the public internet; therefore, it is more secure and provides better performance. To add an extra layer of security, you can use AWS Direct Connect connections that support MACsec to encrypt your data from your on-premises network or collocated device to your chosen AWS Direct Connect point of presence.
7. Network firewalls are typically deployed within CSP customers' transport networks to filter traffic to/from the RAN.
8. The customer-owned on-premises HSM can be used to generate cryptographic keys for importing to AWS KMS and securing an on-premises network, or for equipment purposes.
9. Given the DU functions are typically deployed at the very far edge of the telco network with limited security monitoring (for example, unmanned cell sites), additional security measures (such as disk encryption, secure boot, and so on) should be considered to protect against physical equipment theft or tampering.

Conclusion

Implementing robust security practices is crucial for CSPs to safeguard their network workloads and protect customer data. The constantly evolving threat landscape requires CSPs to proactively address potential threats and risks to maintain the confidentiality, integrity, and availability of their services.

CSPs looking to adopt the AWS Cloud because of its many benefits can take advantage of AWS infrastructure, security services, shared responsibility, and security design principles from the Security Pillar of the Well-Architected Framework. This helps CSPs to design, deploy, and operate telco workloads securely in the AWS Cloud, while adhering to requirements and proactively reducing risk.

From this whitepaper, CSPs can take away various design principles that they should look to incorporate into their workload architectures, such as zero trust, application security using threat modeling, and confidential compute to protect data throughout its lifecycle. This whitepaper outlined example architectures that incorporate these principles, and the recommendations set forth throughout serve as a starting point for CSPs to understand how they can plan and implement their target security posture in the cloud. CSPs should consider this whitepaper with their plans to run telco network workloads on AWS to get insights on how they should define their security on AWS, and how AWS can help fulfill their security needs.

Contributors

Contributors to this document include:

- Rolando Jr Hilvano, Principal Solutions Architect, Telco 5G
- Cheng Liu, Principal Solutions Architect, 5G RAN EKS-A
- Danny Cortegaca, Senior Security Solutions Architect, AWS Industries
- Scott Taggart, Principal Security Solutions Architect

Abbreviations

- ACL - Access Control List
- AMF - Access and Mobility Management Function
- AMI - Amazon Machine Image
- API - Application Programming Interface
- AEAD - Authenticated Encryption with Associated Data
- BYOK - Bring Your Own Keys
- CCoE - Cloud Center of Excellence
- CDR - call detail records
- CEC - cloud enablement center
- CI/CD - continuous integration/continuous deployment
- CNF - containerized network functions
- COTS - commercial off-the-shelf
- CSP - communications service provider
- DNS - Domain Name System
- ETL - extract transform load
- HSM - hardware security module
- HSS - home subscriber service
- HTTP - Hypertext Transfer Protocol
- IAM - identity and access management
- IdP - identity provider
- IMEI - International Mobile Equipment Identity
- IMS - IP Multimedia Core Subsystem
- IMSI - International Mobile Subscriber Identity
- IPsec - Internet Protocol Security
- ISDN - Integrated Services Digital Network
- ISV - independent software vendor
- L2 - Layer 2
- MACsec - Media Access Control Security

- MFA - multi-factor authentication
- MSISDN - Mobile Station International Subscriber Directory Number
- NATGW - Network Address Translation Gateway
- NEF - Network Exposure Function
- NF - Network Function
- NVMe – non-volatile memory express
- OU - organizational unit
- PGW - Packet Data Network Gateway
- PII - personal identifiable information
- RAN - radio access network
- RTP – Real-time Transport Protocol
- SCTP - Stream Control Transmission Protocol
- SDK - software development kit
- SIP - Session Initiation Protocol
- SLDC - Software Development Life Cycle
- TCP - Transmission Control Protocol
- TLS - Transport Layer Security
- TME - Total Memory Encryption
- UDM - Unified Data Management
- UDP - User Datagram Protocol
- UPF - User Plane Function
- VNF - Virtual Network Functions
- VPC - virtual private cloud
- VPN - virtual private network
- XKS - external key store
- XTS-AES - XEX Tweakable Block Ciphertext Stealing- Advanced Encryption Standard

Document revisions

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

Change	Description	Date
Initial publication	Whitepaper published	June 30, 2023

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.

AWS Glossary

For the latest AWS terminology, see the AWS glossary in the *AWS Glossary Reference*.