

SECURITY IN OPEN RAN

WHITE PAPER

January 2021



INTRODUCTION

Open RAN is an open radio access network (RAN) architecture standardized by the O-RAN Alliance based on 3GPP and other standards. O-RAN Alliance's RAN functional split is based on the three key tenets:

- > Decoupling of hardware and software
- > Cloud infrastructure
- > Standardized and open interfaces between the network functions

In the IT world, hardware-software decoupling happened a long time ago. This decoupling led to the emergence of software players that were experts in specific horizontal layers. The software from these players could run on any hardware providing operator customers with a variety of options. An equally rich ecosystem of hardware players emerged.

KEY TOPICS IN THIS WHITE PAPER

This white paper article focuses on the following Open RAN security aspects:

- > Next Generation RAN Architectures
- > Open RAN security based on Zero Trust Architecture
- > Secured communication between Network Functions
- > Secure Framework for RIC
- > Secure platform for Network Elements
- > Key Security Differentiators in Open RAN

TABLE OF CONTENTS

INTRODUCTION	1
TABLE OF CONTENTS	2
1. OVERVIEW	3
2. NEXT GENERATION RAN ARCHITECTURES	4
3. OPEN RAN SECURITY BASED ON ZERO TRUST ARCHITECTURE	7
4. SECURED COMMUNICATION BETWEEN NETWORK FUNCTIONS	8
4.1 Secure communication on all interfaces	8
4.2 Establishing trust based on mutual authentication	9
4.3 Trusted Certificate Authorities.....	11
5. SECURE FRAMEWORK FOR RIC	11
5.1 Security aspects of near-real-time radio intelligent controller (Near-RT RIC)	11
5.1.1 Secure Interface between Near-RT RIC and the O-CU-CP/O-CU-UP/O-DU	11
5.1.2 Conflict resolution and xApp authentication.....	11
RAN software power savings.	12
5.1.3 User identification inside the Near-RT RIC.....	13
5.2 Security aspects of Non-Real-Time Radio Intelligent Controller (Non-RT RIC).....	13
5.2.1 Secure Interface between Non-RT RIC and the O-CU-CP / O-CU-UP / O-DU	14
5.2.3 Conflict resolution between the Non-RT RIC and the O-CU-CP / O-CU-UP / O-DU.....	14
6. SECURE PLATFORM FOR NETWORK ELEMENTS	15
6.1 Secure platform for cloud native network functions	15
6.1.1 Security of a container-based application software.....	16
6.1.3 Security considerations with a cloud native hardware infrastructure	21
6.1.3.1 Secure storage of credentials and data at rest	21
6.1.3.2 Establishing software chain of trust.....	21
6.2 Secure platform for O-RU	23
7. KEY SECURITY DIFFERENTIATORS IN OPEN RAN	24
8. CONCLUSION	25
APPENDIX	26
References.....	26
Acronyms	26
About Mavenir	27

1. OVERVIEW

Virtualization technologies have helped enterprises reduce their TCO through efficient use of compute resources, removal of hardware silos and increased automation. To deliver 5G services, operators need a virtualized network capable of scaling services based on policy-driven service selections for subscribers. Cloud native architecture allows deployment of network functions (NFs) as a cluster of containerized microservices, where each microservice can be deployed, scaled, and upgraded independently. Instead of scaling the whole application, only the required component within the NF is scaled.

Open interfaces between various network functions allow best of breed equipment to be used in networks enabling operators to distinguish themselves from competition by using bespoke network functions, as needed.

In this paper it is demonstrated how, by adopting a zero-trust security framework, an Open RAN architecture provides a path to a more secure open networks and open interfaces over what exists today. Despite misconceptions, open interfaces, defined in the O-RAN technical specifications, provide increased independent visibility and the opportunity for an overall enhanced and more secure system.

5G and Open RAN enable new capabilities and control points that allow suppliers, test equipment manufacturers, wireless carriers, and network operators to assess, mitigate and manage security risks efficiently. This paper details how O-RAN enables operators with full visibility and control of their network's end-to-end security.

There is a vast cloud industry solving security issues, and cloud RAN network functions are similar to other cloud network functions, with similar security requirements and solutions. Cloud architecture ensures resilience, scalability and segmentation and the introduction of features such as AI/ML and Multi-Access Edge Computing (MEC). For example, leveraging MEC, allows collection and processing of sensor traffic at a factory to shift DDoS detection and mitigation to the edge of the network where incidents at the edge can be isolated from the rest of the network. Microsegmentation, containerization, virtualization, and network slicing provide enhanced security and isolation from the hardware up. The security measures are designed into the system rather than being bolted on afterwards as in traditional systems.

2. NEXT GENERATION RAN ARCHITECTURES

3GPP [1] has defined the following architecture for 5G NR gNB as shown in Figure 1.

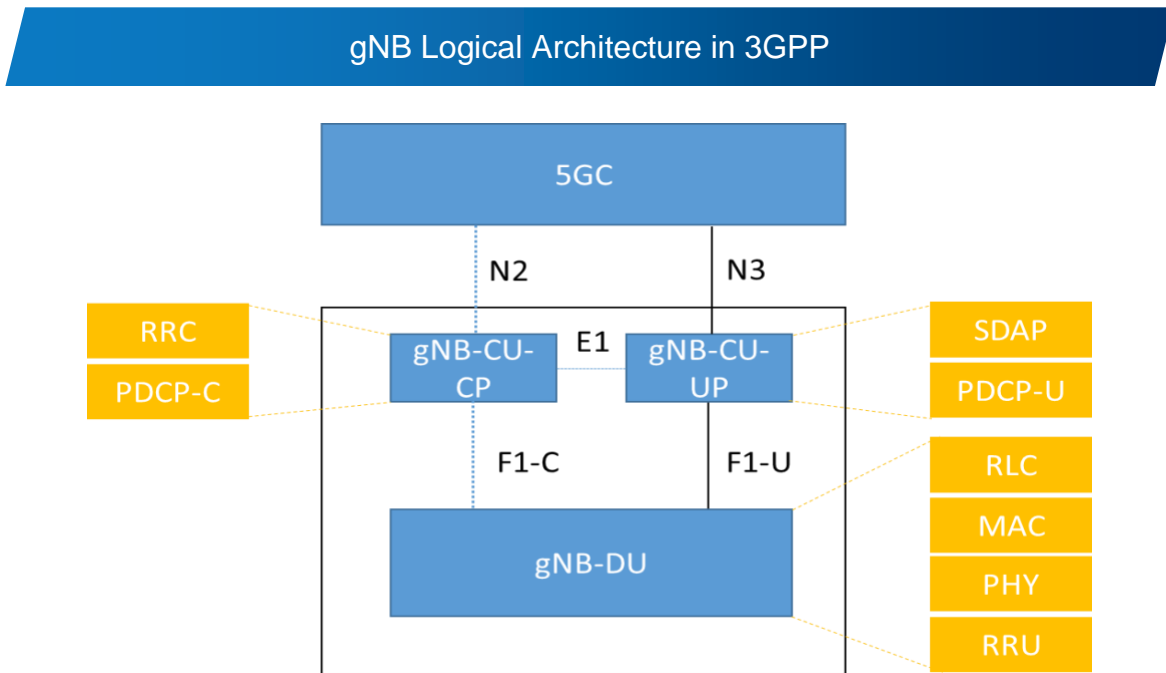


Figure 1

gNB is split into two logical functions called CU (Centralized Unit) and DU (Distributed Unit) as shown in Figure 1 and these two entities are connected by F1-C and F1-U interfaces as defined in 3GPP TS 38.473[2]. It may be noted that the 3GPP architecture does not specify the Remote Radio Unit (RRU) i.e. the interface between PHY and RF layers is left to vendor implementation.

O-RAN Alliance, a group of leading vendors and operators defining Open RAN specifications, further disaggregate CU and DU network functions [3] as defined by 3GPP that are interconnected over open, standardized, secure interfaces as shown in Figure 2.

gNB Logical Architecture in O-RAN

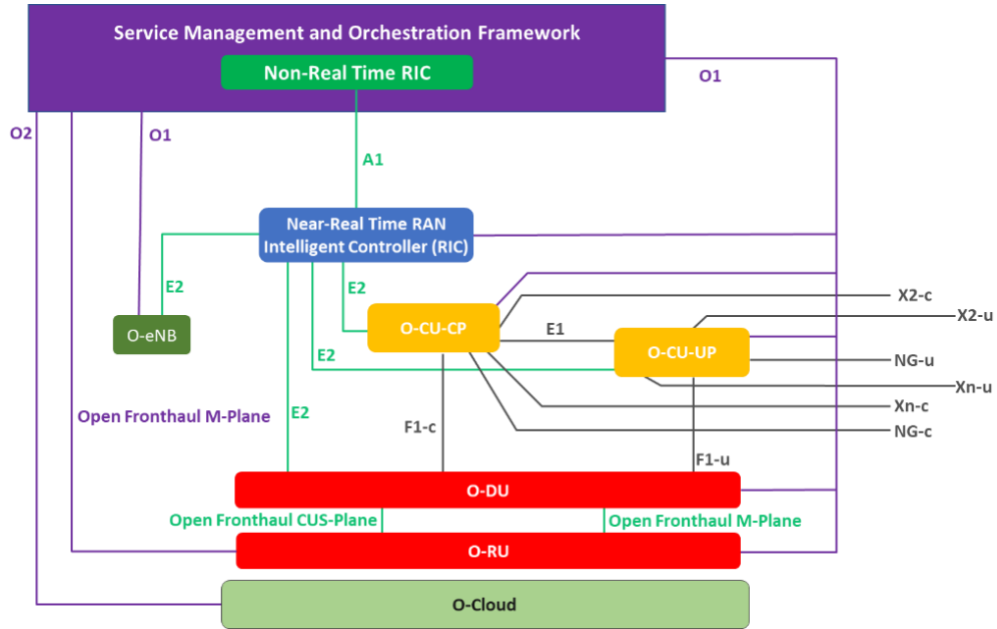


Figure 2

Figure 3 shows the functional and interface split between 3GPP and O-RAN. The O-RAN Alliance adds new interfaces and functions beyond 3GPP's 5G RAN architecture.

Interfaces and Functions split between O-RAN and 3GPP

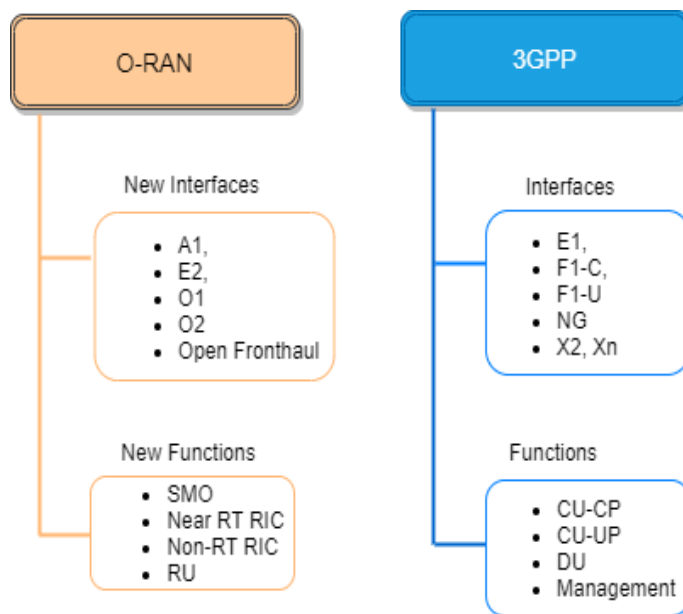


Figure 3

Since O-RAN Alliance builds on 3GPP's 5G NR architecture, it benefits from 3GPP's advanced security features introduced for 5G [4] including:

- > Enhanced user identity privacy i.e., Subscription Concealed Identifier (SUCI)
- > Full protection of control/user plane traffic between the UE and gNB (encryption and integrity protection) over the air interface
- > Full protection of gNB interfaces including the E1 interface between CU-CP and CU-UP and the F1 interface between CU and DU
- > Enhanced home network control (authentication)
- > Additional security for network slices based on SLA

3. OPEN RAN SECURITY BASED ON ZERO TRUST ARCHITECTURE

Rooted in the principle of “never trust, always verify,” Zero Trust is designed to protect modern digital environments by leveraging network segmentation, preventing lateral movement, providing Layer 7 threat prevention, and simplifying granular user-access control.

A zero trust architecture (ZTA) is a cybersecurity architecture that is based on zero trust principles and designed to prevent data breaches and limit internal lateral movement. The following is the relevant text from NIST publication 800-207 - ‘Zero Trust Architecture’ [5]-

A “zero trust” (ZT) approach to cybersecurity is primarily focused on data and service protection but can and should be expanded to include all enterprise assets (devices, infrastructure components, applications, virtual and cloud components) and subjects (end users, applications and other nonhuman entities that request information from resources).

In this new paradigm, an enterprise must assume no implicit trust and continually analyze and evaluate the risks to its assets and business functions and then enact protections to mitigate these risks. In zero trust, these protections usually involve minimizing access to resources (such as data and compute resources and applications/services) to only those subjects and assets identified as needing access as well as continually authenticating and authorizing the identity and security posture of each access request.

Support of a zero-trust architecture requires each O-RAN component to comply with established functionalities and protections. O-RAN Alliance [6] has identified several guiding principles for its ongoing work, including:

1. Support integration with an external identity, credential and access management system (ICAM) using industry standard protocols
2. Require authentication and authorization on all access
3. Support role-based access control (RBAC)
4. Implement confidentiality on connections between O-RAN and external components
5. Implement integrity checking on connections between O-RAN and external components
6. Support encryption of data at rest
7. Support replay prevention
8. Implement security log generation and collection to an external security information and event management (SIEM)

Open RAN security is built on the following tenets:

1. Secured communication between Network Functions
2. Secure framework for the Radio Intelligent Controller (RIC)
3. Secured platform for hosting the Network Functions

The analysis in the following sections assumes a cloud native Open RAN network with Network Functions modeled as containerized microservices.

4. SECURED COMMUNICATION BETWEEN NETWORK FUNCTIONS

This section explores following areas that relate to providing secure communication between all Network Functions in Open RAN.

- a. Secure communication on all interfaces
- b. Ensuring trust based authentication of communicating endpoints
- c. Trusted Certificate Authorities for Identity Provisioning

4.1 Secure communication on all interfaces

O-RAN Alliance specifies an open and secure architecture that includes secure interfaces between all its components. Communications exchanged on these interfaces are cryptographically protected for encryption, integrity protection and replay protection.

Figure 4 depicts the 5G RAN network security architecture.

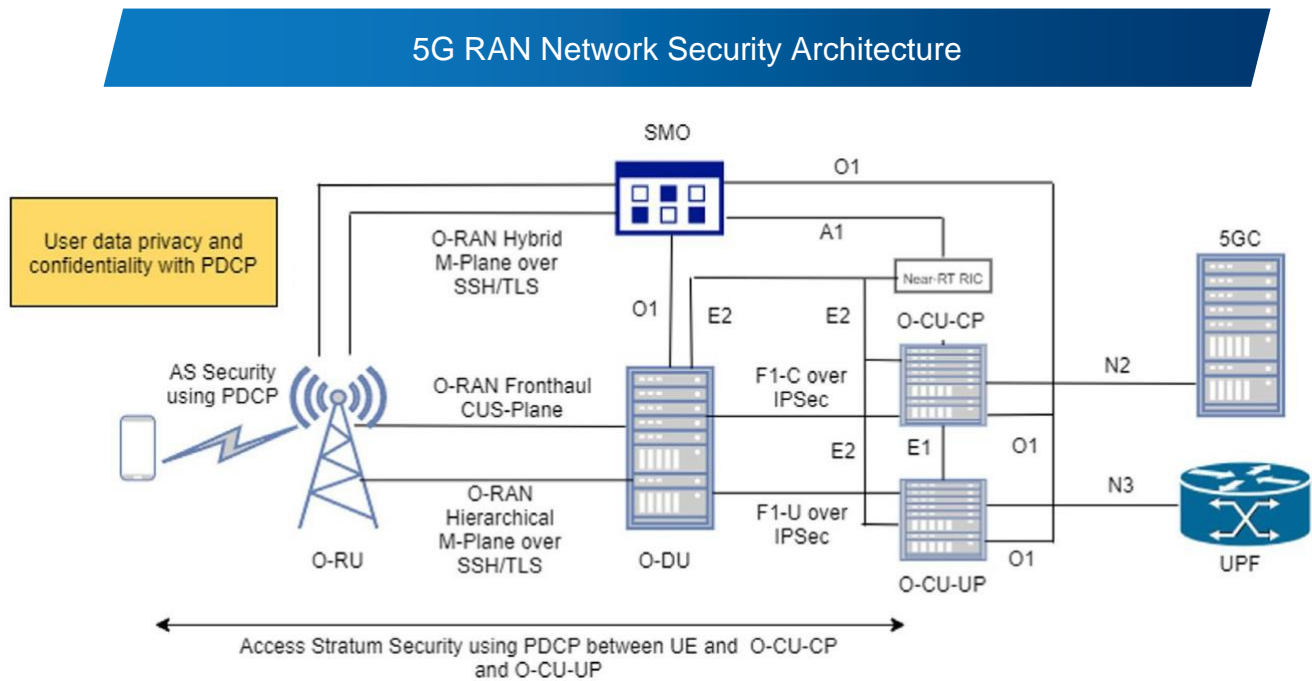


Figure 4

The following table summarizes the protection mechanism used for each interface in an O-RAN based network.

Interface	Between nodes	Security mechanism	Specified by
E1	O-CU-CP and O-CU-UP	NDS/IP (IPSec) or DTLS	3GPP
Xn	Source gNB and Target gNB	NDS/IP (IPSec) or DTLS	3GPP
Backhaul	O-CU-CP and 5GC (N2) O-CU-UP and 5GC (N3)	NDS/IP (IPSec) or DTLS	3GPP
Midhaul (F1)	O-CU-CP and O-DU (F1-C) O-CU-UP and O-DU (F1-U)	NDS/IP (IPSec) or DTLS	3GPP
Open Fronthaul (M-Plane)	O-RU and O-DU/SMO	SSHv2, TLS	O-RAN WG4
Open Fronthaul (CUS-Plane)	O-DU and O-RU	Work in progress(Dec 2020)	O-RAN WG1 STG
O1	SMO and O-RAN Managed elements	Work in progress(Dec 2020)	O-RAN WG1 STG
E2	Near-RT RIC (xAPPs) and O-CU-CP	Work planned(1Q21)	O-RAN WG1 STG
A1	Near-RT RIC and Non-RT RIC	Work planned(1Q21)	O-RAN WG1 STG
O2	SMO and O-Cloud	Work planned(2Q21)	O-RAN WG1 STG

It should be noted that several O-RAN Alliance specifications are still on-going and accordingly security work is happening in parallel. For protection of the CUS-Plane messages [7] on Open Fronthaul LLS interface, O-RAN Alliance is currently in the process of determining all the threats and vulnerabilities, and their impact on the CUS-Plane. O-RAN Alliance plans to complete the analysis and specify security procedures to protect CUS-Plane messages by March 2021.

4.2 Establishing trust based on mutual authentication

Mutual authentication is used for authenticating two entities with each other and setting up a secure encrypted connection between them. Mutual authentication prevents introduction of rogue NFs or xAPPs in the network.

Operator X.509 certificates are used for mutual authentication while establishing secure connections using IPsec and TLS protocols.

All network elements in an Open RAN, i.e. O-CU-CP, O-CU-UP O-DU and O-RU, support X.509 certificate- based authentication and related features such as auto-enrollment and auto-re-enrollment with an operator Certificate Authority (CA) server using a protocol such Enrollment over Secure Transport (EST) or 3GPP- specified CMPv2

The xAPPs in the Near-RT RIC are securely on-boarded like any other microservice and the O-RAN Alliance is expected to use CA signed X.509 certificates to authenticate before communicating over the E2 interface.

Figure 5 illustrates an example flow of how certificate-based authentication is used to authenticate an O- CU, O-DU and O-RU during certificate enrollment with a CA server.

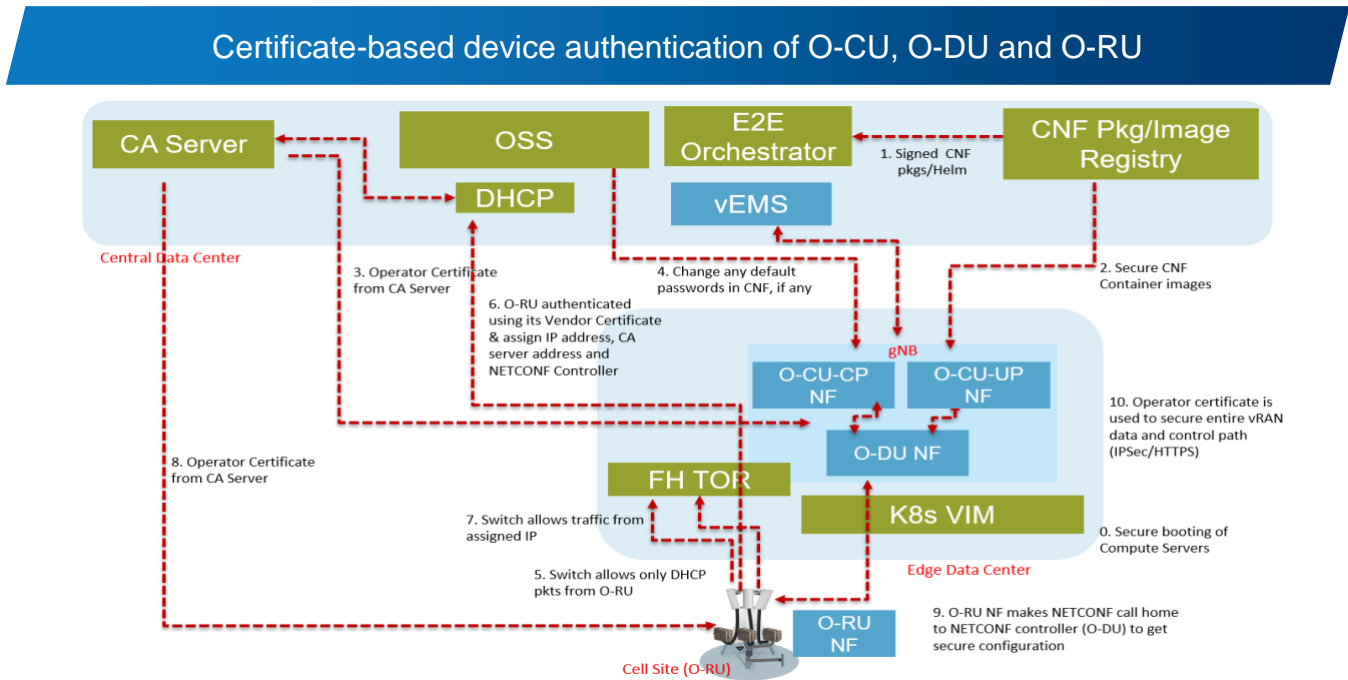


Figure 5

Step 1-2: When the O-RU powers on, the O-CU-CP, O-CU-UP and O-DU instances that are allocated to serve that O-RU are instantiated by the orchestrator, if not already instantiated.

Step 3: an O-CU-CP, O-CU-UP and O-DU performs EST or a CMPv2-based certificate enrollment procedure in compliance with 3GPP with the CA server to obtain an operator certificate. The operator certificate is used for subsequent authentication when establishing an IPSec or a TLS connection.

Step 4: necessary OAM actions are performed on the O-CU, if any, including changing of default passwords.

Steps 5 thru 9 are executed as part of the O-RU power-on sequence. Key security related steps are explained below:

- > The O-RU obtains its IP address, the EMS or OSS address from a DHCP server using one of the DHCP options specified in O-RAN M-Plane specification section 3.1.1 and 3.1.4 [8].

- > The O-RU performs certificate enrollment procedure with the CA server to obtain an operator
- > The O-RU shall notify the EMS or OSS with a NETCONF call home. O-RU's operator certificate is used to authenticate with the EMS. OSS / EMS shall configure the O-RU with the secondary NETCONF controller's address (i.e. the address of the O-DU).
- > The O-RU shall notify the O-DU with a NETCONF call home to securely obtain O-RU's configuration. O-RU's operator certificate is used to authenticate with the O-DU.

4.3 Trusted Certificate Authorities

It is recommended that the certificate authorities (CA) should be audited under the AICPA/CICA WebTrust Program for Certification Authorities.

This promotes confidence and trust in the CA servers used in Open RAN for authenticating network elements

5. SECURE FRAMEWORK FOR RIC

5.1 Security aspects of near-real-time radio intelligent controller (Near-RT RIC)

The Near-RT RIC is an SDN component that contains 3rd party extensible microservices (called xApps) that perform selected radio resource management (RRM) services for the NFs that were traditionally managed inside the gNB. The Near-RT RIC interfaces with the O-CU-CP, O-CU-UP and the O-DU via the O-RAN standardized open E2 interface. The Near-RT RIC also interfaces with the Non-RT RIC and the service management and orchestration framework via the A1 and O1 interfaces.

The key security aspects of the Near-RT RIC include:

- > Secure E2 Interface between the Near-RT RIC and the O-CU-CP / O-CU-UP / O-DU
- > Conflict resolution and xApp authentication
- > User identification inside the Near-RT RIC

5.1.1 Secure Interface between Near-RT RIC and the O-CU-CP/O-CU-UP/O-DU

Interface security is explained in § 4.2

5.1.2 Conflict resolution and xApp authentication

The conflict resolution among the xApps is not necessarily a security issue but can lead to vulnerabilities if not handled properly.

While the xApps in the Near-RT RIC initiate the RIC subscription procedure with the E2 nodes, the subscription manager in the Near-RT RIC platform, enforces the subscription policies and

keeps track of the subscriptions initiated by the xApps and the RAN functions, and event triggers associated with those subscriptions. The subscription manager can resolve signaling conflicts among the xApps by one or more of the following means:

- > The subscription manager will not allow more than one xApp to subscribe to the same NF based on the same event trigger.
- > If more than one xApp subscribes to the same NF and gets the same indication messages from the E2 node, then the subscription manager can allow them to simultaneously control the NF of the E2 node, as long as they do not optimize the same or closely inter-dependent parameters pertaining to the NF.
- > If more than one xApp subscribes to the same NF and gets the same indication messages from the E2 node and if they optimize closely inter-dependent parameters, then the subscription manager can allow them to simultaneously control and optimize those parameters by using locks and backoff timers to retain mutual exclusivity.

RAN software power savings.

As the disaggregated RAN compute resources move to data centers, the power efficiency can take advantage of the global data center power optimization trends. The data center power consumption has increased by 6% since 2010 but at same time the amount of compute in the data center has increased by 550%¹. With centralized baseband processing in the cloud, it is much easier to pool resources taking into account the workload variations across cell sites and time of day and implement usage-based power savings that can be adjusted dynamically. A NGMN study in Europe shows that 80% of a wireless network carry only 20% of the traffic.² and pooling across sites could potentially reduce DU/CU capacity requirements with significant compute and power savings. With scalability and demand-based usage, processors (CPUs or GPUs) that are processing radio software can also run other applications during non-peak times. This is not possible with proprietary baseband systems using dedicated, non-reusable hardware.

Authentication aspects of xAPP is explained in § 4.2

1 <https://www.datacenters.com/news/data-center-power-optimization-increase-efficiency-with-a-data-center-audit>

2 https://www.ngmn.org/wp-content/uploads/NGMN_RAN_EV_D2_Further_Study_on_Critical_C-RAN_Technologies_v1.0.pdf

5.1.3 User identification inside the Near-RT RIC

Maintaining privacy of the users is of utmost importance inside the RIC. ORAN WG3 is working on the UE identification inside the Near-RT RIC that can be addressed by a combination of 3GPP-defined Trace ID, 3GPP-defined RAN UE ID, temporary RAN network interface-specific UE IDs, and by correlating these IEs with one another. Typically, it is ideal for the Near-RT RIC to maintain persistence of UE identification for near-RT granularities, ranging from 10 ms to 1 s. The xApps are not exposed to UE permanent ID. Invalidation of the temporary IDs in the RIC when they are released in RAN nodes will be handled via normal E2 communication. In neither case is this a UE privacy issue or a DoS attack threat.

5.2 Security aspects of Non-Real-Time Radio Intelligent Controller (Non-RT RIC)

The Non-RT RIC is a component in an O-RAN system for non-real-time control of the RAN through declarative policies and objective intents. This is illustrated in Figure 6 below.

1. The Non-RT RIC is deployed in a service management and orchestration framework (SMO) and provides declarative policy guidance for cell-level optimization by providing the optimal configuration values for cell parameters over the O1 interface.
2. The Non-RT RIC also sends declarative policies for UE-level optimization to the Near-RT RIC via the A1 interface.
3. The Near-RT RIC then translates the recommended declarative policy from the Non-RT RIC over A1 interface into per-UE control and imperative policy over the E2 interface.
4. The Non-RT RIC develops ML/AI-driven models for policy guidance and non-RT optimization as rApp microservices. These rApps interface with the xApps over the A1 interface to optimize a set of procedures and functions in the underlying RAN.

Non-Real-Time RIC declarative policies and objective intents

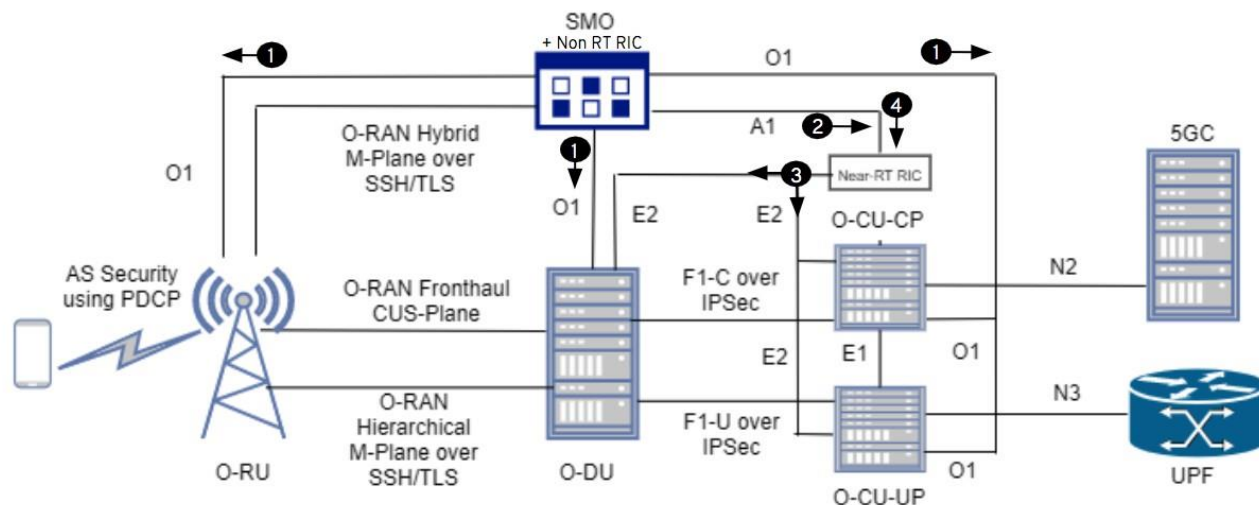


Figure 6

The key security aspects of the Non-RT RIC are the following:

- > Secure interface between Non-RT RIC and the O-CU-CP / O-CU-UP / O-DU
- > Conflict resolution between the Non-RT RIC and the O-CU-CP / O-CU-UP / O-DU

5.2.1 Secure Interface between Non-RT RIC and the O-CU-CP / O-CU-UP / O-DU

Interface security is explained in § 4.2

5.2.3 Conflict resolution between the Non-RT RIC and the O-CU-CP / O-CU-UP / O-DU

Usually, a conflict in RRM arises when the RAN uses policies and objective intents different from the Non-RT RIC to manage the underlying RAN nodes such as the O-CU. This may be the source of rApps causing signaling conflicts with the functioning of the underlying RAN nodes. However, using the RIC subscription policies, mutual exclusivity can be enforced causing the subscribed procedures from the RAN to be managed by the Near-RT RIC, without causing signaling conflicts

6. SECURE PLATFORM FOR NETWORK ELEMENTS

O-RAN Alliance RAN architecture is built on a fully cloud native architecture – the same cloud architecture that is the bedrock of today’s internet and public cloud. The cloud native network functions in the O-RAN network viz. O-CU-CP, O-CU-UP, O-DU, Near-RT RIC and Non-RT RIC, are hosted on a cloud native platform, very similar to the cloud native platform used in the cloud computing industry. The O-RU is a PNF and thus hosted on a non-virtualized platform. In the following sections we take a holistic look at security aspects of these platforms.

6.1 Secure platform for cloud native network functions

The O-RAN architecture uses a cloud-native platform to host O-CU-CP, O-CU-UP, O-DU, Near-RT RIC and Non-RIT RIC network functions. Figure 6 shows a typical cloud native platform with three distinct layers:

1. Container-based application software
2. Cloud native software stack comprising an immutable OS, Kubernetes and Container runtime
3. Cloud native hardware infrastructure

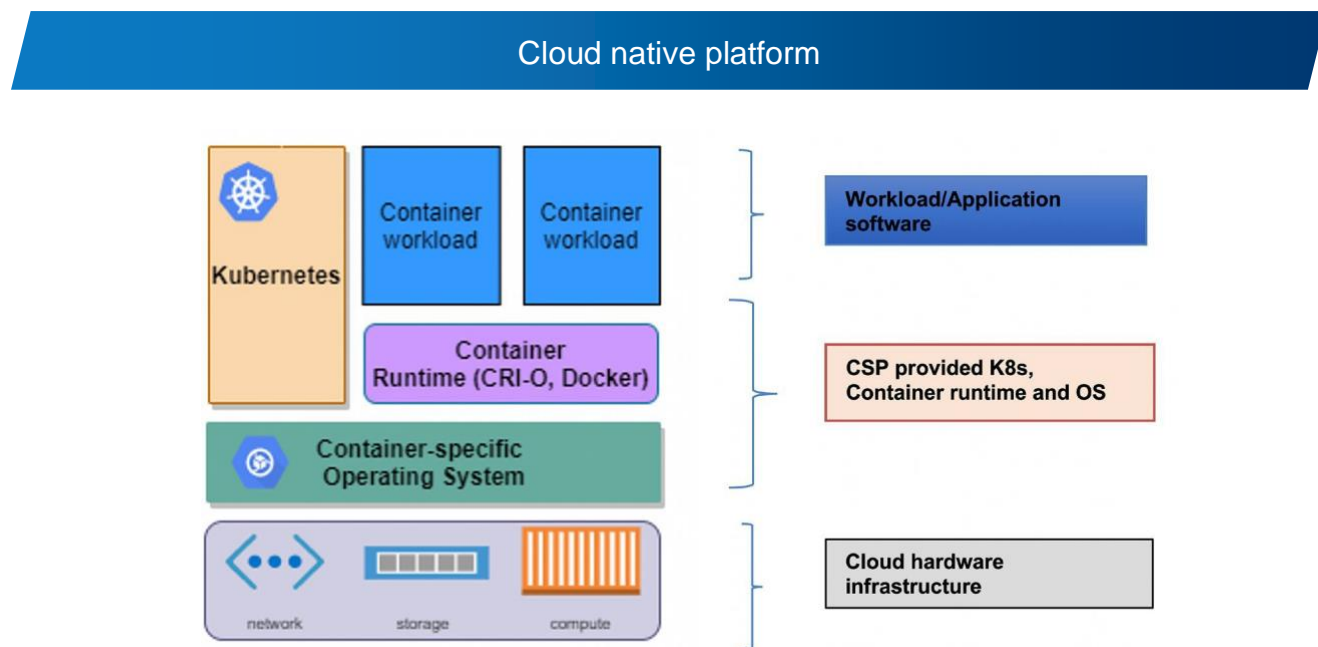


Figure 7

The following sections look at security features of each of the three layers that make up a cloud native platform

6.1.1 Security of a container-based application software

A workload is an application or a service deployed on the cloud. Containers offer a packaging infrastructure in which applications and dependent libraries are abstracted from the environment in which they actually run.

Containers are generally perceived to offer less security than virtual machines. But it's worth noting that containers have been in use in the IT industry to build applications such as for banking which are no less critical than telecom applications in terms of security requirements, and the industry has evolved itself in automating its security and establishing best practices.

The following industry standard practices are used in Open RAN to ensure security of the container-based application software:

- a) Secure software development based on “secure by design” principles
- b) Automating security testing based on DevSecOps
- c) Vulnerability management in Open Source and 3rd party libraries

Secure software development based on “secure by design” principles


A software development life cycle (SDLC) is a framework for the process of building an application from inception to decommission. In the past, organizations usually performed security-related activities only as part of testing—at the end of the SDLC. As a result of this late-in-the-game technique, they wouldn't find bugs, flaws, and other vulnerabilities until they were far more expensive and time-consuming to fix. Worse yet, they wouldn't find any security vulnerabilities at all.

A secure SDLC involves integrating security testing and other security-related activities into an existing development process. Figure 7 shows how a standard SDLC process is augmented with security practices at every stage of software development.

SDLC Process



Secure SDLC Process




Using a secure SDLC process for the workloads deployed in a O-RAN network such as xAPPs in Near-RT RIC, O-CU-CP and O-CU-UP and O-DU microservices, ensures early detection of flaws in the system, awareness of security considerations by all stakeholders involved in designing, development, testing and deployment of containers, and overall reduction of intrinsic business risks for the organization.

Automating security testing based on DevSecOps

Since the beginning of modern computing, security testing has largely been an independent activity from software development. Security focused QA professionals performed testing during the testing phase.

A DevSecOps approach to the container development lifecycle ensures that security is built-in at every stage of the CI/CD pipeline.

Figure 9: Automated security practices based on DevSecOps



The philosophy behind DevSecOps is to begin security testing early in the SDLC. DevSecOps integrates various security controls into the DevOps workflow such as secure coding analysis using static application security testing (SAST), automated unit, functional and integration testing. This enables developers to fix security issues in their code in near real time rather than waiting until the end of the SDLC.

O-RAN Alliance architecture software takes advantage of the advancements in ‘security automation’ and trend in cloud computing towards “shift left.” This ensures that workloads run in the O-RAN network are validated securely (during build/deployment phase) and risk-based timely actions are taken when vulnerabilities are found before they are deployed in operator network.

Vulnerability management of open source and 3rd party libraries

Open source libraries and open source software enable developers to meet the demands of today’s accelerated development timelines. However, they can also open up the platform to attacks due to unaddressed vulnerabilities in the software.

Software component analysis (SCA) is an open source management tool that helps in identifying potential areas of risk from the use of third-party and open-source software. SCA software automatically scans all open-source components, creates an accurate bill of materials (BOM), checks for policy and license compliance, security risks, and version updates. SCA software also provides insights for remedying identified vulnerabilities, usually within the reports generated after a scan.

Specialized container image scanning tools provide automated vulnerability management for containers by identifying and providing remediation paths for all the vulnerabilities in the image. These tools are integrated into the CI/CD pipeline and provide continuous assessment of the container image

Use of software component analysis tools in an O-RAN network allows for deployment of an advanced vulnerability management process that includes automatic tracking, analysis of an application's open source components, identification of component vulnerabilities, and tool-based vulnerability remediation.

Compliance with supply chain risk management requirements from NIST SCRM and CISA ICT SCRM.

6.1.2 Security of cloud native software infrastructure

A cloud native software infrastructure includes the following:

- a) Container-specific operating system – lightweight and purpose-built OS
- b) Container runtime – software that executes containers and manages container images on a node
- c) Container orchestration – software that automates the deployment, management, scaling and networking of containers

Container-specific OS

The cloud native software infrastructure relies, in line with the NIST SP 800-190 recommendations [9], on a host OS built and configured for the sole purpose of running containerized applications instead of general-purpose applications reducing the OS attack surface. In addition, the container-specific OS follows the immutability infrastructure paradigm by preventing any additional individual software package installation protecting against viruses and malware; the entire OS being managed as a single entity. Any additional feature has to be installed as a container. The OS implements strong isolation and mandatory access control (MAC) mechanisms such as SELinux to limit what a container can do and thus protecting the OS from the containers and the containers from each other. The OS also supports inbuilt Linux features such as control groups (cgroups) and namespaces that provide an isolated environment for the application running inside the container. The OS also supports disk encryption including the root partition by leveraging linux unified key setup (LUKS) encryption.

Container runtime

The cloud native software infrastructure includes a lightweight, Kubernetes-specific OCI-compliant container runtime versioned with Kubernetes such as CRI-O to reduce the risk of vulnerabilities.

The cloud native software infrastructure (container -specific OS, container runtime, disk ...) must support running in FIPS mode by using FIPS 140-2 validated cryptography.

Native security with Kubernetes

Kubernetes provides several built-in security capabilities to secure the container environment including network security, resource isolation, access control, logging and auditing. Some of the common Kubernetes built-in controls that help in tightening security include:

- a) Role based access control (RBAC) Use of RBAC in the cluster provides a framework for implementing the principle of least privilege for humans and applications accessing the Kubernetes API.
- b) Configure the security context for pods to limit their capabilities
- c) Pod security policy sets defaults for how workloads are allowed to run in the cluster. These controls can eliminate entire classes of attacks that depend on privileged access.
- d) Use Kubernetes network policies to control traffic between pods and clusters.
- e) Kubernetes' network policies allow control of network access into and out of the containerized applications. In addition to this feature, software-based firewalls may be deployed to control container to container communication within or across different clusters.
- f) Use namespaces to isolate sensitive workloads and create security boundaries – separating workloads into namespaces can help contain attacks and limit the impact of mistakes or destructive actions by authorized users.
- g) Assess the container privileges – Adhering to the principle of least privilege and provide the minimum privileges and capabilities that would allow the container to perform its intended function.
- h) Use mutual Transport Layer Security (TLS) for all inter cluster and intra cluster communications.
- i) Capability to encrypt the etcd datastore to protect infrastructure and application secrets or to support integration with external vaults.

Leveraging Kubernetes operators for security

Kubernetes operators are software extensions to Kubernetes that make use of custom resources to manage services and their components in an automated way. These operators can be leveraged by the cloud native software platform for specific security purposes:

- > Hardware management operators to restrict the need for applications of elevated privileges
- > Compliance operators to continuously monitor the compliance of the cluster
- > File integrity monitoring operators to detect any attacks impacting the platform integrity
- > Platform management operators to fight configuration drift and enforce a secure configuration by eliminating human errors
- > Audit and log operators to manage the audit configuration and the log forwarding to a SIEM

A cloud native-based O-RAN network can leverage native security controls in container runtime and container orchestration platforms such as Kubernetes, to provide defense in depth security for the containerized workload that they host.

Secure configuration of the cloud infrastructure based on industry benchmarks

The cloud infrastructure is configured based on industry best practices such as CIS benchmarks for operating system, Docker and Kubernetes, and Network Equipment Security Assurance Scheme (NESAS) jointly defined by 3GPP and GSMA provides a consistent framework and common external audit program for multiple vendors and operators. This ensures that appropriate security controls are put-in-place in the platform, thus reducing its attack surface.

Some of the common security controls include disabling unused ports and unused service, principle of least privileges (PLoP) for workloads, protecting data in storage, user access control using RBAC, etc.

All virtualized platforms in an O-RAN network are hardened as per 3GPP's security assurance specifications

[10] and other well-known industry benchmarks such as those from CIS [11]. This ensures that security controls are implemented at every layer of the platform thus reducing the platform's attack surface.

Detecting and remediating configuration errors with cloud security posture management

Misconfiguration is the #1 cause of cloud-based data breaches. A mechanism is needed to make sure the configuration of the deployed cloud resources is correct and secure on day one, and that they stay that way on day two and beyond. This is referred to as cloud security posture management (CSPM).

The cloud industry has used CSPM security tools to continuously monitor cloud environments for detection of cloud misconfiguration vulnerabilities that can lead to compliance violations and data breaches.

With the adoption of a cloud native architecture in O-RAN based networks, an operator now has the means to deploy advanced CSPM tools to guard against natural “drift” of on network configuration and reduce the potential for attacks.

Commercial cloud native hybrid platform

Standardizing on a commercial cloud native hybrid platform enables the operator with the following security benefits:

- > A Kubernetes-certified platform with the flexibility to run securely on-prem or in a virtual private cloud, supporting O-RAN topology variations from the SMO, RICs, CUs, and DUs with zero-touch provisioning,
- > Extended software lifecycle with dynamic updates that address new CVEs and optimizations over time into disconnected environments,
- > Support for multi-tenancy so that multi-vendor software can be securely hosted in the same cluster,
- > Support for infrastructure compliance scanning (OpenSCAP) and remediation,
- > A container registry with vulnerability scanning to eliminate vulnerabilities on O-RAN platforms (e.g Near Real-Time RIC) and associated xApps and rApps

6.1.3 Security considerations with a cloud native hardware infrastructure

O-RAN enables decoupling of hardware and software, allowing for a platform to be built from different vendors.

6.1.3.1 Secure storage of credentials and data at rest

It is recommended that O-RAN hardware comes with a hardware-based security module like TPM to manage, generate, and securely store cryptographic keys. Hardware-based security modules are also

meant to provide a hardware root of trust to enable secure computing by providing a secure key storage enclave with minimal cryptographic functions primarily in the signing and signature verification space.

The data at rest must be encrypted using keys generated from hardware-based security modules.

6.1.3.2 Establishing software chain of trust

Zero-trust cannot be achieved without the full participation of all the elements in the trust chain for a network.

Figure 9 illustrates key aspects of establishing chain of trust when adhering to zero-trust in digital systems.

Trusted hardware

The hardware is built with a tamper resistant “hardware root of trust” device that provides a secure environment for storing cryptographic keys and for attestation of certificates and all the software running on that hardware. The device will expose a simple user interface for the application to use when it needs to use the device for storing keys, retrieving certificates etc.

Trusted software

Software signing is enforced at all software layers including the firmware, cloud native software stack and container workloads at time of deployment, as well as authenticated version upgrades to make it more difficult to introduce malicious software into operator- controlled elements.

Establishing end-to-end chain of trust with secure boot

Secure boot requires that every boot up is starting from a piece of software that cannot be updated in the field. This piece of software is referred to as Core Root of Trust for Measurement (CRTM).

Thereafter, during the boot process every software program in the platform will be integrity verified before its execution by the software at the lower layer. This establishes an end-to-end software chain of trust. The trust anchor for the software integrity verification is software signing certificate.

In the O-RAN network, it is recommended to use secure boot based on hardware root of trust and software signing to establish an end-to-end chain of trust.

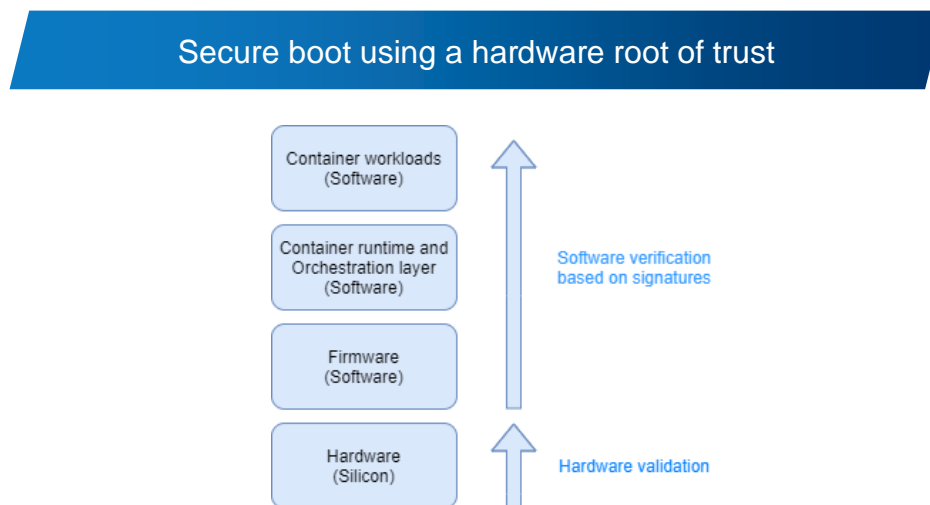


Figure 8

6.2 Secure platform for O-RU

An attacker with unauthorized access to the management interface of an unprotected O-RU could allow an attacker to steal unprotected private keys, certificates, hash values and/or inject malwares and/or manipulate existing O-RU software. An attacker could further launch denial-of-service, intrusion, and replay attacks on other network elements including an O-DU.

Therefore, hardening of the O-RU platform will ensure enough equipment security to substantially reduce the attack surface that would otherwise exist in an unprotected O-RU. Security precautions on the O-RU can be divided into three aspects.

1. Supply chain security
2. Physical security
3. Network security

Supply chain security ensures that throughout the supply chain process of manufacturing, from O-RU to its final installation site and commissioning, a controlled secure chain of custody process is followed. This ensures that the O-RU is properly tracked and tagged.

Physical security ensures that the physical O-RU is sealed with non-tamper-able screws that cannot be easily broken or opened and in the event of tampering or forced opening, all O-RU functionality will be disabled so that the O-RU becomes inoperable. This is in addition to all the physical and logical ports being secured and isolated, so that they cannot be used as a vulnerability entrance into the extended RAN network.

From a network security point of view, O-RU ensures that all authentication and communication security protocols are correctly performed and followed. To ensure reliable and secure software upgrades, the TPM procedures are implemented so that rogue software downloads are prevented. Finally, hardening features, such as disabling unnecessary software components and interfaces when not in use, running software at the correct privilege-level, scrambling/encryption of data in storage, and secure boot and hardware-based security module, are part of the comprehensive security processes on the typical O-RU to ward off as well as prevent unauthorized access to the O-RU

7. KEY SECURITY DIFFERENTIATORS IN OPEN RAN

The following table highlights some of the key differentiators that Open RAN provides compared to a closed RAN or the classical gNB.

Differentiator	Open RAN	Closed RAN
Security of open fronthaul	Provides visibility to the security measure taken to protect this interface. Open, standardized interfaces remove vulnerabilities or risk that comes with proprietary and potentially untrusted implementation.	Protection measure taken to protect CPRI interface in a closed RAN is not known
Operator has full control in building a secure platform	Open RAN's disaggregated architecture allows network operators to build cloud-native platforms by selecting suppliers that meet all the required industry security standards and certifications.	Operator has no control of how the virtualized platform is assembled. It is fully vendor driven.
Better enforcement of security controls in cloud infrastructure	A cloud infrastructure supplier will be directly under an agreement with the operator and will be responsible for security of the cloud infrastructure.	Operator has no direct visibility of the cloud infrastructure provider
Disaggregated platform allows for better visibility and automated monitoring of the network	A cloud native architecture allows operators to deploy the latest security tools for monitoring vulnerabilities and automated remediation measures as required	Operator has no visibility to this information. The operator is fully dependent on the vendor to detect and remediate vulnerabilities in the network
Adoption of industry best practices in development of containerized applications	Allows adoption of industry best practices such as "secure by design" DevSecOps, automated testing in development of containerized applications. Operator also has an option to work with the supplier to determine and influence CI/CD processes used by the supplier.	It is fully vendor driven, and an operator has no mechanism to verify the software development process used by the vendor.
Protection of cryptographic key	NG-RAN cryptographic key (KgNB) is stored in CU, which is located in a centralized data center inside the network.	Stored at the cell site and can be potentially stolen especially when HSM is not implemented in gNBs.

8. CONCLUSION

At the heart of Open RAN is the use of cloud native architecture, the same architecture that is the bedrock of today's internet and public cloud. Security practices in virtualized deployments are mature and used across the cloud computing industry. Virtualized deployment in telecom networks is not new. Operators already have virtualized infrastructure in their data centers and many have deployed virtual workloads for other components in the network including: packet core, IMS, and other applications such as CDN. With a disaggregated architecture, operators will now additionally benefit from security expertise and experience of today's large cloud infrastructure suppliers in managing the security of large IT cloud environments.

Operator regains control as the operator now understands what is required to build and maintain a secure infrastructure. Open RAN is built on a cloud native platform with clear responsibilities and accountability established between hardware/infrastructure suppliers, a hybrid-cloud platform supplier, and RAN software suppliers. It enables network operators to select suppliers that meet all the required industry security standards and certifications.

Open RAN leverages several security industry best practices used in the cloud computing industry. A "shift-left" strategy in the software development process integrates security controls and practices into every phase of the software development. With DevSecOps integrated into the CI/CD pipeline, this also brings automation into secure code reviews and security testing. Use of automated tools for detection, remediation of vulnerabilities in open-source software and detection, and management of secure posture provides an operator with quick detection and resolution of anomalies in the network.

O-RAN Alliance's architecture for RAN is built on the secure foundation of zero trust where network elements mutually authenticate with each other in order to communicate. All communication between them is transported over a secure interface per industry best practices specified by O-RAN Alliance's security specifications. While standards are still evolving, the Open RAN pioneers and ecosystem vendors like AltioStar, Mavenir, Fujitsu and Red Hat, as well as early adopters like Rakuten, Vodafone, Telefonica, NTT Docomo and DISH have ensured that all the interfaces are secured using certificate based security.

Every network element in the Open RAN network undergoes platform hardening as per 3GPP's security assurance specifications and other well-known cloud computing industry benchmarks such as CIS. This protects the network from an attacker gaining unauthorized access and subjecting the network to Denial-Of-Service (DOS) attacks or gaining illegal access.

In summary, open, standardized interfaces remove vulnerabilities or risk that comes with proprietary and potentially untrusted implementation and provides an operator full visibility and control over the cloud environment and network in general.

APPENDIX

References

- [1] 3GPP TS 38.401: NG-RAN; Architecture description
- [2] 3GPP TS 38.473: NG-RAN; F1 Application Protocol (F1AP)
- [3] O-RAN Architecture Description (O-RAN.WG1.O-RAN-Architecture-Description)
- [4] 3GPP TS 33.501: Security architecture and procedures for 5G system (Release 16)
- [5] NIST Special Publication 800-207: Zero Trust Architecture
- [6] O-RAN Architecture Description Chapter X – O-RAN Security
- [7] O-RAN Control, User and Synchronization Plane Specification (O-RAN WG4.CUS)
- [8] O-RAN Management Plane Specification (O-RAN.WG4.MP)
- [9] NIST Special Publication 800-190: Application Container Security Guide
- [10] 3GPP TS 33.511: Security Assurance Specification (SCAS) for the next generationNode B (gNodeB) network product class
- [11] CIS benchmarks: <https://www.cisecurity.org/cis-benchmarks/>

Acronyms

3GPP	3rd Generation Partnership Project	OCI	Open Container Initiative
5G	5th Generation	O-CU	O-RAN Central Unit
CA	Certification Authority	O-DU	O-RAN Distributed Unit
CI/CD	Continuous Integration/Continuous Delivery	O-RAN	Open Radio Access Network
CIS	Center for Internet Security	O-RU	O-RAN Radio Unit
CMP	Certificate Management Protocol	PDCP	Packet Data Convergence Protocol
CNF	Cloud native Network Function	PNF	Physical Network Function
CP	Control Plane	RAN	Radio Access Network
CPRI	Common Public Radio Interface	RBAC	Role Based Access Control
CRI-O	Container Runtime Interface for OCI compatible runtimes	RIC	Radio Intelligent Controller
CRMT	Core Root of Trust Measurement	RLC	Radio Link Control
CSP	Cloud Service Provider	RT-RIC	Real-Time Radio Intelligent Controller
CU	Central Unit	RRM	Radio Resource Management
CUS	Control, User & Synchronization	RRU	Remote Radio Unit
DOS	Denial of Service	SAST	Static Application Security Testing
DDOS	Distributed Denial of Service	SCRM	Supply Chain Risk Management
DTLS	Datagram Transport Layer Security	SDAP	Service Data Adaptation Protocol
DU	Distributed Unit	SDLC	Software Development Life Cycle
EST	Enrollment over Secure Transport	SIEM	Security Information and Event Management
FIPS	Federal Information Processing Standards	SLA	Service Level Agreement
GSMA	Global System for Mobile Communications Association	SMO	Service Management and Orchestration
HSM	Hardware Security Module	SSH	Secure Shell
ICAM	Identity, Credential and Access Management	STG	Security Task Group
LLS	Lower Layer Split	SUCI	Subscription Concealed Identifier
LUKS	Linux Unified Key Setup	TCO	Total Cost of Ownership
MAC	Mandatory Access Control	TLS	Transport Layer Security
MEC	Multi-access Edge Computing	TPM	Trusted Platform Module
MITM	Man-in-the-Middle	UE	User Equipment
NDS	Network Domain Security	UP	User Plane
NESAS	Network Equipment Security Assurance Scheme	VNF	Virtualized Network Function
NF	Network Function	ZTA	Zero Trust Architecture
NIST	National Institute of Standards and Technology		
NR	New Radio		
NR-RIC	Near Real Time RIC		

About Mavenir

Mavenir is building the future of networks and pioneering advanced technology, focusing on the vision of a single, software-based automated network that runs on any cloud. As the industry's only end-to-end, cloud-native network software provider, Mavenir is transforming the way the world connects, accelerating software network transformation for 250+ Communications Service Providers in over 120 countries, which serve more than 50% of the world's subscribers.

For more on Mavenir Solutions please visit our website at www.mavenir.com