P1 SECURITY WHITEPAPER

The O-RAN security gap

How open interfaces and the RIC widen the attack surface, and why controls that are supported but switched off protect nothing.

By Kye Grundy, P1 Security.

P1 SECURITY

contact@p1sec.com

Contents

Executive summary

Open RAN (O-RAN) is one of the most significant structural changes in mobile access networks in a generation. It takes the radio access network, historically a closed, single-vendor black box, and breaks it into open, interoperable, multi-vendor, cloud-native building blocks. The traditional baseband unit is disaggregated into an O-RAN radio unit (O-RU), distributed unit (O-DU) and centralised unit (O-CU). New intelligence layers, the near-real-time RIC, the non-real-time RIC and the service management and orchestration (SMO) framework, add dynamic, data-driven control of the network.

The promise is real: less vendor lock-in, more innovation, and the flexibility of commercial off-the-shelf hardware and cloud deployment.

That same openness reshapes the security picture. Interfaces that used to be hidden inside a proprietary product are now publicly specified, which is good for interoperability and equally good for an attacker who wants to understand where the seams are. This whitepaper walks through the O-RAN attack surface as P1 Security sees it: the open interfaces (open fronthaul with its C, U, S and M planes, plus E2, O1, O2, A1, R1 and Y1), the RIC intelligence layers and the xApps and rApps that run on them, and the multi-vendor software supply chain underneath everything.

The recurring theme is not that O-RAN is insecure by design. The O-RAN Alliance, principally through its WG11 security group, has produced a substantial and improving body of security requirements, largely by reusing proven technologies such as TLS, IPsec and OAuth 2.0, and it is moving the architecture toward zero trust. The problem is the gap between specification and deployment. For most controls, vendor support is mandatory but operator use is optional. A vendor can ship the ability to protect an interface, and an operator can still deploy it turned off. Security controls that are supported but not enabled protect nothing.

The practical conclusion is straightforward. O-RAN vulnerabilities are no longer purely theoretical, and paper compliance is not the same as a secure deployment. Operators should enable and configure the controls that vendors are required to provide, insist on software integrity and SBOMs across every

supplier, adopt zero-trust principles internally, and validate the real deployment through telecom-specific penetration testing rather than trusting the specification alone.

This paper is written for telecom security architects, RAN and core network teams, SOC and detection teams, network equipment vendors, and the regulators and operators who set and verify the security baseline.

Key takeaways

- Most O-RAN interface protections (TLS on O1, A1, R1, O2 and Y1; IPsec on E2) are mandatory for vendors to support but optional for operators to use, so auditing what is actually enabled is the first security control.
- P1 Security has demonstrated a fronthaul control-plane denial of service in the lab: injecting malformed eCPRI (enhanced CPRI) C-plane messages was enough to disrupt an emulated O-RU exchanging legitimate traffic with an O-DU.
- CVE-2023-40998 makes the risk practical, not theoretical: a crafted E2 message crashed the E2 termination of an open-source near-RT RIC reference implementation, causing a denial of service tracked in P1 Security's VKB.
- E2, O1, A1 and Y1 expose subscriber-level data, including UE identifiers, geolocation, RRC states and UE tracing, if deployed without the protections O-RAN Alliance WG11 specifies.
- O-RAN mandates an SBOM for every network-function release and artifact-level software signing verified against a trusted root CA; operators should ingest and enforce both rather than file them.
- O-RAN specifications change roughly every six months, so posture must be validated on the live deployment at its actual specification version through telecom-specific penetration testing.

1. Introduction: why O-RAN changes the security picture

Open RAN represents a fundamental architectural shift. It moves the industry away from proprietary, single-vendor radio access networks toward an open, intelligent, multi-vendor and cloud-native ecosystem. The stated goals are to foster innovation, increase flexibility and improve cost efficiency, largely by breaking the vendor lock-in that has defined the RAN for decades.

To achieve that, O-RAN does several things at once. It disaggregates the baseband into distinct functions: the virtualised O-CU and O-DU, plus the O-RU radio unit at the cell site. It adopts the 7.2x functional split between the O-RU and O-DU. It leans on cloud-native deployment and commercial off-the-shelf hardware. It expects a larger presence of multiple vendors and open-source software inside a single network. And, most distinctively, it introduces new intelligence layers, the service management and orchestration framework, the non-real-time RIC and the near-real-time RIC, that enable dynamic, data-driven and increasingly AI-assisted control of the RAN.

Every one of those choices carries a security consequence. Interfaces that were previously obscured inside a single supplier's product are now openly defined, so their message formats, parameters and procedures are documented for anyone, defenders and attackers alike. The number of software suppliers and open-source components inside one network rises sharply. Control of the RAN becomes something that software applications, potentially from third parties, can exert in near real time. None of this makes O-RAN inherently unsafe, but it does mean the threat surface is wider, more distributed and more openly understood than in a classic RAN.

The O-RAN Alliance is explicit that security is a design goal. Its aim is to let communication service providers deploy and operate O-RAN with the same or higher level of confidence as a 3GPP-specified RAN, and it builds on existing 3GPP and IETF work rather than reinventing it. Its WG11 security group specifies requirements and controls across the attack surface and is steering the architecture toward a zero-trust posture. The rest of this paper takes that framework seriously while being honest about where the real-world risk sits, which is almost always in how the framework is implemented and operated.

2. From proprietary RAN to open, disaggregated RAN

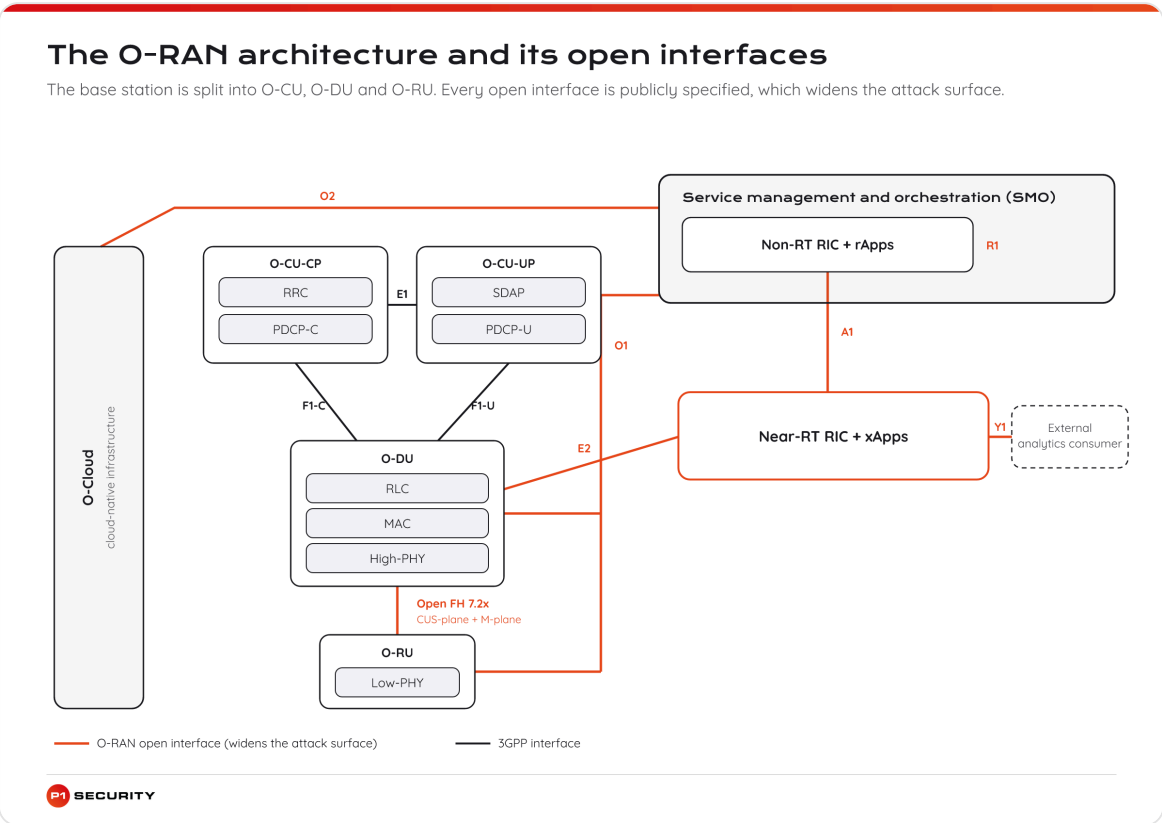

Figure 1. The O-RAN estate: every open interface here is a documented, testable surface, and each needs its controls both supported and switched on

O-RAN differs from a traditional RAN in its degree of openness and disaggregation. It helps to place O-RAN on the spectrum of RAN architectures, because the security implications track the degree of openness and disaggregation.

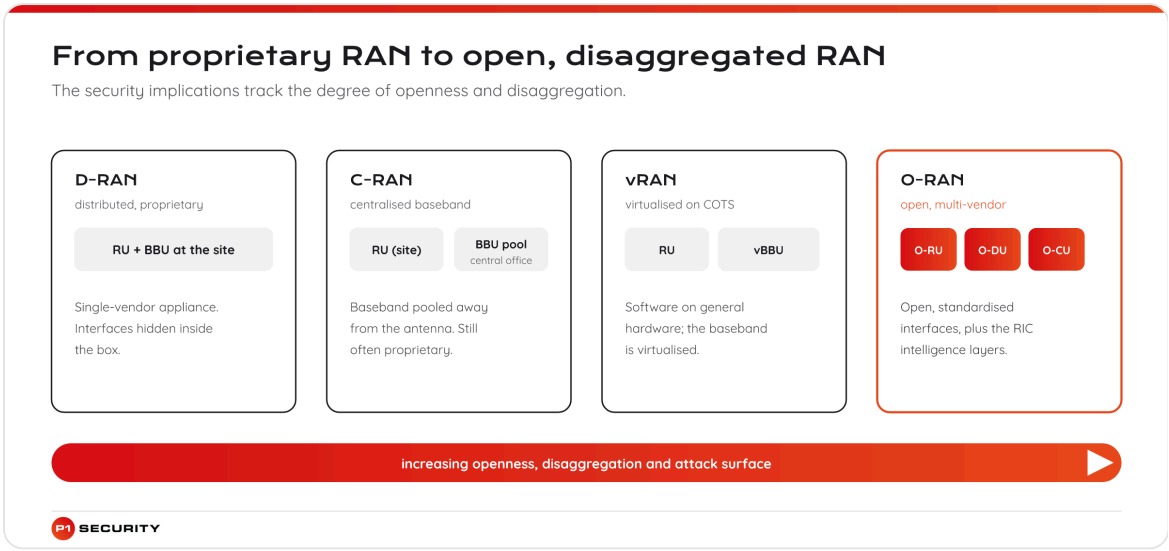

Figure 2. The RAN architecture spectrum, from a closed appliance to open O-RAN

A traditional distributed RAN (D-RAN) uses proprietary hardware and software, with a radio unit at the site and a self-contained baseband unit. Centralised RAN (C-RAN) and virtualised RAN (vRAN) progressively centralise and virtualise the baseband, but the software often remains proprietary even where the hardware opens up. A disaggregated RAN splits the baseband into separate CU, DU and RU functions. O-RAN goes furthest: it defines an O-CU, O-DU and O-RU connected by open, standardised interfaces, adds the RIC intelligence layers on top, and expects interoperable software running on open hardware and cloud infrastructure.

The practical difference for a defender is what becomes reachable and by whom. In a proprietary RAN, the interfaces between baseband components were internal, vendor-specific and effectively invisible. In O-RAN, the equivalent interfaces are public specifications carried over familiar transport such as Ethernet, IP, SCTP and HTTP. That is precisely what enables multi-vendor interoperability, and it is precisely what turns previously hidden internal boundaries into documented, testable and potentially attackable surfaces. Where O-RAN reuses proven 3GPP and IETF technology, it inherits both the strengths and the known weaknesses of those choices. The security analysis therefore has to consider not just individual interfaces but how a foothold on one, often at a physically exposed cell site, can be used to move toward more sensitive functions.

3. How P1 Security approaches O-RAN security

The perspective in this paper comes from hands-on telecom security work rather than a standards read-through. P1 Security's research on O-RAN combines lab-based study of real 4G, 5G and O-RAN equipment with penetration-testing experience across RAN and core networks, and it tracks the O-RAN Alliance specifications as they evolve.

Three points about method shape the analysis. First, telecom security is not general IT security. The interfaces, protocols and failure modes here, eCPRI on the fronthaul, E2AP over SCTP, NETCONF and YANG on the management plane, are specific to mobile networks, and generic tooling does not understand them. Second, the most useful findings are usually about movement, not a single flaw: how access to a physically exposed cell site or a single interface can be leveraged toward the SMO, the core or the wider network. Third, O-RAN is still early. There are relatively few production deployments, and those that exist often do not run the full feature set, so first-hand posture assessment is genuinely difficult. That is an argument for testing real deployments carefully, not an argument for assuming the specification protects you.

4. Securing the open interfaces

Two items are not theoretical. The first is a control-plane denial of service P1 Security has demonstrated in the lab over the open fronthaul; the second is a denial-of-service CVE already tracked in P1 Security's VKB. Both are flagged where they belong below.

O-RAN's open interfaces are the enabler of multi-vendor interoperability and, at the same time, the most visible expansion of the attack surface. Each is publicly defined, each carries sensitive control or subscriber-related data, and each needs its security controls to be both supported and switched on. The sections below summarise the function of each interface and the risk it introduces if left unprotected.

4.1 The open fronthaul and its four planes

The open fronthaul (O-FH) connects the O-RU at the cell site to the O-DU, with optional integration to the SMO in the hybrid model. Its functionality is split across four planes, each with its own security profile.

Control plane (C-plane). The C-plane carries the real-time control information that tells the O-RU how to transmit and receive user data. O-RAN moves from the proprietary, time-division CPRI to packet-based eCPRI (enhanced CPRI), which creates a far more familiar Ethernet and IP environment for an attacker to work in. The eCPRI header and the O-RAN-defined control payload contain parameters sensitive enough to enable a denial of service or user reconnaissance. An attacker with access to the fronthaul link can inject invalid or malformed C-plane messages, which can cause a denial of service for the subscribers served by that O-RU. Because fronthaul links terminate at cell sites that often have limited physical security, this is also a realistic entry point for pivoting deeper into the network.

User plane (U-plane). The U-plane transports the actual user data as IQ (in-phase/quadrature) samples between the O-DU and O-RU, in both directions, carrying everything from 3GPP signalling to application data. Protection for that user data comes from PDCP (Packet Data Convergence Protocol) ciphering applied higher in the stack at the gNB rather than from the fronthaul itself, and 3GPP TS 33.501 makes integrity and confidentiality protection of user data between the UE and gNB optional to use. O-RAN does not itself mandate encryption of the fronthaul CUS-plane. So if PDCP protection is not enabled, and if application-layer protection such as HTTPS is absent, user data traversing the U-plane could in principle be monitored or modified. Extracting valid samples requires significant processing, which raises the effort involved, but effort is not the same as safety.

Synchronisation plane (S-plane). The S-plane provides highly accurate timing to the O-RU, typically using PTP (IEEE 1588) or SyncE, with the timing source depending on the chosen architecture (the LLS-C1 to LLS-C4 options). Accurate timing is essential to a healthy RAN. Where PTP or SyncE run without security controls, an attacker who can degrade the timing distribution can cause an O-RU to lose its timing reference, and an O-RU without a valid timing reference cannot serve its users. The result is a denial of service affecting every device connected to that cell, with the operational and reputational cost that implies for a busy site. The large number of timing configuration options is itself a source of deployment complexity and security gaps.

Management plane (M-plane). The M-plane handles fault, configuration, accounting, performance and security (FCAPS) management of the O-RU from the O-DU or the SMO, using NETCONF with notifications over HTTP or HTTPS. If its controls are implemented incorrectly, configuration and notification messaging can be monitored or modified. Malicious use of management services enables highly targeted attacks: pushing malicious configuration, forcing a firmware rollback to a known-vulnerable version, or watching sensitive management notifications. The O-RU-to-SMO

management option is particularly notable as a pivot point, because it can allow an attacker to escalate from a less-protected cell site toward the SMO and the operator's higher layers.

4.2 The E2 interface and the near-RT RIC

E2 connects the near-real-time RIC to the RAN nodes (the O-CU and O-DU) and enables near-real-time control and optimisation, carried by E2AP over SCTP. The specific actions a controlling application can take are defined by E2 service models (E2SMs). E2 messages carry sensitive information at the RAN, cell and even individual-subscriber level. Examples of what different service models expose include:

- the key performance measurement model (UE ID, beam ID, cell global ID, measurement values);
- the cell configuration and control model (cell global ID, X2 and Xn handover allow and block lists, tracking area information);
- the RAN control model (UE ID, serving cell, RRC states and even the ability to receive copies of a UE's RRC messages).

Because xApps and near-RT RICs may come from different suppliers, each with its own dependencies, access levels and code quality, the effective security profile of E2 depends on components an operator may not fully control. Without adequate protection, E2 message content can be monitored or modified, opening the door to sophisticated attacks such as rogue RAN reconfiguration.

4.3 The O1 management interface

O1 runs between the O1 nodes (O-DU, O-CU and the RICs) and the SMO, for management services such as configuration and notifications. Configuration management uses NETCONF with O-RAN-defined YANG models, and notifications are delivered by RESTful APIs over HTTP, most of them based on 3GPP specifications (provisioning management, fault supervision, trace management), plus O-RAN-defined functions such as physical network function registration. If O1 lacks adequate controls, configuration and notifications can be monitored or modified. That enables malicious reconfiguration of the RAN, and it exposes sensitive notifications, including UE tracing, which reveals detailed information about a subscriber's activity on the network.

4.4 The A1 policy interface

A1 connects the non-real-time RIC to the near-real-time RIC, letting the non-RT RIC guide the real-time behaviour of the network. It offers three services: A1-P for policies (for example, keep a cell's load below a threshold), A1-EI for enrichment information that the near-RT RIC may not have (for example, UE geolocation), and A1-ML for the exchange of AI and ML models and training data. Weak controls on A1 could allow an attacker to modify policies and drive undesired RAN behaviour, or to monitor enrichment information and track subscribers. The parameters at risk include the UE identifier, the cell identifier and UE geolocation. The trust question is sharpened by the fact that third-party xApps can request A1 services through the near-RT RIC, so an operator has to ask how confident it really is that a given xApp is not compromised.

4.5 R1, O2 and Y1

R1 is an internal interface within the SMO framework that lets rApps communicate with the non-RT RIC platform. rApps carry considerable responsibility: through the platform they can manage the RAN via O1 and the O-FH M-plane, generate A1 policies for the near-RT RIC, and reach O-Cloud management via O2. Like xApps, rApps are likely to come from many suppliers with varying security profiles. A weakly protected R1 interface or a malicious rApp could therefore achieve a high level of control over the O-RAN network functions, including monitoring data and configuration, targeted reconfiguration, deployment of malicious backdoors through software updates, or the forced shutdown of cells.

O2 connects the SMO to the O-Cloud and manages both the cloud infrastructure and the O-RAN cloud-native network functions running on it, as a RESTful HTTP interface. It splits into O2ims for the underlying infrastructure and O2dms for deployment of the network functions, and O-RAN defines a Kubernetes-oriented O2dms interface for workload deployment. Malicious action on O2 could cause severe failures, from shutting down network functions to deploying malicious code. The harder question is assurance: it is difficult to guarantee, or even fully understand, the security posture of an external O-Cloud that the operator does not directly control.

Y1 lets the near-RT RIC export RAN analytics information to an external consumer that has subscribed to it. Unlike most O-RAN interfaces, Y1 is externally facing, so it deliberately exposes internal information to outside peers. That raises two concerns: whether a Y1 peer is trustworthy, and whether a malicious peer could trigger a denial of service against the near-RT RIC. Without adequate controls, Y1 metrics could be monitored or modified, disclosing UE identifiers and the corresponding service information for specific subscribers.

4.6 O-RAN interface protections at a glance

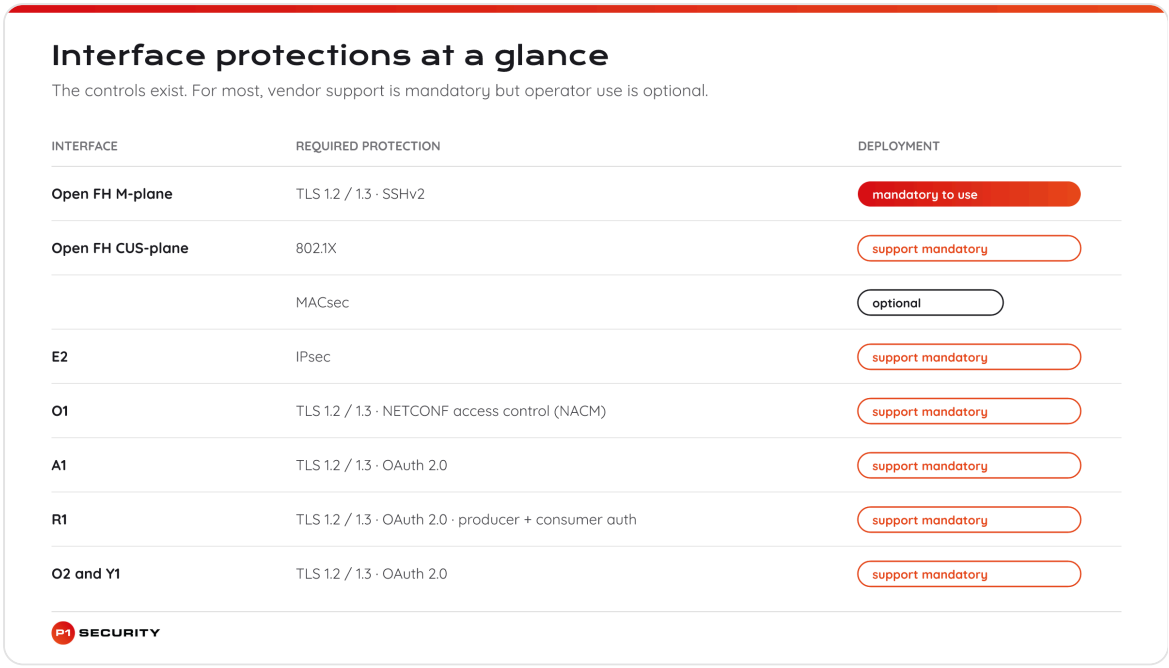

Figure 3. O-RAN WG11 interface protections; a control that is supported is not the same as one that is enabled

The O-RAN Alliance specifies protections for these interfaces in the WG11 security requirements specification, generally reusing proven protocols. In summary:

- **O1:** TLS 1.2 or 1.3, with NETCONF access control (NACM). Shall be supported.
- **A1:** TLS 1.2 or 1.3 and OAuth 2.0, or NACM-based protection.
- **R1:** TLS 1.2 or 1.3 and OAuth 2.0, with both API producer and consumer authenticated.
- **O2 and Y1:** TLS 1.2 or 1.3 and OAuth 2.0.
- **E2:** IPsec.
- **Open fronthaul M-plane:** TLS 1.2 or 1.3 and SSHv2, which are mandatory to use.
- **Open fronthaul CUS-plane:** 802.1X and MACsec, where 802.1X support is mandatory and its use optional, whilst MACsec is optional to support and optional to use (recommended, not enforced). Where 802.1X is not used, authentication and authorisation shall still be provided.
- As a general principle, every interface should handle malformed input safely; this robustness is best confirmed by testing rather than assumed.

The single most important nuance sits underneath this list. For most of these controls, support by the vendor is mandatory but use by the operator is optional. The framework only protects a deployment when the operator turns the controls on and

configures them correctly.

What to verify (open interfaces)

- Are the mandatory-to-use O-FH M-plane protections (TLS, SSHv2) actually enabled, not just supported?
- Where 802.1X and MACsec are optional, is equivalent authentication and authorisation in place on the fronthaul LAN segments?
- Is E2 protected with IPsec, and are A1, O1, O2, R1 and Y1 protected with TLS and OAuth 2.0 in the live deployment?
- Is PDCP integrity and confidentiality protection of user data enabled, given that 3GPP makes it optional?
- Are cell sites and fronthaul links treated as physically exposed, with the O-RU-to-SMO M-plane path hardened against use as a pivot?

5. Securing the RIC intelligence and orchestration layers

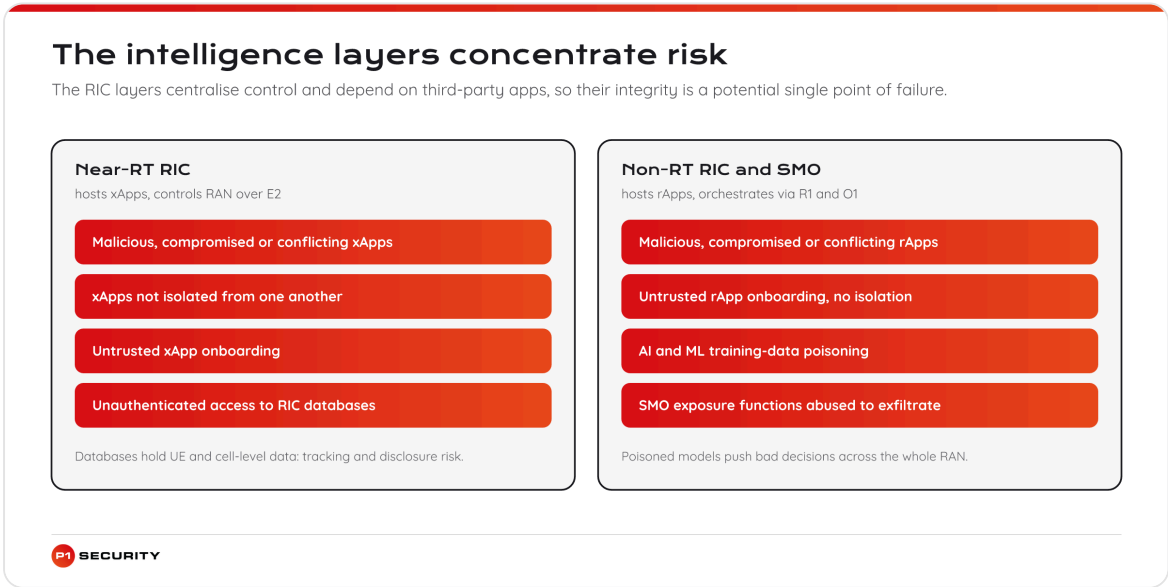

Figure 4. Recurring threats to the O-RAN RIC and orchestration layers, where third-party apps concentrate risk

The RIC intelligence layers and the SMO are the brains of an O-RAN network. They centralise control and optimisation, which makes their integrity a potential single point of failure for large segments of the RAN. They also depend heavily on third-party

applications, open-source software and a web of internal APIs, which concentrates trust, data-access and platform-integrity risk in one place. The threats differ between the near-real-time and non-real-time domains.

5.1 Near-RT RIC and xApps

The near-RT RIC hosts xApps that exert near-real-time control over RAN nodes through E2 and can be influenced by A1 policies from above. The main application-level risks are malicious, compromised or conflicting xApps, xApps that are not isolated from one another, and the onboarding of untrusted xApps. Alongside these sit data risks, principally unauthenticated or unauthorised access to the near-RT RIC databases, which hold sensitive UE and cell-level information and could enable subscriber tracking or information disclosure. There are also interface risks such as a malicious E2 node, a malicious external Y1 consumer, a malicious O1 peer, and APIs that lack proper authentication or authorisation.

5.2 Non-RT RIC, rApps and the SMO

In the non-real-time domain the parallel risks appear with rApps: malicious, compromised or conflicting rApps, rApps that are not isolated, and untrusted rApp onboarding. Beyond the applications, there are threats between the SMO framework and the non-RT RIC in either direction, unauthenticated or unauthorised access to non-RT RIC data stores, and a lack of authentication and authorisation toward non-RT RIC services including O2, O1 and the O-FH M-plane. A specific concern is the malicious use of the SMO's data and service exposure functions to exfiltrate data to external peers, along with malicious external peers consuming SMO services or data.

5.3 AI and ML as a new risk surface

The non-real-time domain introduces a risk that has no real precedent in a traditional RAN: poisoning of the AI and ML data used to train models. Because the intelligence layers can push decisions system-wide, corrupting the training data or the models can translate into incorrect optimisation decisions being applied across the RAN, quietly and at scale. This is why the integrity and confidentiality of models and their training data, at rest, in use and in transit, is a first-class security requirement rather than a footnote.

5.4 What the O-RAN Alliance mandates

WG11 defines a substantial set of controls for these layers.

- **API protection.** The near-RT RIC is mandated to support authentication, integrity and confidentiality using mTLS, plus authorisation using OAuth 2.0, for transactional RESTful APIs, and to support IPsec for time-critical APIs. The non-RT

RIC and SMO are mandated to support TLS and OAuth 2.0, and both API producer and consumer must be authenticated across R1.

- **Database protection.** The near-RT RIC shall authenticate and authorise any xApp access to its internal databases. For the non-RT RIC and SMO there is a notable gap: there is no explicit database-protection requirement. Protection may be achieved implicitly, because an rApp reaching those databases would first be authenticated at R1, but the absence of an explicit control is worth flagging.
- **xApp and rApp protection.** It is mandatory for xApps and rApps to be associated with unique identifiers and to be mutually authenticated with their respective RIC framework, with identifiers unique enough to resist brute forcing. rApps are also required to be able to recover from malicious behaviour over R1.
- **External peering.** Y1 consumers are authenticated through TLS and OAuth 2.0, but the security posture of the peer itself cannot be guaranteed. SMO external interfaces are mandated to support mTLS and OAuth 2.0.
- **AI and ML protection.** Access to models must be authenticated and authorised under least privilege, and models and their training, testing and inference data must be protected for integrity and confidentiality at rest, in use and in transit.

What to verify (intelligence layers)

- Is every xApp and rApp uniquely identified and mutually authenticated with its RIC framework, and are xApps isolated from each other?
- Is there an onboarding process that establishes trust in third-party xApps and rApps before they can act on the RAN?
- Are the near-RT RIC databases protected by authentication and authorisation, and is equivalent protection applied to non-RT RIC and SMO data stores despite the missing explicit requirement?
- Are the SMO exposure functions constrained so they cannot be abused to exfiltrate data to external peers?
- Are AI and ML models and their training data integrity-protected and access-controlled?

6. Supply chain and software integrity

In a multi-vendor O-RAN ecosystem built from commercial and open-source components, software and supply-chain integrity becomes a primary security concern. If an operator cannot establish what is inside a deployment and that it has not been tampered with, none of the interface-level controls can be fully trusted.

Open-source software. O-RAN's use of open-source components brings the usual open-source risks, several of which OWASP catalogues: known vulnerabilities in dependencies, compromise of an otherwise legitimate package, unmaintained software, and untracked dependencies. To manage these, the O-RAN Alliance points to OpenSSF, the Linux Foundation initiative that provides an open-source project security baseline, automated scanning, the Sigstore signing and verification service, and the "SBOM everywhere" effort. O-RAN Software Community projects are adopting OpenSSF best practices and a badging process to demonstrate it.

Software bill of materials. O-RAN makes a software bill of materials (SBOM) mandatory: each network function vendor must create and supply an SBOM for every release, with a minimum set of data (commonly expressed in the SPDX or CycloneDX format), and the SBOM must be securely protected and delivered. The SBOM is what lets an operator answer the question that matters when a new dependency vulnerability is disclosed, which is simply whether the deployment contains the affected component. An operator cannot patch what it does not know it runs.

Software package signing and verification. O-RAN also defines lifecycle-management requirements for software packages, for both closed and open source. A package should be certified by its provider, for example through penetration testing, must include the provider's signing certificate and signatures, and each artifact within it must be individually signed. Verification then follows a clear chain:

1. The provider hashes the package and encrypts that hash with its private key to create a digital signature.
2. The provider sends the package, the signature and its public-key certificate to the operator.
3. The operator's SMO checks the certificate against a trusted root certificate authority it has installed.
4. If the certificate is trusted, the SMO decrypts the signature with the provider's public key to recover the original hash, computes its own hash of the received package, and confirms authenticity and integrity only if the two match.

What to verify (supply chain)

- Does every vendor supply a complete, accurate SBOM for every release, and does the operator actually ingest and use it?
- Are software packages certified, signed at the artifact level, and verified against a trusted root CA before deployment?
- Is there an open-source inventory across xApps, rApps and RIC components, with continuous vulnerability tracking?
- Are OpenSSF best practices required of suppliers, and is compliance checked rather than assumed?

7. Zero trust architecture in O-RAN

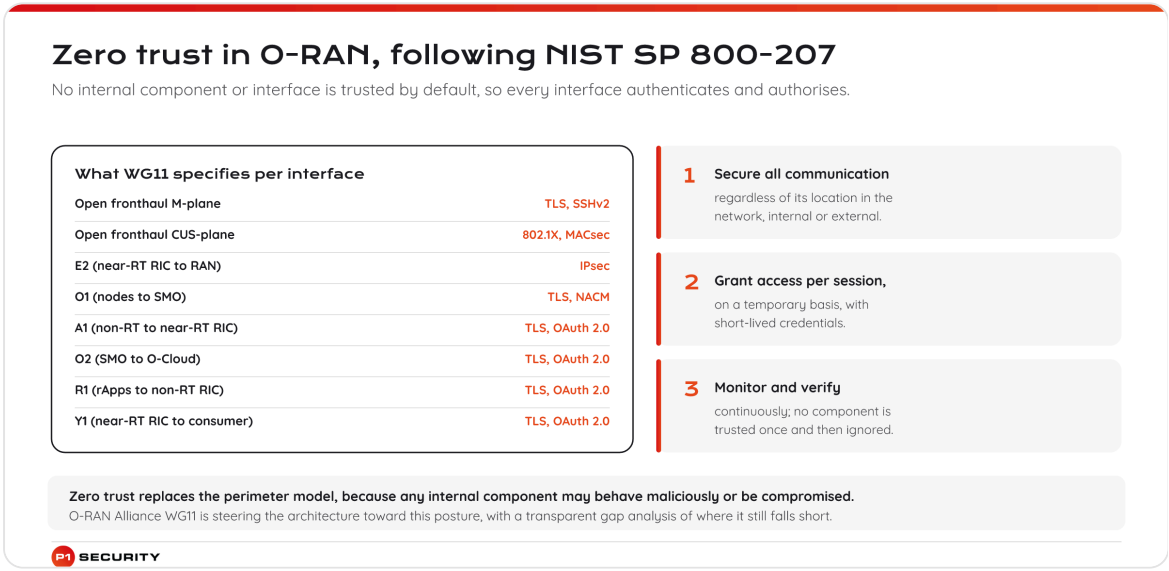

Figure 5. Zero trust in O-RAN, following NIST SP 800-207

Zero trust architecture in O-RAN, following NIST SP 800-207, means no internal component or interface is trusted by default: all communication is secured regardless of location, access is granted per session with short-lived credentials, and verification is continuous. The open, disaggregated, multi-vendor nature of O-RAN makes a traditional perimeter security model unsuitable. When components come from many sources and communicate over open interfaces, the network can no longer assume that internal traffic is inherently trustworthy, because internal components may behave maliciously either deliberately or through compromise.

O-RAN Alliance WG11 has published a technical report on applying zero trust to O-RAN, including a transparent gap analysis of where the current specifications fall short against zero-trust pillars. The Alliance's near-term goal is to reach the "initial" maturity level of the zero-trust architecture, with more advanced levels to follow. This is a credible direction of travel, and the transparency of the gap analysis is a strength, but it also confirms that O-RAN's zero-trust posture is still maturing rather than complete.

8. From theory to reality: known O-RAN vulnerabilities and the limits of paper security

Known O-RAN vulnerabilities now exist on the public record: implementation flaws such as CVE-2023-40998 have been found in O-RAN network functions, and the structural risk beneath them is the gap between what vendors must support and what operators actually enable. For most of O-RAN's short life the security discussion has been largely theoretical, because there have been few production deployments to examine. That is changing, and it is worth being precise about what "real" looks like.

The implementation gap

The most consequential finding is not a single exploit but a structural feature of the specifications: for many controls, support is mandatory for vendors while use is optional for operators. This is stated in the security specifications themselves, so it is not speculation. In concrete terms, a vendor must provide the means to protect the open fronthaul M-plane with TLS and to protect E2 with IPsec, but an operator can still choose to deploy those interfaces with no protection, which nullifies the framework for that deployment. Compliance at the vendor level and security at the deployment level are not the same thing.

A concrete example

Real vulnerabilities have already been identified in O-RAN network functions. One public case, tracked in P1 Security's VKB, is CVE-2023-40998, a denial-of-service issue in an open-source O-RAN near-RT RIC reference implementation. A specially crafted E2 message (an E2SubscriptionRequest) from an xApp caused incorrect length calculations that crashed the near-RT RIC's E2 termination (E2Term), a denial of service within the near-RT RIC. It illustrates the tangible risk of unprotected interfaces combined with untrusted or faulty application components, and it is the kind of implementation flaw that only surfaces through testing rather than specification review.

Existing telecom threats do not disappear

O-RAN sits on top of, and connects into, 3GPP-defined components and interfaces, so the established threats to those components carry over. O-RAN security is an addition to the existing mobile-network threat model, not a replacement for it. For the intelligence and orchestration layers specifically, MITRE's FiGHT knowledge base of adversary techniques against 5G systems is a useful reference for structuring these threats.

A demonstrated fronthaul attack

In a controlled lab, P1 Security demonstrated a control-plane attack over the open fronthaul against an open-source O-RU emulator connected to an O-DU that was exchanging legitimate C-plane and U-plane traffic. Injecting malicious C-plane messages toward the emulated O-RU was enough to disrupt its normal operation, illustrating the denial-of-service exposure described in section 4.1. The demonstration is described here at a conceptual level only; no tooling, payloads or reproduction steps are published.

A moving target

Two further realities complicate matters. O-RAN specifications evolve quickly, with new releases roughly every six months, which makes it hard for vendors and operators to keep deployed components aligned with the latest controls, and existing deployments may still target older specification versions where security controls were weaker or absent. And because deployments are still few and often not fully featured, assessing real-world posture at first hand remains difficult. Both point to the same conclusion: security has to be verified on the actual deployment, at its actual specification version, rather than inferred from the standard.

There are genuine positives to balance this. By reusing proven technologies such as TLS, IPsec and OAuth 2.0, O-RAN builds on years of operational and security scrutiny rather than inventing unproven mechanisms. And the O-RAN Alliance runs a public security-disclosure program on its website, where coordinated vulnerability disclosures can be submitted for review, which is a healthy step for the ecosystem.

To see how the gap between supported and enabled controls maps to your own network, you can benchmark your O-RAN exposure with TelcoASM at telcoasm.p1sec.com.

9. What this paper does not cover yet

This whitepaper focuses on the open interfaces, the intelligence layers and the software supply chain. Several important areas are deliberately out of scope here and are candidates for future P1 Security research and publication:

- O-Cloud and virtualisation security, including VNFs, CNFs and the hardware acceleration layer.
- Shared O-RU scenarios and their security impact, including multi-tenancy, conflicting operator configuration and trusted peering.
- The security implications of the 7.2x split and the resulting increase in O-RU complexity.
- Secure xApp and rApp onboarding, where security validation for applications not supplied by the RIC platform vendor is still an open item in the specifications.
- A deep dive into how individual controls such as TLS and OAuth 2.0 are best deployed in an O-RAN context.
- Pre-O-RAN issues, meaning the 3GPP interface and radio-level threats that O-RAN inherits.

10. A strategic framework for O-RAN security mitigation

Securing an O-RAN deployment does not require exotic defences. It requires turning the specification's optional controls into operational reality and validating the result. The measures below are organised into three tiers and map directly to the risks above.

Tier 1: enable and configure the controls that exist. For most O-RAN controls, vendor support is already mandatory, so the operator's first duty is to enable and correctly configure them. Turn on the mandatory-to-use fronthaul M-plane protections, protect E2 with IPsec, protect A1, O1, O2, R1 and Y1 with TLS and OAuth 2.0, provide authentication and authorisation on the fronthaul LAN segments, and enable PDCP integrity and confidentiality where user-data protection matters. Treating these as optional extras is the single biggest avoidable risk.

Tier 2: establish trust in components and data. Require complete and accurate SBOMs from every supplier, verify software-package signatures against a trusted root CA before deployment, and maintain an open-source inventory with continuous vulnerability tracking. Authenticate and uniquely identify every xApp and rApp, isolate

them from one another, and put a real onboarding and validation process in front of third-party applications. Protect the RIC and SMO databases, and protect AI and ML models and their training data for integrity and confidentiality.

Tier 3: adopt zero trust and test the real deployment. Move away from perimeter assumptions toward the zero-trust principles O-RAN itself is adopting: secure all communication regardless of location, grant access per session with short-lived credentials, and monitor continuously. Treat cell sites and fronthaul links as physically exposed, and harden the pivot paths that lead from them toward the SMO and core. Above all, validate the posture of the actual deployment through telecom-specific penetration testing, at the specification version actually running, because paper compliance does not equal a secure network. For European operators, these controls also align with the EU 5G Toolbox and the NIS2 Directive, which increasingly make demonstrable RAN security an expectation rather than an option.

11. Operator and vendor readiness checklist

A condensed, practical checklist drawn from the sections above.

- [ ] Mandatory-to-use O-FH M-plane protections (TLS, SSHv2) enabled and configured, not merely supported.
- [ ] Authentication and authorisation on the fronthaul LAN segments where 802.1X or MACsec are optional.
- [ ] E2 protected with IPsec; A1, O1, O2, R1 and Y1 protected with TLS and OAuth 2.0 in the live deployment.
- [ ] PDCP integrity and confidentiality of user data enabled where required, given the 3GPP option.
- [ ] Cell sites and fronthaul links treated as physically exposed; O-RU-to-SMO M-plane pivot paths hardened.
- [ ] Every xApp and rApp uniquely identified, mutually authenticated, and isolated from one another.
- [ ] Trusted onboarding and validation process for third-party xApps and rApps before they can act on the RAN.
- [ ] Near-RT RIC databases protected; equivalent protection applied to non-RT RIC and SMO data stores.
- [ ] SMO exposure functions constrained against data exfiltration to external peers.

- [ ] AI and ML models and training data integrity-protected and access-controlled under least privilege.
- [ ] Complete, accurate SBOM ingested from every vendor for every release.
- [ ] Software packages certified, artifact-signed, and verified against a trusted root CA before deployment.
- [ ] Open-source inventory maintained across xApps, rApps and RIC components, with continuous vulnerability tracking.
- [ ] Zero-trust principles applied internally: secured communication regardless of location, per-session access, continuous monitoring.
- [ ] Deployment specification version tracked, with a plan to keep pace with roughly six-monthly releases.
- [ ] Regular, telecom-specific penetration testing of the real O-RAN deployment.

12. Frequently asked questions

What are the main O-RAN security risks?

The widened surface from openly specified interfaces (open fronthaul C, U, S and M planes, plus E2, O1, O2, A1, R1 and Y1), trust in third-party xApps and rApps on the RIC layers, AI and ML data poisoning, and multi-vendor software supply-chain integrity.

Is O-RAN less secure than traditional RAN?

Not by design. O-RAN Alliance WG11 specifies substantial controls that reuse TLS, IPsec and OAuth 2.0. The real risk is the deployment gap: for most controls vendor support is mandatory but operator use is optional, so supported-but-disabled controls protect nothing.

What is the near-RT RIC?

The near-real-time RAN intelligent controller is an O-RAN control layer that hosts xApps and steers RAN nodes over the E2 interface in a near-real-time loop, so its integrity is a potential single point of failure for large RAN segments.

How do you secure O-RAN interfaces?

Enable the mandatory-to-use fronthaul M-plane protections, protect E2 with IPsec and A1, O1, O2, R1 and Y1 with TLS and OAuth 2.0, authenticate the fronthaul LAN, uniquely identify and isolate xApps and rApps, require SBOMs and signed software packages, and test the real deployment.

What is the difference between the near-RT RIC and the non-RT RIC?
The near-real-time RIC hosts xApps and steers RAN nodes over the E2 interface in a near-real-time loop. The non-real-time RIC hosts rApps within the SMO framework and guides the near-RT RIC's real-time behaviour through A1 policies, and it is also where AI and ML model training sits. Both centralise control, which is why their integrity is a potential single point of failure for large RAN segments.

What is CVE-2023-40998?
It is a denial-of-service vulnerability in an open-source O-RAN near-RT RIC reference implementation, tracked in P1 Security's VKB. A specially crafted E2 message (an E2SubscriptionRequest) from an xApp caused incorrect length calculations that crashed the near-RT RIC's E2 termination (E2Term), a denial of service within the near-RT RIC.

What is an SBOM and why does O-RAN require one?
A software bill of materials (SBOM) is an inventory of the components inside a software release, commonly expressed in the SPDX or CycloneDX format. O-RAN makes it mandatory: each network-function vendor must create and supply an SBOM for every release. It lets an operator answer whether a deployment contains a newly disclosed vulnerable component, because an operator cannot patch what it does not know it runs.

13. How P1 Security helps

P1 Security is a specialist mobile network and telecom security company. We help operators, governments, regulators and network equipment vendors find, detect and understand vulnerabilities across telecom signalling, core and RAN infrastructure, using protocol-aware methods that generic IT security cannot provide.

The analysis in this paper comes directly from that work. The services and products that map to it:

- **Audit, vulnerability assessment and red team.** Telecom-specific penetration testing and architecture review across the O-RAN estate: the open interfaces and their four fronthaul planes, the near-RT and non-RT RIC and their xApps and rApps, the SMO, the O-Cloud, and the RAN-to-core paths, tested against real deployments rather than specifications.
- **PTA (P1 Telecom Auditor).** Attack-surface management and telecom-protocol-aware vulnerability scanning and compliance checks across the telecom estate, so exposure is identified with remediation guidance rather than assumed away.

- **PTM (P1 Telecom Monitor).** Real-time intrusion detection for signalling and core networks, so misuse and lateral movement across the signalling and core layers that O-RAN connects into can be detected and acted on.
- **VKB (Vulnerability Knowledge Base).** The largest telecom vulnerability intelligence database, with a continuous feed from our research team, including O-RAN network-function findings such as the near-RT RIC issue referenced in this paper, so teams understand which vulnerabilities affect their network and why they matter.
- **Product security review.** Independent, telecom-specific security testing of O-RAN products and network functions, from RIC platforms and their xApps and rApps to O-CU, O-DU and O-RU software, so vendors in this multi-vendor supply chain can find and fix implementation and integration flaws before operators deploy them.
- **MSSP and TSOC.** For teams without an in-house telecom security operations centre, P1 runs a managed telecom security service and a Telecom Security Operations Centre (TSOC) that converges PTM detection with VKB intelligence, so misuse across the signalling and core layers O-RAN connects into is monitored and triaged continuously rather than only during an engagement.
- **Training.** P1's online training and hands-on Laboratory Playground take security, network and SOC teams from telecom fundamentals to advanced Open RAN, network slicing and RAN-sharing security, with expert-led video, lab simulations and certification, so your own team can find and reason about issues like these.

If you want to test or benchmark your own O-RAN deployment against the risks in this paper, contact us at contact@p1sec.com.

14. References and further reading

Standards and industry guidance:

- **O-RAN Alliance WG11** security specifications: the security requirements specification, threat-modelling work, the near-RT RIC and xApp security study, the non-RT RIC security work, the SMO security analysis, and the zero-trust architecture technical report.
- **O-RAN Alliance WG4** open fronthaul specifications (the control, user and synchronisation plane specification and the management plane specification).
- **O-RAN Alliance WG3** E2 specifications (E2AP and the E2 service models for RAN control, key performance measurement and cell configuration and control) and the Y1 specifications.
- **O-RAN Alliance WG2** A1 and R1 specifications.

- **O-RAN Alliance WG10 and WG6** O1, OAM architecture and O2 specifications. All O-RAN Alliance specifications are published at https://www.o-ran.org/specifications
- **3GPP TS 33.501** security architecture and procedures for the 5G system, including the optional integrity and confidentiality protection of user data between the UE and gNB: https://portal.3gpp.org/desktopmodules/Specifications/SpecificationDetails.aspx?specificationId=3169
- **eCPRI interface specification** (enhanced Common Public Radio Interface).
- **NIST SP 800-207** zero trust architecture: https://csrc.nist.gov/pubs/sp/800/207/final
- **OWASP top 10 risks for open source software**: https://owasp.org/www-project-open-source-software-top-10/, and **OpenSSF** (Linux Foundation) open-source security best practices, Sigstore and SBOM tooling: https://openssf.org/
- **MITRE FiGHT** adversary-technique knowledge base for 5G systems: https://fight.mitre.org/
- **NTIA** minimum elements for a software bill of materials (2021): https://www.ntia.gov/report/2021/minimum-elements-software-bill-materials-sbom
- **CISA/NSA ESF** Open Radio Access Network security considerations (2022): https://www.cisa.gov/news-events/alerts/2022/09/15/cisa-and-nsa-publish-open-radio-access-network-security
- **BSI** (German Federal Office for Information Security) Open RAN risk analysis (2022): the standard European counterpart to the ESF guidance.
- **GSMA NESAS and 3GPP SCAS**: network equipment security assurance and its security assurance specifications, complementing WG11 requirements and the SBOM demands above: https://www.gsma.com/solutions-and-impact/industry-services/assurance-services/network-equipment-security-assurance-scheme-nesas/
- Public CVE record for CVE-2023-40998, the O-RAN Software Community near-RT RIC denial-of-service issue referenced in section 8: https://nvd.nist.gov/vuln/detail/CVE-2023-40998

P1 Security resources:

- SecurityHub TelcoSec glossary entries for O-RAN, O-RU, O-DU, O-CU, near-RT RIC, non-RT RIC, SMO, xApp, rApp, open fronthaul, eCPRI, E2, O1, O2, A1, R1, Y1, SBOM and zero trust architecture.

- Related P1 blog posts: exploring O-RAN security; comprehensive guide to Open RAN; traditional RAN vs Open RAN; Kubernetes in mobile networks.

About P1 Security

P1 Security is a France-based specialist in mobile network and telecom security. It works with mobile operators, governments, regulators and network equipment vendors worldwide, and has been securing operators' and nation-states' critical mobile infrastructure since 2011. Coverage spans SS7, SIGTRAN, Diameter, GTP, SIP, IMS, VoLTE, LTE, 5G Core and SBA, RAN, O-RAN, Cloud RAN, NFVI, roaming and private 5G.

The work is built around three pillars: PTA to identify exposure, PTM to detect signalling attacks in real time, and VKB to understand which vulnerabilities matter. P1 Security is ISO/IEC 27001:2022 certified and a GSMA member.

Explore the P1 Security ecosystem

- p1sec.com: products (PTA, PTM, VKB), services, industries, blog, webinars and advisories.
- SecurityHub: telecom security reference and education, home of the TelcoSec glossary.
- TelcoSec Arsenal: curated tools for telecom security practitioners.
- TelcoASM: confidential telecom attack-surface self-assessment and benchmarking.
- ApexThreats: telecom threat intelligence built from signalling data.
- CVE and security advisories: P1 published CVEs and coordinated disclosures.
- Online training: video theory, hands-on labs and certification.

Talk to us

Ready to look at your own exposure? Write to us at contact@p1sec.com, or benchmark your network with TelcoASM.

Follow P1 Security

- LinkedIn: linkedin.com/company/p1security
- Mastodon: @p1sec on Infosec Exchange
- X: @p1security

Legal and copyright

This document is provided for information and research purposes in good faith, with no warranty of any kind. All findings are anonymised and no client or vendor is named without their consent. Where any CVE or vulnerability is referenced, P1 Security follows a responsible disclosure process with a 180-day vendor window; see cve.p1sec.com.

P1 Security S.A.S., Paris, France.