

Cloud RAN security: the new RAN attack surface

The fronthaul, F1 and E1, and Kubernetes threats P1 Security finds in testing, and how to mitigate them.

By Kye Grundy and Laurent Ghigonis, P1 Security.

Contents

1. Introduction: why Cloud RAN changes the RAN security picture	5
2. Cloud RAN defined: C-RAN, vRAN, Cloud RAN and O-RAN	6
3. The disaggregated, virtualised RAN architecture and its attack surface	8
4. How P1 Security tests Cloud RAN	9
5. Threat classes we see in testing	9
5.1 The fronthaul and forgotten paths from the RU	10
5.2 Legacy fronthaul management and the single-packet DU crash	11
5.3 The "new" F1 and E1 interfaces	12
5.4 Over-the-air paths to the virtualised CU	14
5.5 OAM and the management plane	15
5.6 Virtualisation: Kubernetes, SmartNICs and cluster escalation	16
6. Backhaul, physical security and monitoring	17
7. Cloud RAN risk view across 4G and 5G RAN	18
8. A mitigation framework for Cloud RAN	20
9. Operator and vendor readiness checklist	22
10. Frequently asked questions	23
11. How P1 Security helps	24
12. References and further reading	25

Executive summary

Cloud RAN security starts from an awkward inheritance. The radio access network used to be one of the harder parts of a mobile network to attack: proprietary, purpose-built and physically distributed, it took specialised equipment and skills to reach, so exploitation was less feasible rather than simply ignored. Cloud RAN changes that. By moving RAN functions onto virtualised, cloud-native infrastructure and splitting the old baseband unit into distributed and central units that talk over standard interfaces, operators gain flexibility and cost efficiency. They also inherit a modern, IT-style attack surface on top of telecom protocol stacks that were designed for a more trusting world, one that is far more accessible to capable attackers than the purpose-built RAN it replaces.

This whitepaper sets out what P1 Security's team sees when it tests Cloud RAN deployments. The theme will be familiar from our 5G core work: the most damaging issues are seldom novel. They come from interfaces that trust whatever connects to them, protocol stacks that were never hardened against malformed input, physically exposed cell sites, management planes with weak or default credentials, and Kubernetes clusters where a single compromised pod can become control of the whole cluster. Virtualisation does not remove the old RAN problems. In most of the environments we assess, it adds new ones alongside them.

We group the recurring findings into a set of threat classes drawn directly from our testing and research:

1. Fronthaul and the forgotten paths from the radio unit, including legacy, unauthenticated fronthaul management protocols that can crash a virtualised distributed unit with a single malformed packet.
2. The "new" F1 and E1 interfaces, where a rogue distributed unit or a rogue control/user-plane node can reach protocol stacks in the centralised unit that used to be far harder to touch.
3. Over-the-air paths from the handset toward a virtualised centralised unit, an advanced and largely research-stage scenario that Cloud RAN makes cheaper to attempt.

4. The operations, administration and maintenance (OAM) plane, still one of the most productive places an attacker lands and moves laterally.
5. Virtualisation itself: Kubernetes and OpenShift misconfiguration, SmartNIC and accelerator-card exposure, and pod-to-cluster escalation.

Throughout, we are deliberately careful about impact. Where we describe a crash, a denial of service or an escalation, it is something demonstrated under lab conditions or observed during authorised testing, not a claim about live subscriber harm on any specific network. The worst-case ratings we cite, such as a wide-area outage from a single compromised site, are drawn from P1's cross-vendor audit experience and are presented as what could happen if the weaknesses line up, not as what is happening in the wild.

The closing message is practical. Most of the protections that matter already exist. Fronthaul and midhaul link protection, mutual authentication between RAN network functions, disciplined Kubernetes hardening, real network segmentation, physical-security correlation and continuous robustness testing are all available today. The problem is that they are frequently optional, licensed, hard to operate, or simply skipped. New architecture does not secure insecure operations. It only gives them a more modern wrapper.

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 baseline.

Key takeaways

- A single malformed packet on a legacy, unauthenticated fronthaul management protocol can crash a virtualised DU, producing a local, cell-level denial of service.
- With default configurations the F1 interface often requires no authentication, so an attacker on the midhaul can stand up a rogue DU that the CU accepts and feed crafted messages into its F1AP, GTP-U, RRC and PDCP stacks.
- The E1 interface between CU-CP and CU-UP is typically unauthenticated inside the Kubernetes networking stack; abuse can range from denial of service to disclosure of user-plane keys and manipulation of bearers.
- Kubernetes and OpenShift misconfiguration turns one compromised pod into cluster-wide control via service-account token theft or hostPath mounts, so every network function sharing a cluster shares one blast radius.
- OAM stays the most productive foothold: default credentials and injection flaws in NMS web and API services lead to remote code execution and lateral movement across the management plane.
- The mitigations that matter already exist (fronthaul MACsec/802.1X, midhaul and backhaul IPsec, mutual authentication on F1/E1, least-privilege RBAC, segmentation, continuous fuzz testing) but are frequently optional, licensed, or left switched off.

1. Introduction: why Cloud RAN changes the RAN security picture

Cloud RAN changes the RAN security picture by replacing proprietary, hard-to-reach base stations with virtualised, standardised, cloud-native functions, giving the RAN a modern IT-style attack surface on top of the telecom protocol weaknesses it already carried.

For most of mobile networking history, the RAN was a black box. A base station was a proprietary appliance from a single vendor, deployed at a cell site, connected to the core over a backhaul link, and largely opaque to outsiders. Attacking the RAN meant understanding undocumented protocols and, often, physically standing next to

expensive radio hardware. The economics did not favour the attacker, so the RAN received less security scrutiny than the core, the interconnect or the internet-facing edge.

Cloud RAN removes several of those barriers at once. The monolithic baseband unit is split into a distributed unit and a centralised unit. Those functions are virtualised as containerised network functions running on commercial off-the-shelf (COTS) servers, orchestrated by Kubernetes or OpenShift, and connected by standardised interfaces. The radio unit at the cell site now talks to a distributed unit over a packet-based fronthaul rather than a rigid proprietary link. The result is a RAN that, from an attacker's point of view, looks much more like modern web and cloud infrastructure than like traditional telecom equipment, while still carrying the operational weight and criticality of national mobile connectivity.

That shift matters for three reasons. First, the interfaces between RAN functions are now standardised and reasonably well documented, so an attacker no longer needs proprietary knowledge to understand what a message should look like or how to malformed it. Second, the software runs on general-purpose x86 Linux and cloud-native platforms, which are cheaper to study, cheaper to attack, and share a large body of public exploitation technique with the rest of IT. Third, centralisation concentrates risk: a centralised unit can serve many distributed units, and a cluster can host many network functions, so a single successful compromise has a larger blast radius than the same compromise in a distributed, appliance-based RAN.

None of this makes Cloud RAN a bad choice. The benefits are real and the industry is moving in this direction regardless. But it does mean the RAN can no longer be treated as the quiet, low-priority corner of the network. This whitepaper looks at the specific, ground-level weaknesses P1 Security finds in Cloud RAN deployments, why virtualisation and cloud-native operations amplify them, and what operators and vendors can do to close them.

2. Cloud RAN defined: C-RAN, vRAN, Cloud RAN and O-RAN

Cloud RAN is the cloud-native evolution of vRAN: virtualised RAN functions deployed as containerised network functions on Kubernetes or OpenShift in a disaggregated CU/DU architecture. The terminology in this space is genuinely confusing, and opinions differ on where one term ends and the next begins. Before discussing threats it helps to state the definitions we use in this paper.

Traditional (distributed) RAN. A proprietary base station where the radio unit (RU) and the baseband unit (BBU) sit together at the cell site, typically as single-vendor hardware and software.

C-RAN (centralised RAN). An architecture that allows a RAN site to be split up and distributed between locations, moving baseband processing away from the antenna and centralising it, for example in a "BBU pool" or a central office. C-RAN was first explored publicly by China Mobile in its 2011 paper "C-RAN: The Road Towards Green RAN", which centralised baseband processing (BBU pooling). The formal distributed-unit and centralised-unit (gNB-DU / gNB-CU) split that modern Cloud RAN relies on came later, standardised by 3GPP in Release 15. Importantly, C-RAN describes where functions run and how they are split. It does not mandate virtualisation.

vRAN (virtualised RAN) and Cloud RAN. vRAN is the practice of running RAN functions as software on virtualised, general-purpose infrastructure rather than on purpose-built appliances. Cloud RAN is its cloud-native evolution, running those functions as CNFs in a Kubernetes or OpenShift cluster. Cloud RAN does not have to follow a centralised C-RAN topology, but it makes centralisation practical, which is why C-RAN and Cloud RAN are now often used interchangeably. In this paper we use a working definition that reflects what we actually see in the field:

Cloud RAN today = virtualised RAN functions (CNFs) deployed in a C-RAN-style, disaggregated architecture.

O-RAN. Open RAN, as defined by the O-RAN Alliance, is a further step: it standardises open, multi-vendor interfaces (the Open Fronthaul with its control, user, synchronisation and management planes, plus the E2, A1 and O1 interfaces) and adds intelligence layers such as the Non-Real-Time and Near-Real-Time RAN Intelligent Controllers. O-RAN is usually deployed cloud-native, so it overlaps heavily with Cloud RAN, but the two are not the same thing. A Cloud RAN deployment can be built on a single vendor's virtualised software without any of O-RAN's open interfaces, and an O-RAN deployment layers additional open interfaces and new components on top of the virtualised base.

The practical consequence is that Cloud RAN and O-RAN share a common virtualised, disaggregated foundation and therefore a large part of their attack surface. Where O-RAN adds its own open interfaces and RIC-based intelligence, it adds its own surface as well, and the O-RAN Alliance has published a security framework to address it. This paper focuses on the Cloud RAN foundation that most operators are deploying now, and draws on O-RAN work where it is directly relevant.

3. The disaggregated, virtualised RAN architecture and its attack surface




Figure 1. The Cloud RAN split across fronthaul, midhaul and backhaul, running on shared cloud-native infrastructure

A Cloud RAN deployment typically breaks the old base station into three logical pieces connected by three logical link domains.

The **radio unit (RU)** stays at the cell site with the antenna. The **distributed unit (DU)**, the gNB-DU in 3GPP terms, handles the lower and middle layers of the stack and is increasingly virtualised (vDU), sometimes co-located at the cell site and sometimes further back. The **centralised unit (CU)**, the gNB-CU, handles the higher layers and is commonly virtualised (vCU) and hosted centrally, potentially in a data centre. In deployments that use the E1 interface, the CU is split further into a control-plane part (CU-CP) and a user-plane part (CU-UP).

These functions communicate over three domains:

- **Fronthaul (FH):** between the RU and the DU, carrying radio in-phase/quadrature (IQ) samples and signalling using protocols such as eCPRI (enhanced Common Public Radio Interface), and, on legacy paths, CPRI or OBSAI, alongside management and timing traffic. Older deployments and vendor-specific extensions also carry legacy, often unauthenticated management protocols.
- **Midhaul:** between the DU and the CU, carrying the F1 interface, and where the CU is split, the E1 interface between CU-CP and CU-UP. The midhaul is typically exposed through a cell-site router (CSR).

- **Backhaul:** between the CU and the mobile core, and the OAM and NFV infrastructure links used to manage everything.

Two things drive the expanded attack surface. The first is that these interfaces, previously buried inside a single appliance, are now carried over networks that can be reached from a cell site, from an adjacent compromised network, or from a virtual networking stack. The second is that the software now runs as CNFs on a shared, cloud-native platform, so the security of the RAN depends on the security of Kubernetes, the host nodes, the SmartNICs and accelerator cards that provide packet acceleration, and the OAM and NFVi management planes that orchestrate them.

A useful way to think about it is that Cloud RAN inherits three overlapping problem sets at the same time: the physical and protocol weaknesses of the RAN, the cloud-native and orchestration weaknesses of Kubernetes-based infrastructure, and the management-plane weaknesses that have always plagued telecom OAM. The sections below walk through the threat classes where those problem sets meet.



4. How P1 Security tests Cloud RAN

The findings in this paper come from hands-on work on real Cloud RAN environments, not from a standards review. Our method combines telecom-specific penetration testing, product and protocol analysis, physical red-team engagements at cell sites, and controlled lab reproduction of the equipment and interfaces in scope. We test from the perspectives an attacker actually uses: physical access at the edge, an adjacent compromised network, and the internal cloud-native platform.

Everything below is described conceptually. Where we mention a crash, a denial of service or an escalation, it is something we have reproduced under lab conditions or observed during authorised testing, with no payloads or reproduction detail published.



5. Threat classes we see in testing

As with our 5G core work, the most impactful issues here stem less from the standards themselves and more from implementation flaws, insecure configuration, weak physical security and missing foundational hygiene. We describe demonstrated behaviour in conceptual terms only.

5.1 The fronthaul and forgotten paths from the RU

The weakest point of any RAN deployment is the cell site itself, and that is not unique to Cloud RAN. Cell sites are numerous, geographically dispersed, maintained by a growing number of subcontractors, and expensive to physically secure. Locks, alarms and CCTV are inconsistent at best, and rolling out physical fixes across thousands of sites is costly. Historically most attacks on cell sites have been destructive or motivated by theft of raw materials. The more interesting risk for a security team is the silent compromise: an attacker who reaches the fronthaul link, either through the fronthaul network or through physical ports, and uses it to move laterally into more trusted parts of the network.

Several trends have made the fronthaul a more prominent attack surface than it used to be:

- The move from CPRI to the packet-based eCPRI makes L2/L3 layers easier to access with ordinary networking tools.
- Better public knowledge of the protocols means an attacker no longer needs proprietary documentation.
- Increased responsibility pushed onto the RU, driven by performance demands, encourages shortcuts in security.
- General-purpose hardware needs acceleration components (SmartNICs and data processing units, DPUs) to keep up, and those components add their own attack surface.

The fronthaul also carries a messy mixture of protocols, including some that are effectively legacy and barely maintained. A cell site with weak physical security can expose an old, under-maintained management protocol that has almost no public presence any more, running on hardware that was never designed to resist a hostile network. In testing we have seen fronthaul-adjacent services that offer, for example, an unauthenticated root shell on RU or DU hardware, simply because the design assumed nobody would ever be positioned to reach them.

The fronthaul typically exposes three kinds of service toward an RU: the eCPRI, CPRI or OBSAI transport of user data and signalling to and from the radio domain; management traffic such as NETCONF over SSH for RU fault, configuration, accounting, performance and security (FCAPS) management; and timing distribution using PTP or SyncE. Each of these can be attacked for destructive effect, for information disclosure, or as a stepping stone for lateral movement. The recurring pattern is that these protocols were treated as an afterthought, protected mainly by the assumption of physical isolation, and that assumption no longer holds.

A representative fronthaul attack chain (as observed in assessments and lab work):

1. An attacker gains access to a cell site with weak physical security.
2. From there they reach the fronthaul link, through the fronthaul network or via physical ports.
3. Exposed services in the RU, DU or OAM are reachable, often with no protection or with default administrative credentials.
4. A non-robust service can be crashed with a malformed message, degrading service, or the access can be used to compromise OAM services and move laterally upstream.

What to verify

- Is fronthaul link protection (for example MACsec or 802.1X) actually enabled, not just supported by the vendor?
- Are legacy and vendor-specific fronthaul management protocols inventoried, and are their listening services authenticated?
- Are RU and DU management interfaces free of default and password-less access?

5.2 Legacy fronthaul management and the single-packet DU crash

Some deployments still manage RU-to-DU settings over a legacy, UDP-based control protocol that carries an XML payload and a set of standard procedures alongside vendor-specific extensions, typically with no encryption or authentication. This class of protocol deserves attention because it illustrates the wider Cloud RAN problem so cleanly.

That combination, a complex XML document parsed over UDP by an old parser, embedded in a legacy process that has been "containerised" and dropped into a modern Kubernetes environment, is a robustness failure waiting to happen. A single malformed message can crash the DU process, producing a local, cell-level denial of service. In other words, under lab conditions a single crafted packet can take a cell offline.

The root causes are the point:

- A protocol designed decades ago, never modernised, with a parser that has not been meaningfully tested against malformed input.

- A vendor implementation bug in how that input is handled.
- Poor integration of a legacy process-management model into a cloud-native, containerised environment.
- Misuse of Kubernetes high-availability features, so that a single process crash propagates upward instead of being contained.
- Binaries that are typically not hardened against exploitation.

This is Cloud RAN's core tension in miniature. The underlying weakness is old. The virtualised, cloud-native deployment did not create the fragile parser, but it changed the failure mode: a local process crash became a cascading restart that a badly configured high-availability layer amplified rather than absorbed. The lesson is not that any one protocol is uniquely bad, but that legacy telecom protocol stacks need robustness and fuzz testing before they are wrapped in a container and treated as production-grade cloud-native software.

5.3 The "new" F1 and E1 interfaces




Figure 2. The F1 and E1 interfaces: a rogue DU the CU trusts

The F1 interface connects the DU and CU, carrying F1AP signalling and GTP-U user data, while the E1 interface splits the CU into control-plane (CU-CP) and user-plane (CU-UP) parts using E1AP. In real Cloud RAN deployments, implementations tend to follow the 3GPP Release 15 CU/DU split. F1 runs over the midhaul network by way of a cell-site router, and touches essentially every UE procedure toward the RAN and the core, including registration and RRC security activation.

The problem we see repeatedly is trust. With default security implementations, the CU trusts all signalling received over the F1 link, no authentication is required for a DU to connect, and the cell-site router that fronts the midhaul is often weakly secured. Where that holds, an attacker who reaches the midhaul, either by breaking into a cell site or by pivoting from an adjacent compromised network, can stand up a **rogue DU** toward the CU, which accepts the new DU. In our assessments we have been able to do exactly this. Once connected, the attacker can send crafted messages into the CU's F1AP and GTP-U processing stacks, and can also exercise the RRC and PDCP stacks that a UE would normally drive.

Telecom protocol stacks like these are, in our experience, frequently fragile, with complex message structures and state machines that were not built to resist hostile, malformed input. Under lab conditions both simple and complex malformed messages have been enough to cause a denial of service in the CU. Simple, completely broken messages are easy for a defensive implementation to drop; the harder cases are valid-looking messages that are expensive to encode but that the CU still mishandles. Either way, in testing the attacker does not need to be sophisticated to knock the CU over.

The impact framing here needs care. A single CU can serve many DUs, and through them a large number of subscribers, so a CU outage has a wide blast radius in principle. The worst-case rating in our risk view treats a single-site F1AP compromise as potentially wide-area for some architectures, but that is a worst-case drawn from cross-vendor audit experience, not a statement about any specific live network or about routine real-world exploitation.

The **E1 interface**, introduced above, is less commonly deployed than F1, and where it exists it is usually carried inside the Kubernetes virtual networking stack (for example, over a secondary container-network interface), rather than exposed over a cell-site router, which makes it somewhat more trusted. That trust is the weakness. In most cases we see, the CU-CP and CU-UP blindly process signalling arriving on the E1AP interface with no authentication on either side. An attacker with connectivity to E1 can claim to be the CU-CP or the CU-UP, which opens a range of attacks from purely destructive to quietly observational. Beyond a denial of service against the E1AP processing stack, abuse of the E1AP procedure logic can, in principle, allow:

- monitoring of session status,
- timing analysis to build a pattern of life,
- malicious manipulation of user-plane bearers, and
- disclosure of sensitive material such as the keys that protect user data, or false reporting of a subscriber's usage.

What to verify (F1 and E1)

- Is any authentication or integrity protection enforced on F1 and E1, or does the CU trust anything that connects?
- Is the midhaul link protected (for example with IPsec), and is that protection actually turned on in production rather than merely licensed?
- Are the F1AP, E1AP, GTP-U, RRC and PDCP stacks subjected to robustness and fuzz testing, not just functional testing?
- Is there network segmentation preventing an adjacent compromised network from reaching the midhaul at all?

5.4 Over-the-air paths to the virtualised CU

A more advanced scenario concerns the software stacks that face the handset directly. In a typical Cloud RAN processing flow, the lowest physical-layer functions may run on a SmartNIC, higher physical-layer and lower L2 functions (MAC and RLC) run in the vDU, and the higher layers (PDCP and, at L3, RRC) run in the vCU. Historically, attacking the RRC or PDCP layers over the air required expensive, complex and specialised radio equipment, which kept the scenario out of reach for most attackers.

Cloud RAN lowers that barrier. Because the target software now runs on general-purpose x86 Linux rather than bespoke silicon, it is cheaper for an attacker to study and to attack, and the same over-the-air message that reaches the RRC or PDCP stack can potentially reach code running in a centralised, virtualised CU that may sit in a data centre, which increases the potential impact if the stack can be crashed or exploited.

We are deliberately cautious here. Over-the-air compromise of a virtualised CU from a handset is an advanced scenario that requires expensive and specialised capability, and it is an active research topic for our team rather than something we present as routinely observed in the wild. The usual operator strategy for this class of threat is to contain a compromise rather than assume it can be fully prevented, and it is genuinely hard to justify defensive investment against a threat this specialised. We raise it because Cloud RAN measurably changes the economics of the attack, and because the direction of travel, cheaper hardware targets and more centralised impact, is exactly the sort of shift that turns a theoretical attack into a practical one over time.

5.5 OAM and the management plane

The presence of an OAM plane is not specific to Cloud RAN, but it remains one of the most important attack surfaces and one of the most productive places an attacker lands. RAN network functions are typically connected to a network management system (NMS) that lets administrators act on them over protocols such as SSH, SNMP and NETCONF. At the cell site, the RU commonly reaches OAM through the DU.

An NMS is a complex collection of services: web interfaces, command-line interfaces, KPI and metrics pipelines, logging, databases, registries, FCAPS functions and links into the NFV infrastructure. In a CNF-based Cloud RAN, the complexity multiplies. There are hundreds to thousands of traffic flows within and between clusters, plus as many CNFs and related Kubernetes resources such as pods, stateful sets, config maps and policies.

With that many moving parts, mistakes in network segmentation and service configuration, including roles, accounts and authentication, are easy to make and hard to catch. The result is often a myriad of services, some of them ones the operator had forgotten were there, exposed and listening.

The recurring OAM attack pattern is depressingly consistent. An attacker who has compromised a RAN network function communicates with the OAM services that function is allowed to reach. Weak or missing network segmentation lets them cross into a more trusted internal OAM zone. There they find the usual problems: default credentials, weak input sanitisation, and web or API services vulnerable to injection that can lead to remote code execution. From there they compromise administrative accounts, gain access to highly sensitive NMS functionality and move laterally. OAM abuse also runs in the other direction: legitimate OAM functionality, once controlled, can be used maliciously to push configuration changes that degrade the RAN or to exfiltrate sensitive information such as subscriber data.

A representative OAM attack chain (as observed in assessments):

1. An attacker breaks into a cell site with weak physical security.
2. They reach the backhaul or management link, for example via the cell-site router.
3. They communicate with upstream OAM systems through a management service exposed toward the RAN, such as an API or an SNMP endpoint.
4. That service is unsecured or vulnerable, for example an injection flaw in a web UI element leading to remote code execution.
5. The result is compromise of OAM administrative accounts and lateral movement across the management plane.

5.6 Virtualisation: Kubernetes, SmartNICs and cluster escalation

Cloud RAN's defining characteristic, running network functions as CNFs on Kubernetes or OpenShift, is also where a contained problem can become a catastrophic one. The platform brings an explosion of complexity, with thousands of Kubernetes resources to configure correctly, and securing it is a distinct skill that traditional RAN teams have had to acquire quickly, often while just trying to get the system operational. Security tends to be the last skill learned, not the first.

The escalation pattern we see mirrors what we describe in our 5G core work. Network functions and OAM components run inside pods that frequently carry insecure configurations, and a remote code execution in a single pod can become full cluster compromise through well-known cloud-native techniques:

- **Service-account token theft.** If a compromised pod's service account has broad rights (the informal test is whether it can effectively do anything in the cluster), stealing its token gives the attacker cluster-wide control.
- **Host filesystem access via hostPath.** If the account running the compromised service is permitted to mount host paths, a pod can mount the underlying node's filesystem and, from there, steal the node's kubeconfig and pivot to the host.

Kubernetes role-based access control (RBAC) is often not configured to best practice, and where domain-level tokens are in use, a single stolen token can hand an attacker complete infrastructure control. Pods running in privileged mode with host mounts make it straightforward to gain a foothold on the nodes that run the majority of the network functions. And because the OAM IT nodes that manage the cluster are frequently linked to telecom equipment, a compromise there commonly leads to lateral movement into the telecom domain.

The Kubernetes and OpenShift control plane is itself a target: if it is not adequately secured, an attacker who reaches it gains complete control over the cluster and, with it, every network function the cluster hosts. This is the same amplification we warn about in the 5G core. When multiple critical functions share a cluster, taking the cluster takes all of them at once.

SmartNICs, DPUs and accelerator cards deserve specific attention because they are easy to overlook. General-purpose hardware needs these components to match the performance of legacy purpose-built equipment, and they are increasingly common in virtualised telecom networks. They are not dumb hardware. They run small operating systems and expose their own services on constrained hardware, which encourages security shortcuts, and they process network packets with custom implementations that may be fragile. Because designers assume nobody would try to

reach them, security is sometimes not considered at all. Examples we have seen in the field include a configuration shell with no authentication on a physically exposed serial port, and password-less root SSH left open. The management links for the cloud platforms themselves add another layer: dedicated new platforms for cloud management introduce IT-style risks directly into the operational-technology world of the RAN.

What to verify (virtualisation)

- Is Kubernetes/OpenShift RBAC scoped to least privilege, with no broad domain-level tokens that unlock the cluster if stolen?
- Are privileged pods and hostPath mounts eliminated or tightly justified and monitored?
- Is the cluster control plane hardened and segmented from the telecom functions it hosts?
- Are SmartNICs and accelerator cards inventoried, with their exposed management services authenticated and their serial and SSH access locked down?

6. Backhaul, physical security and monitoring

Backhaul links are frequently under-protected and cell-site physical security is not affordable at scale, so the realistic defence for both is detection and correlation rather than prevention.

Two cross-cutting issues sit alongside the threat classes above.

The **backhaul** between the CU and the core is not specific to Cloud RAN, but it is frequently under-protected. Encryption, where present, is often applied to the control plane and not the user plane. Physical access to a backhaul link can be enough to intercept or manipulate traffic: an inexpensive fibre switch is sufficient to perform an on-site man-in-the-middle on a backhaul link, giving an attacker a path toward the core and the OSS. Lack of filtering and link protection on the backhaul is one of the higher-impact issues we track, precisely because it points at the crown jewels.

Physical security is the uncomfortable foundation under most of this. RAN deployment and maintenance are hugely expensive, so security is often what is left over. Locks are inconsistent, the number of subcontractors keeps rising, the number of sites is large, physical fixes are expensive to roll out, and realistic CCTV coverage

generally does not exist. Compounding this, public and regulator-facing open data sources can reveal cell-site locations, which lowers the effort an attacker needs to plan an approach.

Since perfect physical security is not affordable, the realistic improvements are about detection and correlation rather than prevention:

- analysing unusual events from existing firewalls,
- correlating planned operations-team interventions with alarms,
- correlating physical alarms with firewall events, and
- focusing scarce resources on the highest-value sites.

The RAN is also known to have weaker monitoring than other parts of the network, so investment in RAN-specific alarms and network-event detection pays off disproportionately.

7. Cloud RAN risk view across 4G and 5G RAN




Figure 3. Where Cloud RAN risk concentrates: relative impact and likelihood across the threat classes

Pulling the threat classes together, it helps to look at them as a rough risk picture rather than a single ranked list. The relative ratings below are based on P1 Security's experience across multiple vendors and worldwide audit results, and they express potential impact and likelihood, not observed activity on any particular network.

- **Exploitation of F1AP vendor bugs.** Potentially the widest-area impact, since a single physically compromised site can, in a worst case, translate into a wide-area RAN denial of service or exploit path. Not present in traditional integrated RAN, since the F1 stack only becomes reachable once the CU and DU are split, so it applies to Cloud RAN and to other disaggregated RAN deployments that share the same split. Rated high.
- **Compromise of a RAN site shelter.** Local RAN site compromise with potential to propagate. Relevant to both traditional and Cloud RAN. Rated high, driven by weak physical security.
- **Exploitation of over-the-air RRC/PDCP vendor bugs.** Local RAN site denial of service or compromise originating from the radio side. Relevant to both architectures, but cheaper to attempt against Cloud RAN. Rated high in principle, though advanced to execute.
- **Compromise of OSS or cluster management.** Potentially wide-area impact from a single entry point, because cluster control means control of the functions the cluster hosts. Specific to Cloud RAN. Rated medium, reflecting the effort required to get there.
- **Exploitation of DU eCPRI or fronthaul management protocol bugs.** Local RAN site denial of service or compromise from the antenna side, as in the single-packet DU crash. Relevant to both architectures. Rated medium.
- **Lack of filtering and link protection on the backhaul.** Higher impact because it enables attacks toward the core and the OSS. Relevant to both architectures. Rated medium.

Taken together, the highest-impact items, exploitation of F1AP vendor bugs, compromise of a RAN site shelter, and over-the-air RRC or PDCP bugs, are where testing and hardening effort should concentrate first.

8. A mitigation framework for Cloud RAN




Figure 4. Defence in depth across the Cloud RAN, from fronthaul to core

The protections that matter mostly already exist. The recurring problem is that they are optional, licensed, operationally awkward, or simply not switched on. The framework below is organised from foundational hygiene to proactive operations and maps directly onto the threat classes above.

Tier 1: link and interface protection

- **Fronthaul protection.** Where the deployment supports it, enable MACsec or 802.1X on the fronthaul. Outside O-RAN Alliance deployments this is not always well defined, and in practice it is either a proprietary solution or nothing at all, which leaves an easy path for an attacker at a weak cell site. Treat "no fronthaul protection" as a finding, not a default.
- **Midhaul and backhaul protection.** IPsec is generally mandatory for a vendor to implement, but implementing it is not the same as an operator operating it. In practice it is often not enabled, because PKI, certificate lifecycle and troubleshooting are harder, and because the capability may sit behind an additional licence. The awkward but necessary question is who decides that a given environment is "secure enough" to run without it. High-throughput RU-to-DU and DU-to-CU links in particular often lack even off-the-shelf integrity protection. Enable link protection end to end and verify it is actually in operation.
- **Internal mutual authentication.** Require authentication and integrity on F1 and E1 rather than letting the CU trust anything that connects. Mutual TLS on internal RAN

and OAM interfaces, as part of a zero-trust posture, turns "the CU accepts any DU" into "the CU accepts only DUs it can cryptographically verify".

Tier 2: platform, credential and segmentation hygiene

- **No default or well-known credentials.** Change or remove every default and password-less access on RUs, DUs, CUs, SmartNICs, OAM systems and management interfaces as a mandatory, verified post-deployment step. This is the single most common enabler in our findings.
- **Kubernetes and NFVi hardening.** Follow cloud-native security best practice: least-privilege RBAC with no broad domain-level tokens, Pod Security Standards enforced through admission control, elimination of privileged pods and unnecessary hostPath mounts, a hardened and segmented control plane, and awareness of which network functions share a cluster and therefore share a blast radius.
- **Network segmentation.** Enforce strict filtering between the RAN, the OAM, the IT and the core domains, and between clusters, so that a foothold at a cell site or in an adjacent network cannot reach the midhaul, the OAM plane or the cluster control plane unimpeded.

Tier 3: robustness testing, vendor engagement and detection

- **Robustness and fuzz testing.** Subject the fronthaul management protocols, F1AP, E1AP, GTP-U, RRC and PDCCP stacks to continuous robustness and fuzz testing, not only functional testing. The single-packet vDU crash and the F1/E1 CU crashes are robustness failures, and robustness failures are found by robustness testing.
- **Vendor engagement and bug containment.** Test new products for security issues before and during deployment, encourage vendors to fix not just individual bugs but whole bug categories, and, where bugs cannot be immediately fixed, contain them: harden the system and toolchain against exploitation, and test availability behaviour so that a single process crash does not cascade into a cell-level outage through misused high-availability logic.
- **Detection, correlation and RAN monitoring.** Accept that physical prevention is unaffordable at scale and invest in detection instead. Deploy network intrusion detection, correlate firewall anomalies, planned interventions and physical alarms, prioritise high-value sites, and close the RAN's well-known monitoring gap with RAN-specific alarms and network-event detection. The better the hardening and segmentation, the more noise an attacker must make to move, and the more likely detection becomes.

9. Operator and vendor readiness checklist

A condensed, practical checklist drawn from the sections above.

- ☐ Fronthaul protection (MACsec or 802.1X) enabled where supported; "no protection" treated as a finding.
- ☐ Midhaul and backhaul link protection (IPsec or integrity protection) enabled end to end and verified in operation, not merely licensed or supported.
- ☐ Authentication and integrity enforced on F1 and E1; the CU does not trust arbitrary DUs or CU-CP/CU-UP peers.
- ☐ Legacy and vendor-specific fronthaul management protocols inventoried, authenticated, and robustness-tested.
- ☐ No default, well-known or password-less credentials on RU, DU, CU, SmartNIC, OAM and management interfaces.
- ☐ SmartNICs, DPUs and accelerator cards inventoried; exposed serial, shell and SSH access locked down.
- ☐ Kubernetes/OpenShift hardened: least-privilege RBAC, no broad domain-level tokens, no unnecessary privileged pods or hostPath mounts, hardened and segmented control plane.
- ☐ Awareness of which network functions share a cluster, and of the blast radius if that cluster is compromised.
- ☐ Strict segmentation and filtering between RAN, OAM, IT, core and inter-cluster domains.
- ☐ Continuous robustness and fuzz testing of fronthaul management, F1AP, E1AP, GTP-U, RRC and PDCP stacks.
- ☐ High-availability behaviour tested so a single process crash is contained, not amplified.
- ☐ Vendor security testing and disclosure engagement in place, with a focus on fixing bug categories and containing what cannot be fixed immediately.
- ☐ Physical-security detection and correlation: network IDS, firewall-anomaly analysis, alarm correlation, and prioritisation of high-value sites.
- ☐ Regular, telecom-specific penetration testing and product security review across fronthaul, midhaul, OAM, virtualisation and physical perspectives.

10. Frequently asked questions

What is Cloud RAN and how does it differ from O-RAN?

Cloud RAN is the cloud-native evolution of vRAN, running virtualised RAN functions as containerised network functions on Kubernetes or OpenShift in a disaggregated CU/DU architecture. O-RAN adds standardised open multi-vendor interfaces and RIC intelligence on top, so the two overlap but are not the same.

What are the main Cloud RAN security threats?

Fronthaul and legacy management protocols (for example a malformed packet crashing a virtualised DU), unauthenticated F1 and E1 that let a rogue DU or CU-CP/CU-UP reach the CU, over-the-air paths to a virtualised CU, OAM compromise, and Kubernetes misconfiguration enabling pod-to-cluster escalation.

Why are the F1 and E1 interfaces risky in Cloud RAN?

With default implementations the CU often trusts any DU that connects, with no authentication and a weakly secured cell-site router, so an attacker on the midhaul can stand up a rogue DU and send crafted messages into the CU's F1AP, GTP-U, RRC and PDCP stacks.

How do you secure a Cloud RAN deployment?

Enable fronthaul (MACsec or 802.1X) and midhaul and backhaul (IPsec) link protection, enforce mutual authentication on F1 and E1, remove default credentials, harden Kubernetes RBAC and remove privileged pods, segment the RAN, OAM, IT and core domains, and continuously fuzz-test the protocol stacks.

Why are legacy fronthaul management protocols a Cloud RAN risk?

Some deployments still manage RU-to-DU settings over a legacy, UDP-based control protocol that carries an XML payload typically with no encryption or authentication. A single malformed message can crash the DU process and take a cell offline, a robustness failure amplified when a legacy process is containerised and dropped into a cloud-native environment with misused high-availability logic.

What is a rogue DU attack?

With default F1 security implementations the CU requires no authentication for a DU to connect, so an attacker who reaches the midhaul, by breaking into a cell site or pivoting from an adjacent compromised network, can stand up a rogue DU that the CU accepts. Once connected, the attacker can send crafted messages into the CU's F1AP and GTP-U stacks and exercise the RRC and PDCP stacks a UE would normally drive, which under lab conditions has been enough to cause a denial of service in the CU.

Is Cloud RAN less secure than traditional RAN?

Cloud RAN is not a bad choice, but it widens the attack surface. Virtualisation does not remove the old RAN problems; in most environments P1 assesses it adds new ones alongside them, inheriting the physical and protocol weaknesses of the RAN plus the cloud-native, orchestration and management-plane weaknesses of Kubernetes-based infrastructure. Centralisation also concentrates risk, so a single compromise has a larger blast radius than in a distributed, appliance-based RAN.

How does Kubernetes misconfiguration escalate to full cluster compromise?

Network functions run inside pods that frequently carry insecure configurations, so a remote code execution in a single pod can become full cluster compromise. Common paths are service-account token theft, where a pod's broadly privileged token grants cluster-wide control, and hostPath mounts, where a pod mounts the node's filesystem and steals its kubeconfig to pivot to the host. Because multiple critical functions often share a cluster, taking the cluster takes all of them at once.

11. 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 findings in this paper come directly from our offensive work and research. The services and products that map to them:

- **Audit, vulnerability assessment and red team.** Telecom-specific penetration testing, product analysis and physical red-team engagements across the RAN, from the fronthaul and cell site through F1 and E1, OAM and the NFV infrastructure, tested from the physical, adjacent-network, and internal perspectives an attacker actually uses.
- **Product security review for vendors.** Independent, telecom-specific security testing of RAN products and network functions, so vendors can find and contain implementation and robustness flaws, including the protocol-stack fragility and cloud-native misconfigurations described here, before they ship. This complements formal equipment-security assurance under GSMA NESAS and 3GPP SCAS, which set a baseline rather than an adversarial test.
- **PTA (P1 Telecom Auditor).** Attack-surface management and telecom-protocol-aware vulnerability scanning and compliance checks across the telecom estate, so

RAN and core exposure is identified with remediation guidance rather than guessed at.

- **PTM (P1 Telecom Monitor).** Real-time intrusion detection for signalling and core networks, so the lateral movement and misuse described here can be detected and acted on, complementing the RAN-specific monitoring and correlation this paper recommends.
- **VKB (Vulnerability Knowledge Base).** The largest telecom vulnerability intelligence database, with a continuous feed from our research team, so teams understand which vulnerabilities affect their network, why they matter and where risk concentrates.
- **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 the RAN and signalling activity described here is monitored and triaged continuously rather than only during an engagement.
- **Training.** P1's online training and hands-on Laboratory Playground take security, network, RAN and SOC teams from telecom fundamentals to advanced Open RAN, cloud-native RAN and virtualisation 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 benchmark or test your own Cloud RAN deployment against the threats in this paper, contact us at contact@p1sec.com.

12. References and further reading

Standards and industry guidance:

- **3GPP TS 38.401:** NG-RAN architecture, including the CU/DU functional split: <https://portal.3gpp.org/desktopmodules/Specifications/SpecificationDetails.aspx?specificationId=3219>
- **3GPP TS 38.470 series:** the F1 interface (F1AP is TS 38.473).
- **3GPP TS 38.460 series:** the E1 interface (E1AP is TS 38.463).
- **3GPP TS 38.331:** Radio Resource Control (RRC) protocol specification.
- **3GPP TS 33.501:** security architecture and procedures for the 5G system.
- **O-RAN Alliance WG4 and WG11:** Open Fronthaul specifications and O-RAN security requirements, published at <https://www.o-ran.org/specifications>

- **eCPRI specification:** the enhanced Common Public Radio Interface (eCPRI) specification.
- **CPRI Cooperation:** the Common Public Radio Interface specification (legacy fronthaul).
- **GSMA NESAS and 3GPP SCAS:** network equipment security assurance and its security assurance specifications: <https://www.gsma.com/solutions-and-impact/industry-services/assurance-services/network-equipment-security-assurance-scheme-nesas/>
- **GSMA FS.31:** baseline security controls.
- **Cloud-native platform hardening references:** CIS Kubernetes Benchmark (<https://www.cisecurity.org/benchmark/kubernetes>), NSA/CISA Kubernetes Hardening Guidance (<https://www.cisa.gov/news-events/alerts/2022/03/15/updated-kubernetes-hardening-guide>) and MITRE ATT&CK for Containers (<https://attack.mitre.org/matrices/enterprise/containers/>).
- **China Mobile Research Institute**, "C-RAN: The Road Towards Green RAN" (2011): the original centralised-RAN concept.

P1 Security resources:

- P1 Security security advisories and coordinated disclosures at cve.p1sec.com.
- SecurityHub TelcoSec glossary entries for RAN, C-RAN, vRAN, Cloud RAN, O-RAN, CU, DU, RU, F1, E1, eCPRI, fronthaul, midhaul, backhaul, GTP-U, F1AP, E1AP, RRC, PDCP, and Kubernetes in mobile networks.
- Related P1 blog posts: [traditional RAN vs Open RAN](#); [exploring O-RAN security](#); [Kubernetes in mobile networks](#); [securing NFVI](#); [fuzzing the mobile core](#); [RAN-to-core compromise red-team lessons](#).

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

© 2026 P1 Security S.A.S. All rights reserved. PTA, PTM and VKB are products of P1 Security.

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.