P1 SECURITY WHITEPAPER

Top 5G core security vulnerabilities

The five weak spots P1 Security finds most in real 5G pentests, and how to close them.

By El Mehdi Regragui, P1 Security.

P1 SECURITY

contact@p1sec.com

Contents

Executive summary

For 5G core security, the move from 5G Non-Standalone (NSA) to 5G Standalone (SA) is not just a performance upgrade. It replaces a familiar telecom core with a cloud-native, service-based architecture (SBA) that speaks HTTP/2 and JSON, runs as containerised network functions on Kubernetes, and exposes a far wider and more web-like attack surface. That surface does not remove the old problems. In most of the assessments we run, it adds new ones on top of them.

This whitepaper distils what P1 Security's team finds when it tests production and pre-production 5G core environments. The uncomfortable headline is that the most damaging issues are rarely exotic. They are basic:

- Weak and default credentials.
- Reused keys.
- Remote code execution through unsafe implementation shortcuts.
- Privilege escalation that turns one container into control of a whole cluster.
- Server-side request forgery that abuses the SBA's own routing logic.
- Denial of service triggered by malformed signalling.

When these combine with weak segmentation and slow patch cycles, they become exactly the entry points that advanced persistent threat (APT) groups and other capable attackers need.

We group the recurring findings into five vulnerability classes:

1. Authentication and access control failures
2. Remote code execution (RCE)
3. Privilege escalation, including Kubernetes pod-to-cluster and pod-to-host escalation
4. Server-side request forgery (SSRF) in the service-based interface (SBI)
5. Denial of service (DoS) in the SBI

Section 5 then examines three emerging risks that mature 5G SA and inter-operator roaming will make more prominent: SSRF reaching across the roaming interconnect through the SEPP, the latent danger of continued HTTP/1.1 support, and the compounding effect of poor software supply-chain visibility and slow telecom patch cycles.

The closing message is practical. Preventing a serious compromise does not require impossibly complex defences. It requires disciplined security hygiene, real network segmentation, internal zero-trust between network functions, continuous robustness testing, and detection and response capability that turns an attacker's lateral movement into noise you can catch. New architecture does not secure insecure operations. It only gives them a more modern wrapper.

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

Key takeaways

- The five recurring 5G core vulnerability classes in P1 Security's SA pentests are authentication and access-control failures, remote code execution, privilege escalation (including Kubernetes pod-to-cluster and pod-to-host), server-side request forgery in the SBI, and denial of service in the SBI.
- The most damaging findings come from implementation flaws and insecure configuration, not from the 3GPP standards themselves, and rarely require a zero-day.
- Because converged dual-mode cores often pack 5G and EPC control-plane functions into one Kubernetes cluster (the AMF and SMF alongside the 4G MME, SGW-C and PGW-C), a single pod-to-cluster escalation can hand an attacker control of all of them at once.
- SSRF against the SBI can turn an internal SCP, or a SEPP across the N32 roaming boundary, into an internal network scanner by manipulating the `3gpp-Sbi-Target-apiRoot` or `Host` header when routing validation is weak.
- OS command injection in custom SSH authentication logic yields unauthenticated RCE. Separately, CVE-2025-32433 (an Erlang/OTP SSH pre-authentication flaw, CVSS 10.0, in CISA's KEV catalogue) shows the real-world stakes for telecom host management.
- Residual HTTP/1.1 backward compatibility re-opens request smuggling and desync attacks, so HTTP/2 should be enforced end to end across the SBI.

1. Introduction: why 5G SA changes the security picture

5G standalone (SA) changes the security picture because it swaps telecom-specific signalling for a cloud-native service-based architecture running HTTP/2, JSON and containerised network functions on Kubernetes, inheriting web and cloud-native risk on top of existing telecom exposure. 5G is arriving in two very different shapes, and the security picture depends on which one an operator is actually running. 5G NSA, which reuses the 4G evolved packet core, is still the dominant real-world deployment. Many operators have not yet built a strong business case for full 5G SA, and a standalone core rollout is genuinely complex. But SA is where 5G is heading, and it introduces a new paradigm of network technologies:

- A cloud-native **service-based architecture (SBA)**, where control-plane network functions such as the AMF, SMF, AUSF, UDM, NRF, PCF, SCP and SEPP communicate as services over the service-based interface rather than through rigid point-to-point interfaces. The **UPF** remains the user-plane function, controlled by the SMF over N4 (PFCP) and carrying subscriber traffic over N3, N6 and N9 rather than over the SBI.
- **HTTP/2** as the transport for the service-based interface (SBI), replacing telecom-specific signalling with a web protocol.
- **JSON** for data serialisation across those interfaces.
- **Containerised network functions** orchestrated by Kubernetes, often with several functions packed into a single cluster.

These choices are what make 5G SA powerful. They also make it look, from an attacker's point of view, a lot more like modern web and cloud infrastructure than like a traditional telecom network. The 5G core inherits a whole legacy of application-layer and cloud-native security risks while still carrying the operational weight and criticality of national mobile infrastructure.

There is a second, quieter problem that runs through almost every engagement we do: the line of security responsibility between the network equipment vendor and the operator is often dangerously blurred. Who removes default credentials after deployment? Who confirms the system was hardened to guidelines? Who owns the third-party components inside a network function? When the answer is unclear, insecure defaults survive into production, and production becomes the real attack surface.

This whitepaper deliberately moves past the high-level threat-landscape overview. Drawing directly from our 5G penetration testing work, it analyses the specific, ground-level vulnerabilities that give adversaries a path into the heart of the 5G core, and what operators and vendors can do about them.

2. The real 5G attack surface

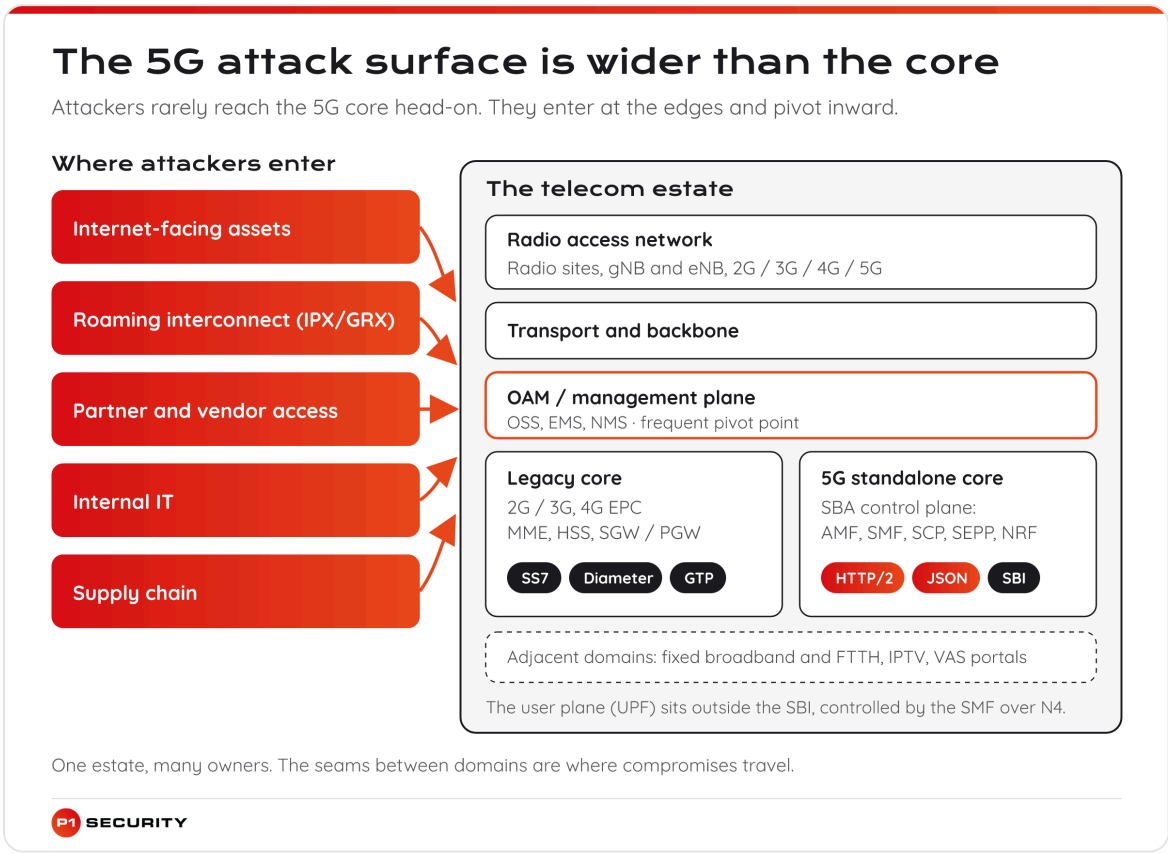

Figure 1. The 5G core sits inside a wider telecom attack surface; attackers enter at the edges and pivot inward

The real 5G attack surface reaches well beyond the core: it spans the RAN, core network, OAM management plane, roaming interconnects, internet-facing assets and adjacent IT, and capable attackers typically enter at the edges and pivot inward. Before looking at individual weaknesses, it helps to be honest about the environment. The 5G core does not exist in isolation. It sits inside a much larger telecom estate, and that whole estate is in scope for a capable attacker.

A typical operator infrastructure exposes several primary domains:

- **Radio access network (RAN):** the gNodeB base stations and radio-site equipment that form the first point of contact for devices, plus the backhaul and fronthaul that connect them.
- **Core network:** the brain of the mobile network, evolved from the 4G EPC (MME, HSS, SGW/PGW) to the 5G core: the control-plane service-based functions such as the AMF and SMF, plus the user-plane UPF.
- **Operations, administration and maintenance (OAM):** the management plane, including network management systems (NMS) and OSS (operations support

system), used to configure and monitor every element.

- **Roaming interconnects:** the IPX and GRX networks that link operators together, a trusted-but-exposed channel that has been abused for cross-network attacks for years.
- **Internet-facing assets:** externally exposed services, staff access portals and VPNs, and partner-operated jump boxes, which are frequently the initial point of entry.
- **Adjacent domains:** fixed broadband and FTTH, IPTV and value-added services, and ordinary enterprise IT, all of which can provide a foothold that eventually pivots toward telecom functions.

The important point is that attackers do not need to walk up to the 5G core and knock on the front door. They come from roaming, from an internet-facing asset, from partner access, from a supply-chain compromise, or from an internal IT path, and then they pivot. 5G security is therefore a subset of a much wider telecom security problem, not a self-contained one.

This is not theoretical. Public threat-intelligence reporting has documented sophisticated, often state-aligned groups active against telecommunications networks, with names such as LightBasin, Liminal Panda, Phantom Taurus and Salt Typhoon among those publicly associated with telecom-targeting activity. These actors have shown a deep understanding of operator infrastructure and a willingness to enter through ordinary enterprise-style paths, such as internet-facing routers and VPNs or roaming interconnects, before establishing long, undetected presence and moving toward far more sensitive telecom functions for intelligence gathering, surveillance and, potentially, disruption. That is the concrete case for testing those paths before an adversary uses them.

3. How P1 Security tests: methodology and perspectives

The findings in this paper come from hands-on penetration testing and red-team work in real 5G environments, not from a standards review or a paper exercise. A few points about method are worth stating, because they explain why the findings look the way they do.

Telecom penetration testing is not IT penetration testing. The protocols, the interfaces (SBI, N32, the OAM plane), the equipment and the failure modes are specific to mobile networks. Generic IT scanners and generic web testing miss most of

what matters here, because they do not understand the service-based architecture, the 3GPP-mandated headers, or the way network functions trust one another.

We test from multiple perspectives, because attackers do:

- A **partner or vendor perspective**, for example through a jump box provided for partner access, which is a common real-world entry point.
- An **internal perspective**, modelling an attacker who has already reached the IT environment and is trying to cross into telecom domains.
- A **roaming or interconnect perspective**, where the threat arrives from another operator across a trusted channel.

We assume breach. The most valuable findings are usually not the single flashy exploit but the chain: how a small initial foothold expands, step by step, into control of a critical function or an entire cluster. Where segmentation is weak and privileges are reused, that chain is short.

The rest of this paper is organised around what those chains repeatedly reveal.

4. The five weak spots: the most common 5G core vulnerability classes

Our findings consistently show that the most damaging vulnerabilities rarely come from the 3GPP standards themselves. They come from implementation flaws, insecure configuration and missing foundational hygiene. The five classes below are the ones we see most often and that carry the most impact.

4.1 Authentication and access control

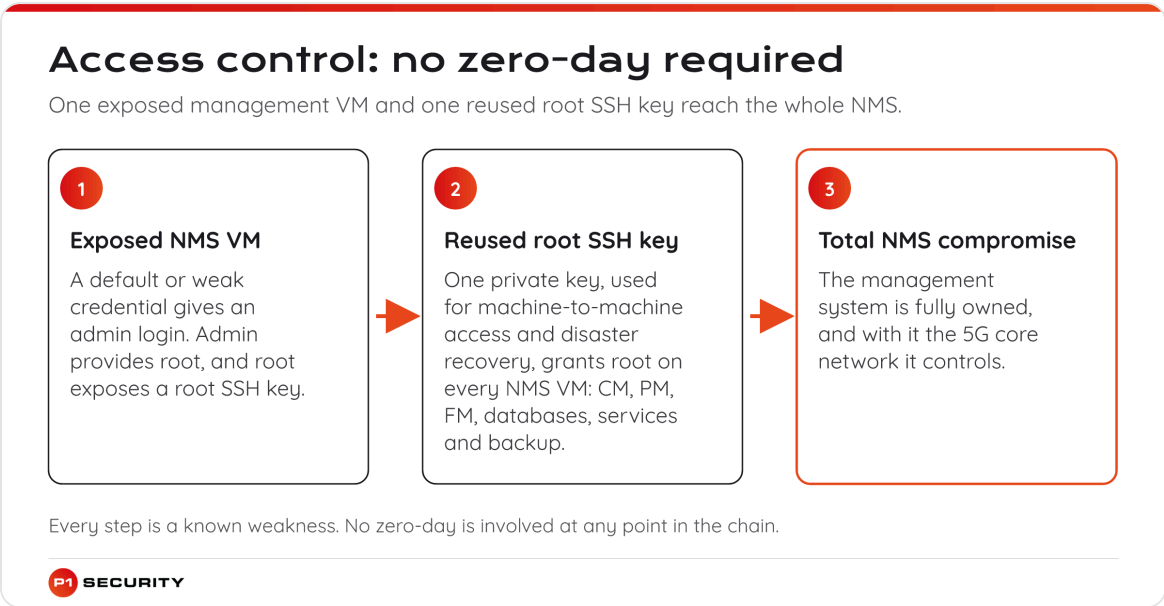

Figure 2. An authentication and access-control kill chain: ordinary credential failures chain from one exposed management VM to the whole NMS and the 5G core behind it

Authentication and access-control failures, default and weak credentials, hardcoded secrets and reused SSH keys are the most common and most effective ways into a 5G core, and they rarely need a zero-day.

Despite everything new in 5G, the most common and most effective attack vectors are still rooted in basic authentication and access control failures. In assessment after assessment we uncover:

- **Default and weak credentials** left on network functions and management systems after deployment.
- **Hardcoded credentials and secrets** embedded directly in application binaries, acting as persistent backdoors the operator often does not know about.
- **Password and private-key reuse**, especially reused SSH private keys used for machine-to-machine communication, and unrestricted root SSH access.

None of this is advanced, and that is precisely the point. Attackers do not need to burn a zero-day when they can simply log in.

These weaknesses are most dangerous when they appear in the management plane, where a small foothold expands quickly. A representative chain we have seen in the field, from an internal (IT environment) perspective:

1. Initial access to an exposed NMS virtual machine through a **default or weak credential**. Logging in as an admin user provides root privileges, which in turn

expose a root SSH private key.

2. That **reused root SSH private key**, used for machine-to-machine communication, grants unrestricted root access across every NMS virtual machine: configuration (CM), performance (PM) and fault management (FM), databases, application services and backup. The same key is often relied on by disaster-recovery processes so the NMS can be restored.
3. From there, **total compromise of the NMS, and therefore of the 5G core network it controls**.

No zero-day was involved at any step. The compromise was built entirely from insecure passwords, reused keys and weak boundaries.

This is where the accountability problem becomes concrete. If it is not documented and verified who is responsible for changing default credentials and hardening the system after deployment, the vendor and the operator each assume the other did it, and neither does. The practical lesson is simple: post-deployment hardening cannot be assumed; it has to be verified.

What to verify

- Is there a documented owner for changing default and weak passwords, the vendor or the operator, as part of the security policy?
- Is there a complete inventory of all users and credentials for the new system?
- Is post-deployment hardening confirmed in writing by the vendor and independently checked, ideally as a specific scope item in a penetration test?
- Are hardcoded credentials actively hunted for? This is a task penetration testing is well suited to.
- Do disaster-recovery or backup processes depend on a single shared root SSH key across the management estate, and if so, what replaces it?

4.2 Remote code execution

Remote code execution (RCE) lets an attacker run operating-system commands on the server hosting an application, typically leading to full compromise of that application and its data. In a 5G core, a successful RCE against a network function can hand an adversary control of a critical piece of the network.

RCE is still more common in the OAM domain, because of the large number of management and configuration services there. But the adoption of HTTP-based service communication in the 5G SBI raises the relevance of the signalling side too. The 5G core has inherited the application-layer risks already familiar from the web, while still carrying telecom's operational weight.

A real example is **OS command injection inside a custom SSH authentication function**. To extend SSH beyond local users toward centralised authentication (LDAP, RADIUS, Keycloak and similar), the correct approach is to use Linux PAM modules or secure APIs. Sometimes, though, developers take a shortcut and build a shell command directly from user-supplied input, instead of using a vetted authentication library.

Because the user-supplied input is concatenated straight into a shell command, a value crafted to break out of its intended field is executed by the system as a command rather than treated as data. The result is **unauthenticated RCE**, since the attacker needs no valid account, only network reachability. This anti-pattern is not language-specific: it can appear wherever user input is concatenated into a shell command during authentication. The fix is to never pass user input to a shell, and to rely on parameterised APIs or a vetted authentication library instead.

This class of risk is not academic. Unauthenticated RCE in SSH-facing components is exploited in the wild, by a range of mechanisms. **CVE-2025-32433** is one public example: a flaw in how the Erlang/OTP SSH daemon handles pre-authentication protocol messages, which allows code execution without a valid account, rated CVSS 10.0 and listed in CISA's Known Exploited Vulnerabilities catalogue. Erlang/OTP is used in telecom and operational-technology environments for SSH-based host management, which makes this directly relevant to the 5G estate. The functionality that introduces these flaws is usually legitimate. It is the implementation that creates the compromise path, and when the flaw is unauthenticated, the attacker only needs to reach the service.

4.3 Privilege escalation

Once an attacker has an initial foothold, often as a low-privileged user obtained through an RCE, the next objective is privilege escalation: elevating access to administrator or root, and thereby to unauthorised administrative operations and sensitive data. In 5G environments we distinguish two forms.

Local privilege escalation, from a standard user to root on a single host, driven by problems defenders already know well:

- Misconfigured `sudo` rules that grant excessive permissions, often straightforward to exploit.
- Command injection in proprietary, vendor-delivered binaries.
- Known vulnerabilities in outdated kernels or end-of-life operating systems.

Kubernetes privilege escalation, which is where cloud-native 5G becomes dangerous at scale:

- **Pod-to-cluster escalation**, where a pod's misconfigured identity and access management permissions are abused to seize administrative control of the cluster.
- **Pod-to-host escalation**, where a vulnerability is used to escape the container and compromise the underlying host node.

The strategic impact of a cluster compromise is hard to overstate. It matters most in converged core deployments, sometimes called dual-mode or multi-mode cores, where one vendor solution packages both 5G and EPC control-plane functions, the AMF and SMF alongside the 4G MME, SGW-C and PGW-C, inside a single cluster. Taking over the cluster gives an attacker administrative control over all of them at once, across both generations. A single-system breach becomes a catastrophic failure scenario in which an adversary could simultaneously de-register subscribers, tear down data sessions and reroute user traffic. This is exactly where 5G's clustered, cloud-native complexity amplifies a security failure: one escalation path can create far more blast radius than teams expect.

What to verify (RCE and privilege escalation)

- Are critical vulnerabilities patched quickly once discovered? In telecom this is often not the case, because cyber risk is weighed against availability risk, vendors can be slow, supply-chain components are poorly understood, and end-of-life systems linger. The Kubernetes layer adds further complexity.
- Are detection capabilities such as EDR and XDR deployed on the relevant hosts and nodes?
- Do vendors integrate security resources and secure-coding references into their development lifecycle?
- Is regular, expert-led penetration testing scheduled? It remains one of the most reliable ways to find these implementation-specific flaws first.

4.4 Server-side request forgery in the SBI

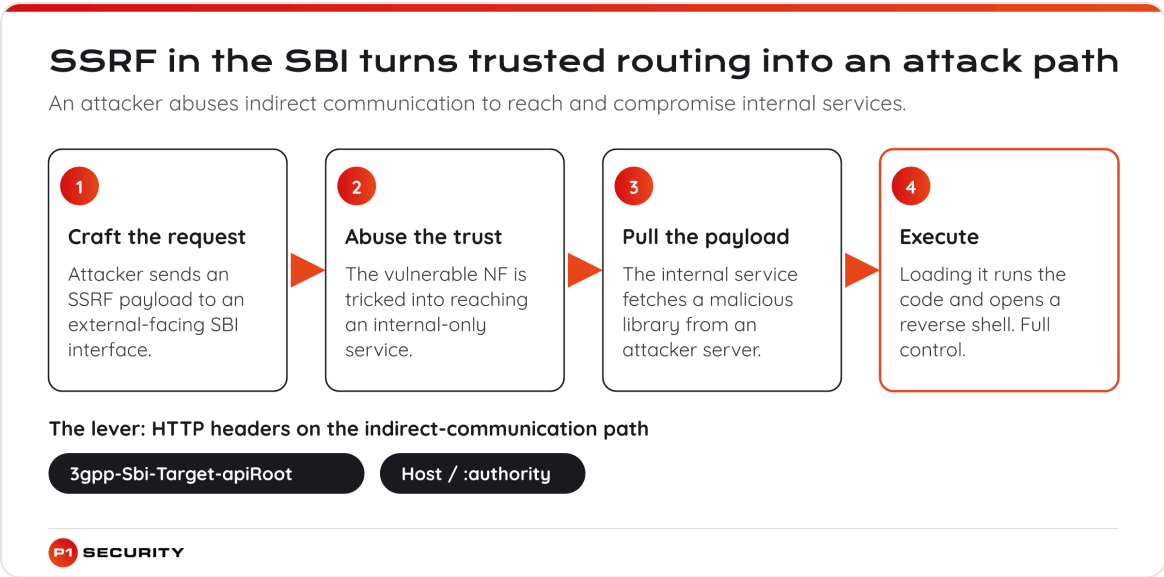

Figure 3. Server-side request forgery in the SBI: the four steps from a crafted request to a reverse shell

Server-side request forgery (SSRF) forces a server-side application to make requests to an unintended location, effectively turning the server into a proxy for the attacker to reach internal systems it could not touch directly. In ordinary infrastructure, that is already serious. In the highly interconnected 5G SBA, where indirect communication between network functions is a design feature, it becomes especially potent. We see two main patterns.

Business-logic SSRF. In one engagement, an external-facing SBI interface accepted a crafted request that let us drive a multi-step chain to full compromise:

1. The attacker sends a crafted request containing an SSRF payload to an external-facing SBI interface.
2. The vulnerable network function is tricked into sending a request to an internal-only service.
3. That internal service is instructed to retrieve and load a malicious library from an attacker-controlled server.
4. On loading the library, the system executes the attacker's code and establishes a reverse shell, granting full control.

Routing-based SSRF and abuse of indirect communication. The 5G architecture allows network functions to talk indirectly through a Service Communication Proxy (SCP). An attacker can manipulate the `3gpp-Sbi-Target-apiRoot` HTTP header (or the `Host` / `:authority` header) in a request to the SCP, supplying the address of an internal system instead of a legitimate network function. If the SCP's routing rules are

not properly validated, it forwards the request, and the attacker has effectively turned the SCP into an internal network scanner, mapping live hosts and open ports inside the cluster that should have been unreachable.

The reason this matters so much in 5G is that it turns trusted telecom control logic into a relay. Rather than attacking an internal service directly, the attacker abuses the network's own service-communication model to reach it, enabling internal reconnaissance, access to hidden resources, and in some cases compromise of a downstream system.

4.5 Denial of service in the SBI

A denial-of-service (DoS) attack makes a system unavailable to its legitimate users, either by overwhelming it or by triggering a crash. The 5G SBI is unusually exposed to the second kind, because its APIs are complex. SBI requests carry many parameters and attributes, structured JSON bodies, and a set of custom 3GPP HTTP headers that functions are mandated to support, such as `3gpp-Sbi-Target-apiRoot` used for indirect communication via the SCP and SEPP. That richness is a large attack surface for robustness failures, and it is often a robustness and input-handling problem rather than a brute-force traffic problem.

Common root causes we find through fuzzing and robustness testing:

- **JSON parser issues**, where a single malformed or unexpected JSON value crashes the backend network function.
- **Missing input validation**, where unexpected data types trigger unhandled exceptions that bring a service down.
- **Regular-expression DoS (ReDoS)**, where an inefficient regular expression used to validate subscriber identifiers or callback URIs is exploited with a crafted string, causing catastrophic backtracking that consumes CPU and crashes the application.
- **Buffer overflows**, where memory-unsafe backend components can be crashed or, worse, escalated toward code execution.

In a telecom context these are not just application crashes. Depending on the function hit, a failed network function can affect subscriber service, session handling and mobility, or broader core behaviour, translating directly into subscriber-facing outages, failed calls and data sessions, and reputational damage. The lesson is that SBI endpoints need continuous robustness and fuzz testing, not only feature testing.

What to verify (SSRF and DoS)

- Is continuous robustness and fuzz testing performed against SBI interfaces?
- Is network segmentation and application-level access control enforced between functions and domains?
- Are authentication and authorisation required on all internal services, for example mutual TLS plus OAuth 2.0 access tokens as part of a zero-trust approach?
- Are the routing configurations of the SCP and SEPP validated and tested against header manipulation?

5. Emerging risks in mature 5G SA and roaming

Beyond today's recurring weak spots, three risks are set to become more prominent as 5G SA and inter-operator roaming mature. They exploit the deeply interconnected and technically complex nature of the full 5G ecosystem, extending the attack surface beyond a single operator.

5.1 The roaming interconnect as an SSRF gateway

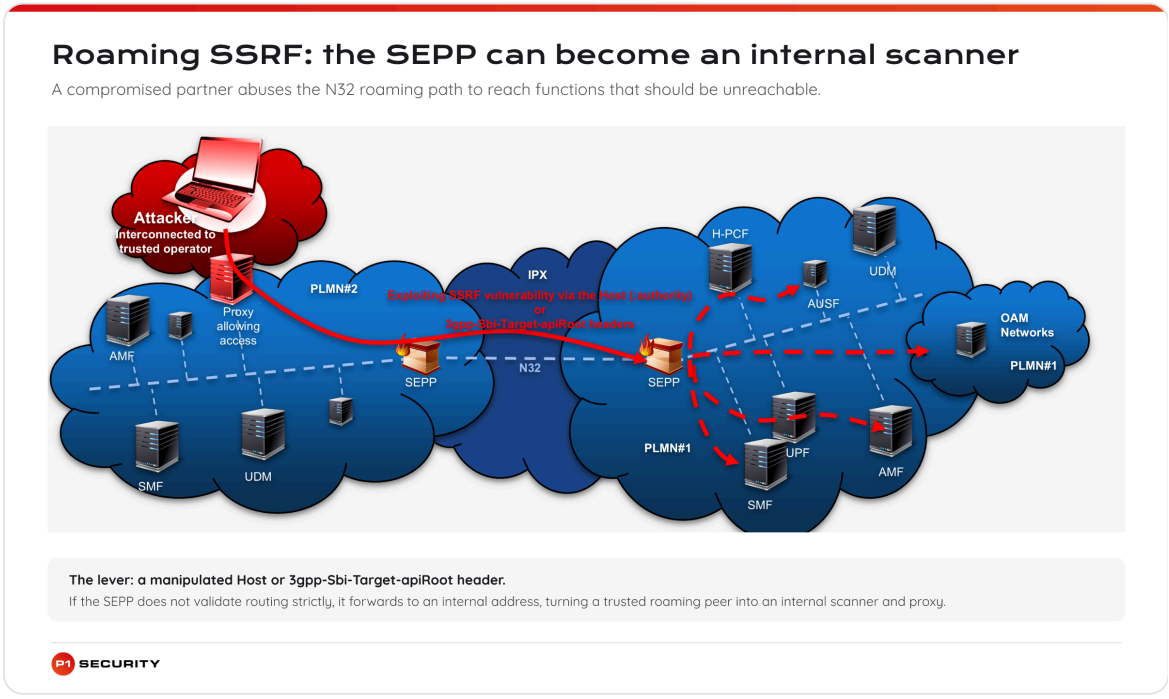

Figure 4. Roaming SSRF: a compromised partner can turn the victim's SEPP into an internal scanner

When operators interconnect, they do so through a Security Edge Protection Proxy (SEPP) over the N32 interface. An attacker at a malicious or compromised partner operator could send a request to a victim's SEPP and manipulate the `Host` header or the `3gpp-Sbi-Target-apiRoot` header. If the victim's SEPP lacks strict routing validation, it could be tricked into forwarding the request to an internal IP address inside the victim's network rather than to a legitimate roaming partner function.

The trusted SEPP then becomes a powerful internal scanner and proxy, reachable across the roaming boundary, capable of attacking sensitive internal segments including the OAM network.

We have not yet observed 5G SA roaming deployed widely in the wild, which is exactly why now is the time to validate it. Waiting until 5G roaming is deployed at scale is the wrong moment to discover a weak routing assumption.

5.2 The latent danger of HTTP/1.1 support

3GPP specifies HTTP/2 for SBI communication. In practice, many network functions still maintain backward compatibility with HTTP/1.1. That reintroduces an unnecessary and significant class of risk, effectively re-opening attack vectors that HTTP/2's binary framing was designed to reduce, such as **HTTP request smuggling (desync)**, which exploits a discrepancy in how a front-end proxy (such as a SEPP or SCP) and a back-end network function interpret an ambiguous request. By smuggling a second request inside a benign one, an attacker can bypass security controls, poison the request pool, exfiltrate sensitive user data and circumvent firewall rules.

Smuggling can also arise on HTTP/2-to-HTTP/1.1 downgrade paths, which is why ensuring that all upstream communication uses HTTP/2 end to end matters. It is a small change from a compatibility point of view and a large one from a security point of view.

5.3 Supply chain and patch management

Two systemic industry problems combine into a compounding risk.

First, **poor supply-chain visibility.** Vendors rely heavily on open-source tools and libraries, which is normal, but many provide inaccurate or incomplete software bills of material (SBOMs). If neither vendor nor operator fully knows what third-party components are inside a deployment, patching becomes slower and less complete. An operator cannot patch a vulnerability it does not know it has.

Second, **dangerously slow patch cycles.** In telecom it is common for critical core-network vulnerabilities to take a year or more to move from CVE assignment to public fix, leaving a long window in which effectively known weaknesses remain exploitable.

The exposure belongs to anyone who discovers the flaw independently in the meantime. Closing that window is one of the hardest problems in telecom security: even when the industry knows what needs fixing, operational pressure, dependency chains and availability concerns often slow real remediation to an unsafe pace. It depends on both accurate SBOMs and a disciplined patch process.

6. How to mitigate 5G core vulnerabilities: defence in depth

Preventing a compromise by a capable APT group does not require exotic defences. It requires discipline and defence in depth. The measures below are organised into three tiers, from foundational hygiene to proactive operations, and they map directly to the vulnerability classes above.

Tier 1: foundational security hygiene

- **System hardening.** Change or remove every default credential and insecure configuration as a mandatory, verified post-deployment step. This is the concrete fix for the accountability gap: it turns "who was supposed to do this" into a documented, checkable action in the deployment lifecycle.
- **Patch management.** Apply security patches rigorously and promptly, informed by a complete and accurate SBOM for every network function so third-party flaws are not missed.
- **Password management and MFA.** Enforce unique, complex credentials everywhere, and deploy multi-factor authentication wherever possible to blunt credential theft and brute forcing.

Tier 2: network and user security

- **Network segmentation.** Enforce strict firewalling and filtering between segments, especially between the OAM, IT and core signalling networks, so a single foothold cannot reach high-value targets.
- **Application-level access control.** Adopt zero trust internally. No traffic between network functions should be trusted automatically; require strong authentication and authorisation on internal services, for example mutual TLS (mTLS) that validates both client and server on every transaction, and OAuth 2.0 token-based authorisation between network functions as specified in 3GPP TS 33.501.
- **User role segmentation.** Separate duties and limit administrative reach so no single account or key unlocks the estate.

Tier 3: proactive security operations

- **Detection and response.** Deploy EDR and XDR on hosts and nodes. The better a network is hardened and segmented, the more noise an attacker must make to move, and the more likely detection becomes.
- **Deceptive technology.** Place honey tokens, such as fake credentials, as tripwires. Because no legitimate process uses them, any access attempt is a high-fidelity intrusion signal and an early warning.
- **Penetration testing.** Schedule regular, expert-led, telecom-specific penetration testing for both vendors and operators. It is one of the most effective ways to find the implementation-specific flaws described in this paper before an adversary does, alongside the hygiene, segmentation and detection measures above.

The logic is cumulative: better hygiene and tighter segmentation force attackers to work harder, more effort creates more noise, more noise improves the chance of detection, and better detection improves the chance of ejecting the attacker before a compromise becomes catastrophic.

7. Operator readiness checklist

A condensed, practical checklist drawn from the sections above.

- [ ] Documented ownership for post-deployment hardening and credential changes (vendor vs operator), and independent verification that the hardening was actually done.
- [ ] Complete inventory of users, credentials and keys for every new system; no default or reused credentials in production.
- [ ] No reused SSH private keys for machine-to-machine access; unrestricted root SSH access removed.
- [ ] Secure-coding practices confirmed with vendors, especially no user input concatenated into OS commands during authentication.
- [ ] Timely patch process, informed by an accurate SBOM per network function; end-of-life operating systems and kernels retired.
- [ ] Kubernetes hardening: least-privilege pod RBAC and service accounts, container-escape mitigations, and awareness of which functions share a cluster.
- [ ] SCP and SEPP routing configuration validated and tested against `3gpp-Sbi-Target-apiRoot` and `Host` header manipulation.

- [ ] Continuous robustness and fuzz testing of SBI interfaces, not only feature testing.
- [ ] HTTP/2 enforced end to end; HTTP/1.1 support removed from production paths.
- [ ] Strict segmentation and filtering between OAM, IT and core signalling networks.
- [ ] Zero trust on all internal service-to-service traffic: authentication and authorisation, ideally mutual TLS plus OAuth 2.0 tokens between network functions.
- [ ] EDR/XDR deployed; honey tokens placed; an incident-response plan able to eject a persistent adversary.
- [ ] Regular, telecom-specific penetration testing across partner, internal and roaming perspectives.

8. Frequently asked questions

What are the most common 5G core security vulnerabilities?
In P1 Security's 5G SA penetration tests, the recurring classes are authentication and access-control failures, remote code execution, privilege escalation (including Kubernetes pod-to-cluster and pod-to-host), server-side request forgery in the SBI, and denial of service in the SBI.

Why does 5G standalone widen the attack surface?
5G SA moves core control-plane signalling from telecom-specific protocols onto a service-based architecture over HTTP/2 and JSON, with containerised network functions usually running on Kubernetes. Telecom protocols such as NAS, NGAP, PFCP and GTP-U remain in place, so the SBA adds web and cloud-native risk on top of existing telecom exposure rather than replacing it.

Is HTTP/1.1 a risk in 5G?
Yes. 3GPP specifies HTTP/2 for the SBI. Network functions that keep HTTP/1.1 backward compatibility re-open request smuggling and desync attacks, so HTTP/2 should be enforced end to end.

How can operators reduce 5G core risk?
Verified removal of default credentials, patching against an accurate SBOM, network segmentation, internal zero trust with mutual TLS and OAuth 2.0 authorisation between network functions, continuous SBI fuzzing, and regular telecom-specific penetration testing.

9. 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. The services and products that map to them:

- **Audit, vulnerability assessment and red team.** Telecom-specific penetration testing and architecture review across the 5G core and SBA, OAM, RAN, roaming and interconnect, NFVI and cloud, tested from the partner, internal and roaming perspectives described in section 3.
- **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 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.
- **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 signalling misuse and lateral movement described here are 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 5G 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 5G core against the weak spots in this paper, contact us at contact@p1sec.com.

10. References and further reading

Standards and industry guidance:

- **3GPP TS 23.501**: system architecture for the 5G system (the SBA and the network functions).
- **3GPP TS 29.500**: technical realization of the service-based architecture (HTTP/2 as the SBI transport, indirect communication via the SCP, and the `3gpp-Sbi-Target-apiRoot` header).
- **3GPP TS 33.501**: security architecture and procedures for the 5G system (SEPP and N32 protection, inter-NF authentication, TLS and mutual TLS on the SBI, OAuth2 authorisation).
- **GSMA FS.36**: 5G interconnect security (SEPP and N32).
- **GSMA NG.113**: 5GS roaming guidelines.
- **GSMA FS.31**: baseline security controls.
- **GSMA FS.21**: interconnect signalling security recommendations, with per-protocol guidance in **FS.11** (SS7), **FS.19** (Diameter) and **FS.20** (GTP).
- Public threat-intelligence reporting on telecom-targeting threat actors (for example CrowdStrike on LightBasin and Liminal Panda, and Unit 42 on Phantom Taurus and on CVE-2025-32433 exploitation; Microsoft and CISA/FBI advisories on Salt Typhoon).

Security frameworks and threat references:

- **MITRE FiGHT (5G Hierarchy of Threats)**: the 5G-specific, ATT&CK-style adversary technique framework that SOC and detection teams can map these vulnerability classes to: https://fight.mitre.org/
- **OWASP API Security Top 10**: the SBI is a set of HTTP/2 REST APIs, so the SSRF, RCE and DoS classes here map directly onto standard API security risks: https://owasp.org/API-Security/
- **ENISA threat landscape for 5G networks**: the standard European reference for the 5G threat picture, aimed at the operator and regulator audience: https://www.enisa.europa.eu/publications/enisa-threat-landscape-for-5g-networks
- **GSMA NESAS and 3GPP SCAS**: the network equipment security assurance scheme and its security assurance specifications, the independent baseline that speaks directly to the vendor-versus-operator accountability question this paper raises.

- **CIS Kubernetes Benchmark**: the authoritative hardening baseline behind the Kubernetes and pod IAM recommendations in this paper: https://www.cisecurity.org/benchmark/kubernetes
- **NVD, CVE-2025-32433**: the Erlang/OTP SSH pre-authentication remote code execution flaw (CVSS 10.0) cited in section 4.2: https://nvd.nist.gov/vuln/detail/CVE-2025-32433
- **CISA Known Exploited Vulnerabilities (KEV) catalogue**: lists CVE-2025-32433 among vulnerabilities observed in active exploitation: https://www.cisa.gov/known-exploited-vulnerabilities-catalog

P1 Security resources:

- SecurityHub TelcoSec glossary entries for SBA, SBI, SCP, SEPP, NRF, AMF, SMF, UPF, N32, HTTP/2 and Kubernetes in mobile networks.
- Related P1 blog posts: pen-testing 5G core networks; SBA API security in the 5G core beyond TLS; HTTP/2 in 5G networks; SCP security in the 5G core; SEPP security in 5G roaming; NRF security in the 5G core; AMF security hardening; secure 5G core design patterns; zero-trust architecture in 5G.

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, operator or vendor is identified as the subject of a P1 Security finding without their consent. Third-party research and advisories are cited as published sources. 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.