Palindrome Technologies

ASSURANCE | TRUST | CONFIDENCE

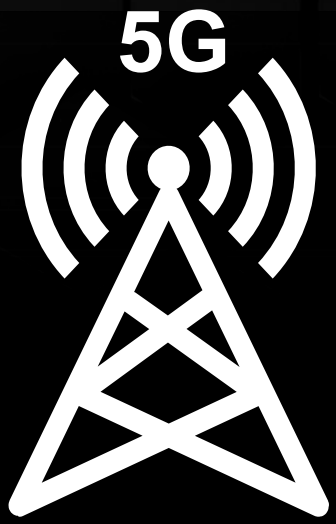

Zero Trust

Implementing Zero Trust in 5G Networks

An Introduction for a Maturity-Model-Based Path Using CISA, NIST, 3GPP, and O-RAN Guidance

Shashank Murali, Peter Thermos
PALINDROME TECHNOLOGIES
www.palindrometech.com

Contents

Tables and Figures

Abstract

Zero trust has become a dominant cybersecurity paradigm for enterprise information systems, yet its practical application to fifth-generation (5G) networks remains underdeveloped. Existing guidance explains zero trust principles and identifies threats affecting 5G infrastructure, but it does not provide a sufficiently concrete implementation path for mobile network operators, integrators, private 5G adopters, or government stakeholders. This paper proposes a 5G-specific zero trust implementation model that adapts the CISA Zero Trust Maturity Model to the realities of a distributed, cloud-native, API-driven, and multi-domain 5G ecosystem. The approach integrates concepts from NIST SP 800-207, NIST SP 1800-35, 3GPP standards, DHS/CISA and NSA/CISA 5G guidance, and O-RAN security guidance. The paper maps zero trust principles to the 5G core, service-based architecture, transport, radio access network, O-RAN interfaces, edge computing environments, network slicing, management and orchestration systems, and software supply chain. It further explains how defense-in-depth and zero trust are complementary rather than competing models: defense-in-depth provides layered preventive, detective, and resilient controls, while zero trust constrains trust decisions and enforces continuous verification at identity, workload, device, API, and data boundaries. A phased maturity model is presented with technical controls, deployment considerations, and examples for each implementation stage. The result is a more operational path for implementing zero trust in 5G systems that is aligned with well-established standards while remaining practical for real deployments.

Keywords: Zero Trust, 5G Security, 5G Core, Service-Based Architecture, CISA Zero Trust Maturity Model, NIST SP 800-207, 3GPP TS 33.501, O-RAN Security, Network Slicing, MEC, Cloud-Native Security, Telecom Cybersecurity.

1 Executive Summary

Securing large-scale communication networks has historically been difficult because these environments evolve faster than the trust models used to protect them. Telecommunications ecosystems have long combined heterogeneous equipment, multiple protocol layers, geographically distributed infrastructure, administrative silos, vendor dependencies, roaming relationships, and strict availability requirements. As these systems became increasingly software-defined and interconnected with enterprise, cloud, and operational environments, the traditional security challenge expanded from perimeter defense to the protection of a deeply interdependent ecosystem. The result is a long-standing tension between openness and interoperability on one hand, and assurance, containment, and resilience on the other. In practical terms, communication networks have often been secured through defense-in-depth strategies that distribute preventive, detective, and resilient controls across infrastructure, transport, applications, and operations, yet still leave room for excessive implicit trust between components, domains, and management planes.

Recent threat activity has made these structural weaknesses more difficult to ignore. The Salt Typhoon[15] campaign demonstrated how state-sponsored actors can maintain persistent access in U.S. critical infrastructure environments, including the communications sector, using stealthy techniques designed to pre-position for disruptive effects rather than merely collect intelligence. CISA and partner agencies assessed that such activity was intended to enable lateral movement toward operationally significant assets in the event of geopolitical crisis. More broadly, official reporting on PRC-linked compromises of telecommunications and backbone infrastructure has underscored that communications networks are not only high-value intelligence targets, but also strategic terrain in which trusted connections, provider edge systems, and management paths can be abused for long-term persistence and cross-domain pivoting. These examples reinforce a central lesson: in complex communications ecosystems, security cannot rely on the assumption that internal location, trusted peering, prior authentication, or vendor integration are sufficient indicators of legitimacy.

This paper argues that modern 5G security must be understood as the combined application of defense-in-depth and zero trust rather than the replacement of one by the other. Defense-in-depth remains essential because a 5G system is a distributed, cloud-native, API-driven, and multi-domain ecosystem that requires layered protections across user equipment, radio access functions, transport, service-based core functions, cloud and edge platforms, orchestration, and inter-operator boundaries. Zero trust is needed because those layers alone do not adequately govern how trust should be evaluated across the interfaces that connect them. By adapting the CISA Zero Trust Maturity Model [2] into a 5G implementation model aligned with NIST SP 800-207[1], NIST SP 1800-35[3], 3GPP

standards, and O-RAN security guidance, this paper offers a more concrete path for applying zero trust principles within a defense-in-depth architecture for modern communications systems.

Zero trust is used in this paper in the architectural sense reflected in NIST SP 800-207 publication, where no subject, device, workload, network function, interface, or partner domain is trusted solely because of its location, prior admission, or membership in an internal network. Each access request is authenticated, authorized, encrypted where appropriate, evaluated against dynamic policy, and reassessed as conditions change. This approach shifts from perimeter reliance to resource-centric control where the network is treated as potentially compromised ("*assumed breach*") and the design objective confines lateral movement, constrains privilege, reduces dwell time, and produces sufficient telemetry for timely detection and response.

2 Defense-in-Depth and Zero Trust in 5G

The core objective of defense-in-depth is to reduce the probability and impact of compromise by distributing security controls across multiple layers of an ecosystem rather than relying on any single safeguard, trust boundary, or enforcement point. In practical terms, defense-in-depth seeks to ensure that failure at one layer does not directly translate into systemic failure, because additional preventive, detective, responsive, and resilient mechanisms remain available elsewhere in the environment. The core objective of zero trust is different but complementary and aims to minimize implicit trust by requiring that access and communications be explicitly verified, narrowly authorized, context-aware, and continuously reassessed throughout the system lifecycle.

Defense-in-depth and zero trust should therefore not be treated as competing security philosophies. Rather, defense-in-depth is the broader layered security strategy for an ecosystem, while zero trust is a more specific architectural discipline that governs how trust decisions are made within and across those layers. Defense-in-depth provides the structural distribution of controls across physical infrastructure, platforms, networks, identities, applications, data, and operations. Zero trust strengthens that layered model by requiring explicit verification, least-privilege access, contextual policy enforcement, and continuous reassessment of trust instead of allowing trust to be inferred from network location, prior authentication, vendor association, or internal connectivity.

For 5G systems, this distinction matters because a defense-in-depth architecture may already implement security controls such as hardened radio nodes, transport protections (e.g., IPsec or TLS), subscriber authentication, orchestration controls, logging, and network segmentation. But these controls alone do not ensure zero-trust outcomes. For example insecure conditions such as a compromised workload inside the 5G core, a vulnerable API, a misconfigured xApp in the Near-RT RIC, or an overprivileged orchestration account, may

still be trusted too broadly. Therefore, zero trust does not eliminate layered security but rather changes the trust assumptions inside those layers.

Figure 1 Defense-in-depth as layered security strategy and zero trust as the cross-layer trust-governance model in 5G

Figure 1 emphasizes that defense-in-depth and zero trust are interdependent rather than merely adjacent concepts. Defense-in-depth supplies the layered control environment within which zero trust decisions can be made and enforced. Those layers provide the mechanisms for isolation, cryptographic protection, logging, resilience, and recovery. Zero trust, in turn, determines how those mechanisms should behave by requiring that access and communications be explicitly verified, narrowly authorized, continuously reassessed, and bounded by context. In this sense, defense-in-depth answers the question of ***where***

controls exist across the ecosystem, while zero trust answers ***how*** trust should be evaluated within and across those controls. Without layered security, zero trust lacks sufficient enforcement and resilience mechanisms; without zero trust, defense-in-depth may still contain excessive implicit trust between layers, workloads, or domains.

In 5G systems, these interdependencies are especially important because the same transaction often depends on multiple layers at once. For example, secure service-based communication between network functions depends on platform integrity, including certificate issuance and secure storage, transport protection, API authorization, and observability. If any of these layers is vulnerable, zero trust enforcement becomes unreliable even if certain controls are implemented (e.g., authentication, TLS using self-signed certs). Another example is where an xApp may be onboarded through a software supply chain process, deployed onto O-Cloud infrastructure, authenticated to the Near-RT RIC, and granted access to selected E2 functions. Defense-in-depth provides separate controls that protect each stage of this chain, while zero trust constrains how much the xApp can do at each step and whether that trust should persist over time. Similarly, in MEC and network slicing environments, layered protections such as segmentation, workload isolation, secrets management, and telemetry are only effective when zero trust policies define which tenants, applications, devices, or management functions may interact and under what conditions. The relationship is therefore not hierarchical but operationally recursive, where each layer supports zero trust decisions, and zero trust improves the security value of each layer.

3 Why 5G Requires a Distinct Zero Trust Model

A mature 5G deployment differs from legacy mobile networks in several important ways. First, the 5G core is service-based and heavily API-oriented. Network functions expose services to one another over service-based interfaces, which means the compromise of one function can propagate through authenticated but over-trusted service relationships if fine-grained authorization and workload identity are not enforced. Second, the platform is increasingly cloud-native where container orchestration, CI/CD pipelines, registries, service meshes, and software supply chain artifacts become part of the telecom trust model. Third, traffic and control planes are distributed across central sites, edge locations, transport infrastructure, radio domains, and management systems. Fourth, O-RAN introduces open interfaces and application-driven control structures that increase interoperability while also broadening trust boundaries but also expand the attack surface.

The result is a system in which trust decisions occur simultaneously across several layers such as subscriber and device identity, network function identity, application authorization, API mediation, slice isolation, edge workload admission, management-plane access, and

cross-domain federation. A perimeter-centric design is inadequate because the most consequential interactions often occur between internal software components or between trusted infrastructure domains. This is precisely the type of environment that zero trust was meant to address, but only if its principles are translated into telecom-specific control patterns.

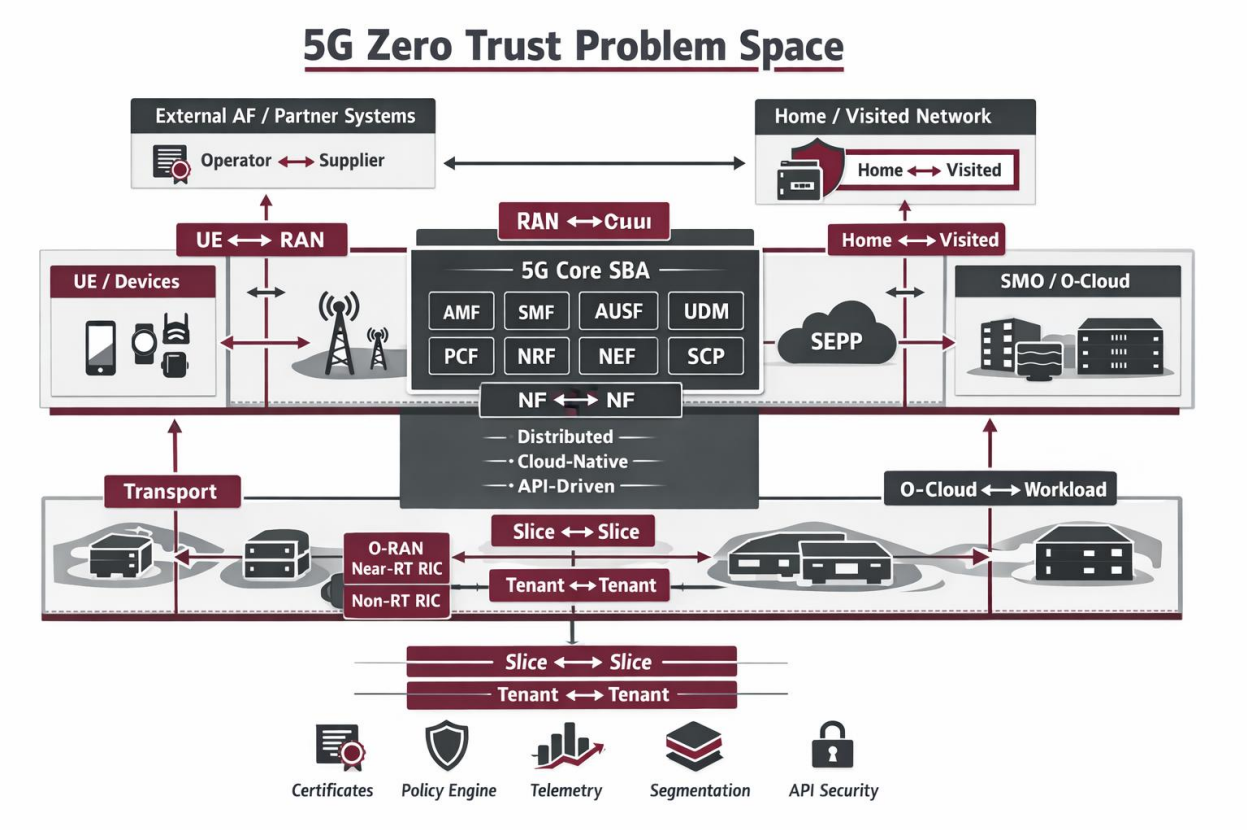

Figure 2 5G zero trust problem space

Figure 2 illustrates that the challenge of implementing zero trust in a 5G ecosystem is not limited to a single access-control plane or a single trust boundary but rather, it encompasses the interaction of multiple concurrent trust relationships across identities, workloads, interfaces, tenants, and operator domains. For example, in the 5G core, network functions communicate through service-based interfaces that may implement authentication but are still over-permissive if discovery, authorization, and API exposure are not tightly constrained. Similarly, in the RAN and O-RAN environment, trust extends across the O-RU, O-DU, O-CU, Near-RT RIC, Non-RT RIC, SMO, and O-Cloud, where software-driven control loops can influence radio behavior and performance. At the edge, we have MEC platforms that introduce third-party or tenant-specific applications with different business owners, different risk tolerances, and different data-handling requirements. For example, a manufacturing analytics application deployed on MEC should not inherit broad

internal reachability merely because it is hosted on operator-controlled infrastructure but rather it should be restricted to using declared APIs and explicitly define approved data stores and tenant-specific policy. In an O-RAN deployment [11,14], a compromised xApp with broad permissions on E2-connected functions, could allow an adversary to manipulate optimization behavior or impact availability, if its privileges are not bounded and its integrity is not continuously validated. These examples underscore that the 5G zero trust problem space is not defined by implementing a single perimeter, but rather by the need to enforce explicit trust decisions at every meaningful interaction point in a highly distributed system.

Although 5G is architecturally capable of supporting zero trust, the guidance in the 3GPP standards remains optional. For example, the 3GPP TS 33.501[5] standard defines security capabilities for the service-based architecture, including mutually authenticated TLS and NRF-mediated authorization patterns, but standards compliance does not by itself guarantee that these controls are enabled, scoped, continuously validated, or applied uniformly across all network functions. A core may support certificate-based authentication and token-based authorization while still permitting broad service discovery, static local authorization, weak token validation, incomplete segmentation, or over-permissive east-west reachability. This optionality underscores the difference between a standards-capable 5G core and a zero-trust-operated 5G core.

In order for a network operator to begin enforcing zero-trust they must identify and inventory all network functions, workloads, certificates, subscribers and administrative identities, exposed APIs, management paths, slice dependencies, interconnect relationships, data repositories, and transaction flows. This discovery phase establishes the factual basis for policy and creates the baseline against which anomalous behavior can later be identified.

4 A 5G-Adapted Zero Trust Maturity Model

The CISA Zero Trust Maturity Model provides five pillars, namely identity, devices, networks, applications and workloads, and data. In addition, it incorporates cross-cutting capabilities in visibility and analytics, automation and orchestration, and governance. For 5G systems, these pillars should be preserved but translated into telecom-specific entities and operational dependencies.

The value of the CISA model is that it makes zero trust measurable without implying that an operator must mature every domain at the same rate. Each pillar can be evaluated independently, which is important in telecom environments where identity, certificate management, or subscriber authentication may be more mature than segmentation, telemetry correlation, or automated response. The model helps produce a maturity profile,

rather than a binary conclusion, where the next practical increment of improvement is identified.

Table 1 DHS-CISA Zero Trust Pillars

CISA Pillar	5G Interpretation
Identity	Subscriber identity, workforce identity, privileged administrator identity, network function identity, workload identity, API client identity, certificate-based machine identity
Devices	User equipment, SIM/eSIM, radio units, distributed units, centralized units, servers, edge nodes, administrator endpoints, IoT devices attached through 5G access
Networks	RAN, transport, service-based interfaces, management networks, slice overlays, interconnect boundaries, east-west cloud communications, O-RAN interfaces
Applications and Workloads	5G core network functions, CNFs/VNFs, xApps, rApps, OSS/BSS components, SMO, CI/CD services, observability and policy engines
Data	Subscriber data, policy data, telemetry, charging data, secrets, software artifacts, model data, operational logs, configuration state

This paper further refines the maturity journey into five implementation levels:

1 **Level 0 – Traditional Telecom Trust.** Trust is inferred from network location, vendor domain, management zone, or prior authentication. Segmentation is coarse. Internal APIs are over-trusted. Platform identity and software provenance controls are limited.
2 **Level 1 – Foundational Zero Trust Controls.** Strong identity, certificate management, authenticated service-to-service communications, baseline segmentation, device inventory, and logging are established for critical domains.
3 **Level 2 – Policy-Driven 5G Segmentation.** Access is based on explicit policy, workload identity, device posture, and interface-specific authorization. Microsegmentation is introduced for core functions and management planes.
4 **Level 3 – Adaptive and Continuously Verified 5G.** Telemetry, analytics, and automated response influence trust decisions in near real time. Workload admission, behavioral analytics, and slice-specific controls become integrated.
5 **Level 4 – Optimized and Autonomous Zero Trust Operations.** Cross-domain policy orchestration, continuous assurance, strong software provenance, closed-loop containment, and highly automated trust reevaluation are operationalized.

Table 2 shows how the general CISA Zero Trust maturity stages translate into a 5G operating environment, where a traditional 5G deployment is characterized by static trust, coarse segmentation, and siloed visibility, while more mature stages progressively introduce authenticated interfaces, workload identity, scoped authorization, per-function or

per-slice segmentation, centralized telemetry, and automated response. The table also reinforces that zero trust in 5G is not a single control or vendor feature, but rather a staged progression from inherited telecom trust assumptions toward dynamic, policy-driven, and continuously validated trust decisions across the 5G ecosystem.

Table 2 Mapping of 5G Implementation Levels to CISA Zero Trust Maturity Stages

5G Implementation Level	CISA Stage Alignment	Alignment Rationale
Level 0 – Traditional Telecom Trust	Traditional	Trust is inferred from network location, vendor domain, management zone, or prior authentication; segmentation, telemetry, and policy enforcement remain coarse and siloed.
Level 1 – Foundational Zero Trust Controls	Initial	Inventories, strong identity, certificate management, authenticated service communications, baseline segmentation, and logging establish the first movement away from static telecom trust.
Level 2 – Policy-Driven 5G Segmentation	Advanced, early	Access is governed by explicit policy, workload identity, device or platform posture, scoped authorization, and micro-segmentation across core, cloud, RAN, management, and interconnect domains.
Level 3 – Adaptive and Continuously Verified 5G	Advanced, mature	Telemetry, analytics, posture signals, behavioral baselines, and automated response influence trust decisions in near real time across selected high-value domains.
Level 4 – Optimized and Autonomous Zero Trust Operations	Optimal	Cross-domain policy orchestration, continuous assurance, automated containment, software provenance enforcement, and dynamic trust re-evaluation are operationalized across the 5G ecosystem.

The five 5G implementation levels are a telecom-specific refinement of the four CISA maturity stages rather than a separate maturity model. The additional granularity is useful because the transition from explicit policy enforcement to adaptive and continuously verified operations is operationally significant in mobile networks. In particular, the CISA Advanced stage is divided into two 5G levels, one focused on policy-driven segmentation and one focused on telemetry-driven adaptation and continuous verification.

5 Implementation of the Maturity Model Across 5G Domains

Across the 5G core, RAN, O-RAN, MEC, slicing, cloud platform, and interconnect, zero trust is implemented through a recurring set of primitives including mutual TLS to establish cryptographic identity between communicating functions and prevent unauthenticated impersonation, OAuth2 or equivalent authorization mechanisms to constrain what an authenticated function may do, micro-segmentation to enforce approved communication paths and limit the blast radius of a compromised workload or slice, workload identity to separate machine identity from unstable network attributes such as IP address and short-lived and posture-bound administrative access to remove the standing-privilege model that can enable credential reuse and long dwell time.

5.1 5G Core and Service-Based Architecture

The 5G core is built around a service-based architecture defined by 3GPP. In this model, network functions such as the Access and Mobility Management Function, Session Management Function, Unified Data Management, Authentication Server Function, and Policy Control Function expose services over service-based interfaces. This architecture improves flexibility and programmability, but it also means that a security failure in API authentication, authorization, certificate management, or service registration can have systemic impact. A zero trust implementation must therefore treat each network function as a workload with its own identity, policy context, and trust boundary.

At foundational maturity, all SBI traffic should be strongly authenticated and encrypted, with certificate-based network function identity and strict lifecycle management. At policy-driven maturity, authorization should become service- and operation-specific rather than binary reachability-based. For example, an NF that is permitted to discover a service should not automatically be trusted to invoke all operations exposed by that service. A practical implementation can combine mTLS, service registration controls, API mediation, and policy enforcement tied to workload identity claims. This is especially important in roaming and interconnection scenarios, where Security Edge Protection Proxy functions and home/visited network trust boundaries already reflect the need for constrained trust under 3GPP security architecture.

Example. Consider a compromised analytics microservice co-hosted in the 5G core platform. In a traditional design, east-west trust and broad service discovery privileges may allow that workload to query policy or subscriber-related services beyond its business purpose. In a zero trust design, the workload is bound to a narrowly scoped machine identity, receives only the minimum service discovery privileges required, is permitted to call only approved APIs, and is continuously evaluated against runtime posture signals. If

its process lineage, image provenance, or network behavior changes, policy can reduce or terminate access in-session rather than waiting for manual intervention.

5.2 3GPP Security Concepts as Zero Trust Building Blocks

Although 3GPP does not use zero trust terminology as its organizing concept, there are several 5G security mechanisms that provide zero trust building blocks when interpreted through a policy-driven trust model. The TS 23.501[5] standard defines the service-based architecture and the roles of core functions such as the NRF, NEF, SCP, PCF, UDM, AUSF, and AMF along with the corresponding security architecture and procedures, including trust boundaries, service registration and discovery protections, network function authorization, subscriber privacy protections, home and visited network security, and requirements for monitoring and configurability. These specifications do not by themselves produce a complete zero trust architecture, but they provide normative control points that can be made more explicit, granular, and continuously verified.

A useful interpretation is that the SBA turns network functions into API-addressable resources and therefore collapses the distinction between traditional telecom signaling trust and application-layer trust. Under a zero trust model, the NRF should not be treated merely as a service directory but as a sensitive control point whose registration and discovery decisions influence the attack surface of the entire core. Similarly, the NEF should be viewed as a high-assurance policy boundary for exposing network capabilities to external application functions. If the NEF is over-permissive, poorly segmented, or weakly authenticated, it can become the mechanism by which external influence is projected into sensitive core behavior. Likewise, the SCP should be treated as a policy enforcement and observability point for service-to-service traffic rather than just a routing convenience.

From a zero trust perspective, 3GPP mechanisms suggest several concrete implementation patterns. First, network function identity should be cryptographically strong, lifecycle-managed, and bound to narrowly scoped authorization policy. Second, service registration and discovery should be policy-constrained so that only approved functions can advertise, discover, or invoke specific services. Third, trust decisions should consider more than successful authentication. Runtime context such as workload provenance, configuration state, platform posture, and observed behavior should influence whether access is maintained. Fourth, external and inter-PLMN boundaries should be treated as explicit trust transitions even when standards-based interconnection protections such as SEPP are present.

The 3GPP security model also reinforces an important distinction between subscriber- or device-facing trust and internal network-function trust. At the access edge, 5G-AKA and EAP-AKA' establish mutual authentication between the UE and the network with protections for subscriber identity and key derivation. However, successful UE

authentication does not justify broad trust elsewhere in the system. A zero trust interpretation of 3GPP therefore separates subscriber legitimacy from downstream authorization decisions involving policy control, service exposure, slice assignment, application access, and management-plane operations. Similarly, signaling plane integrity between network functions should not be inferred solely from membership in the 5G core. Network function registration, discovery, and service invocation should be treated as distinct decisions subject to different policy constraints, because an authorized NF is not necessarily authorized for every service, every operation, or every data domain it can technically reach.

This distinction becomes especially important in practical 3GPP deployment scenarios. For example, the NRF should enforce not only the presence of valid credentials but also the integrity of registration metadata, service claims, and discovery scope. The NEF should be segmented from internal management and core control paths so that exposure of network capabilities to application functions does not create an implicit bridge into broader control authority. The SCP, where used, can support zero trust objectives by centralizing policy enforcement, routing control, and observability for SBI traffic, but it also becomes a high-value target whose compromise could distort policy decisions at scale. At inter-PLMN boundaries, SEPP protects signaling exchanges between networks, yet a zero trust design still requires explicit service allow-lists, tight certificate governance, protected attribute handling, and detection of abnormal transaction patterns that may indicate signaling abuse, misbinding of partner trust, or policy drift. Thus, the 3GPP standards provide the necessary primitives, while zero trust determines how strictly and continuously those primitives are enforced.

For example, in roaming scenarios, 3GPP already recognizes the need to constrain signaling and protect home-visited interconnection through functions such as the Security Edge Protection Proxy. A mature zero trust implementation should build on this by applying explicit allow-lists for exposed services, strong certificate validation, protected attributes, telemetry-driven anomaly detection, and policy separation between roaming-related control flows and unrelated internal services. The design objective is not merely standards compliance, but minimization of the blast radius if a partner relationship, interconnect component, or exposed API path is abused.

5.3 Cloud-Native 5G Platform

NSA/CISA guidance on 5G cloud infrastructures correctly emphasizes lateral movement prevention, secure isolation of network resources, data protection, and platform integrity. These themes align directly with zero trust. In a cloud-native 5G environment, the policy decision point is not limited to traditional access management; it extends to container admission, registry trust, service mesh identity, runtime protection, secrets handling, and orchestration permissions. A concrete path begins with signed images, software bill of

materials, registry controls, hardened cluster baselines, namespace isolation, and least-privilege service accounts. It then progresses toward workload attestation, service-mesh-enforced identity, policy-as-code, and automated quarantine actions. Secure CI/CD is not optional. If a pipeline can publish unauthorized images, modify Helm charts, or change secrets and policies without strong verification, then the entire telecom trust model can be subverted upstream. Zero trust in 5G therefore requires software supply chain controls as a first-class implementation domain rather than an auxiliary DevSecOps concern.

5.4 Network Slicing

5G network slicing is often described as a mechanism for logical separation and differentiated services, but slice isolation does not automatically create zero trust. NSA/CISA guidance on slicing correctly highlights threats involving denial of service, misconfiguration, life-cycle management weaknesses, and cross-domain dependencies. A zero trust maturity model should therefore treat slices as policy domains requiring explicit trust constraints rather than as self-securing constructs. At foundational maturity level, operators should establish clear slice inventories, dependency maps, tenant ownership, and control-plane boundaries. At more advanced maturity levels, slice-specific policy enforcement should constrain which identities, workloads, and management functions may interact with slice resources. Telemetry should also be slice-aware so that anomalous behavior can be evaluated in the context of that slice's threat model and service objectives. In practical terms, a public safety slice, an enterprise private network slice, and a consumer broadband slice should not share identical trust assumptions, even if some physical infrastructure is shared.

5.5 MEC and Edge Environments

Multi-access edge computing complicates zero trust because workloads, policy enforcement, and data processing are pushed closer to users and devices, often at sites with different operational controls and physical security assumptions than central core environments. Low-latency requirements may also encourage local breakouts or exception pathways that bypass centralized security patterns.

A 5G zero trust model should require edge workload identity, policy synchronization with central orchestration, strong admission controls, secure attestation where feasible, and explicit governance over which data must remain local, be exported, or maybe shared with third-party applications. An example is an industrial MEC deployment hosting quality-control analytics for a manufacturing site. In a non-zero-trust design, the analytics application may inherit broad platform reachability because it is "internal" to the edge site. In a zero trust design, that application is restricted to its declared APIs, its service account, its tenant context, and its approved data stores, with policy aligned to both operational technology sensitivity and telecom infrastructure risk.

5.6 RAN and O-RAN

O-RAN creates one of the strongest cases for a 5G-specific zero trust architecture. Disaggregation of the RAN into the O-RU, O-DU, O-CU, Near-RT RIC, Non-RT RIC, SMO, O-Cloud, and application ecosystems increases the number of interfaces, software components, and trust transitions. O-RAN security work has increasingly moved toward zero trust concepts by defining security requirements, protocol protections, certificate management approaches, threat studies, and explicit zero trust analysis for key interfaces and components. The June 2026 O-RAN study on Zero Trust Architecture for O-RAN reinforces this direction by treating zero trust as an architectural posture for the disaggregated RAN rather than as an isolated access-control feature.

A practical O-RAN zero trust approach should treat the A1, E2, O1, O2, and open fronthaul interfaces as explicit policy boundaries. xApps and rApps must not be trusted merely because they are onboarded into the RIC ecosystem. They require signed provenance, lifecycle controls, bounded API permissions, behavioral monitoring, and revocation paths. Likewise, the SMO should be treated as a high-impact control plane subject to strict privilege separation, strong authentication, and extensive logging because compromise there can cascade across multiple domains.

A technically meaningful zero trust treatment of O-RAN requires interface-specific enforcement rather than generic segmentation alone. The open fronthaul between the O-RU and O-DU carries latency-sensitive control and user traffic and therefore benefits from strong mutual authentication, certificate lifecycle management, equipment identity validation, and strict configuration control, even when inline security functions must be minimally intrusive. The E2 interface between the Near-RT RIC and E2 nodes is especially sensitive because it exposes control influence over optimization and near-real-time behavior; policy should therefore constrain which xApps can subscribe to, observe, or modify which E2 service models and under what operating conditions. The A1 interface between the Non-RT RIC and Near-RT RIC should be treated as a high-consequence policy boundary because model outputs, enrichment information, and optimization intent entering through A1 can indirectly affect large portions of the radio environment. Likewise, O1 and O2 interfaces should be separated according to management function, lifecycle authority, and infrastructure domain so that observability, provisioning, and cloud resource control are not collapsed into a single overprivileged administrative path.

The software and infrastructure lifecycle of O-RAN components introduces an additional zero trust requirement that is often understated in radio security discussions. xApps and rApps should be onboarded through signed and verifiable supply-chain processes, with bounded permissions, runtime isolation, and rapid revocation paths if behavior deviates from expected policy. The O-Cloud should support workload isolation, attestation where practical, hardened orchestration, and measured trust in the hosting substrate so that

compromise of the platform does not silently undermine the integrity of the RIC or SMO control plane. A concrete example is an energy-optimization rApp that receives wide telemetry visibility and the ability to influence resource allocation policies across multiple sites. In a weakly governed design, the rApp may become an indirect pivot into broader operational control or a source of service degradation through flawed recommendations. In a zero trust design, its data access is scoped, its actuation authority is bounded, its software provenance is verified, and its interactions are continuously evaluated against performance and policy baselines. This is essential in multi-vendor O-RAN environments, where interoperability increases flexibility but also multiplies trust transitions, software dependencies, and opportunities for privilege propagation.

The June 2026 study[14] is particularly useful because it aligns O-RAN zero trust implementation with the actual O-RAN architectural decomposition. In that framing, the RIC, SMO, O-Cloud, open fronthaul, and RAN application ecosystem should be treated as separate trust zones connected through policy-governed interfaces. This reinforces three implementation priorities in a maturity model. First, O-RAN identity must extend beyond human users to include RAN applications, workloads, infrastructure components, and interface endpoints. Second, authorization must be bound to interface function and operational role, so that an xApp, rApp, or management component is not granted broad authority merely because it has been onboarded. Third, visibility and analytics must be able to correlate behavior across O-RAN interfaces and lifecycle events, because misbehavior may appear as legitimate optimization, telemetry access, or management activity unless evaluated against expected policy and baseline behavior.

For example, a malicious or defective xApp that is granted broad E2-related influence could degrade performance, manipulate optimization decisions, or leak telemetry. A zero trust architecture constrains the xApp to a minimal operational scope, verifies its identity and integrity before deployment, monitors its interactions with the Near-RT RIC, and separates its privileges from unrelated RAN or management functions. This is not conceptually different from controlling a cloud-native microservice, but the telecom impact is higher because control decisions may affect radio performance, availability, or regulated service delivery.

5.7 Recommended Phased Roadmap

A practical roadmap for 5G zero trust should begin with high-value trust boundaries rather than attempt to transform the entire environment at once. The first phase should focus on identity, certificates, asset inventories, privileged access, and authenticated communications for the 5G core, management planes, and cloud platform. The second phase should introduce workload-aware segmentation, API authorization, slice dependency mapping, and software supply chain controls. The third phase should integrate telemetry-driven adaptation, continuous posture evaluation, and automated response for selected

domains such as MEC, RIC platforms, and interconnect boundaries. The final phase should emphasize closed-loop assurance, cross-domain policy orchestration, and measurable security objectives tied to service resilience and operational risk.

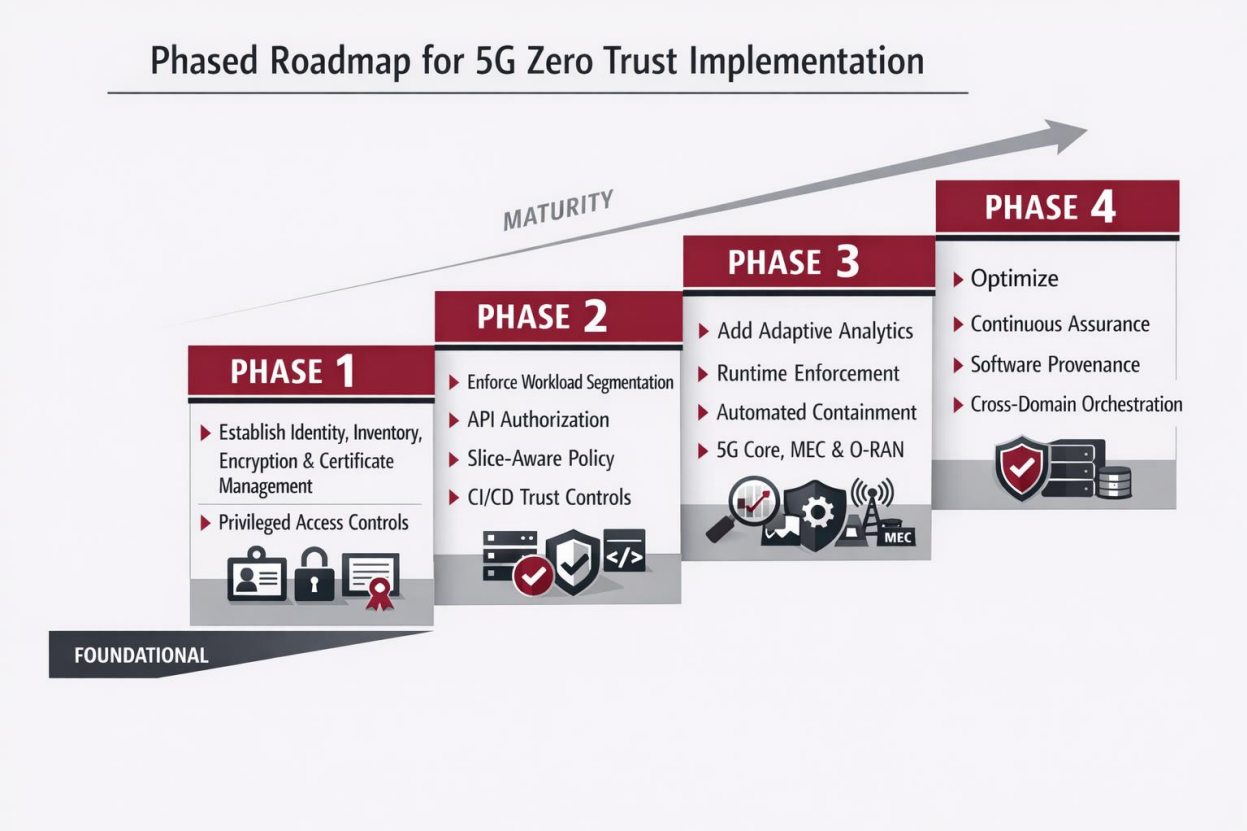

Figure 3 Phased roadmap for 5G zero trust implementation

The relevant standards and guidance should be treated as complementary rather than interchangeable. For example, the NIST SP 800-207 provides the reference architecture and the PDP/PEP enforcement logic where the CISA ZTMM v2.0 provides the maturity rubric and the 3GPP TS 33.501 standard provides 5G security primitives for SBA protection, network function authentication, authorization, and inter-PLMN security. The O-RAN security specifications extend zero trust considerations into disaggregated RAN interfaces, RIC applications, SMO, and O-Cloud environments, and provide a more direct O-RAN-specific approach of how zero trust applies to those components and interfaces. The GSMA interconnect guidance supports deny-by-default signaling protection, monitoring, and filtering for SS7, Diameter, and 5G interconnect scenarios. The implementation task is therefore not to invent a separate telecom zero trust framework, but to integrate these sources into a coherent, measurable, and enforceable operating model.

6 Maturity Levels for APT-Resilient 5G Networks

A practical zero trust target for mobile network operators must be stated in terms of adversarial resistance rather than architectural aspiration. Traditional and initial maturity conditions are not sufficient for high-value 5G domains against state-sponsored telecommunications actors, including campaigns that use credential theft, trusted administrative paths, interconnect abuse, living-off-the-land techniques, supply-chain weaknesses, and exploitation of newly disclosed or unknown vulnerabilities. These initial maturity levels may establish useful hygiene, but they do not prevent lateral movement, long-lived credential reuse, over-permissive east-west communications, or delayed detection.

For the 5G core, the minimum defensible target should be an Advanced-equivalent maturity level with selected Optimal characteristics applied first to the most consequential trust boundaries including the management plane, edge interconnect, O-RAN control environment, cloud-native hosting platform, privileged administration paths, and software supply chain. This target should not be interpreted as a guarantee that advanced persistent threats can be kept out of mobile infrastructure because no maturity model can eliminate zero-day exploitation, vendor defects, insider misuse, or sophisticated supply-chain compromise. The more realistic objective is to deny adversaries the conditions on which long-term telecom intrusions often depend such as static trust relationships, reusable administrative credentials, weak segmentation, incomplete telemetry, loosely governed partner access, and implicit authorization after initial authentication. The purpose of Advanced maturity is therefore to reduce the adversary's freedom of movement and dwell time, force repeated re-authentication and re-authorization at meaningful control points, and generate sufficient telemetry to detect abnormal behavior before a localized compromise becomes systemic.

Table 4 Recommended Minimum Zero Trust Maturity Targets for High-Risk 5G Domains

5G Domain	Minimum Target	Required Characteristics
5G core and SBA	Advanced	mTLS enabled across service-based interfaces, scoped NRF-mediated authorization, service-specific discovery, validated token audience and expiry, per-function policy, and centralized SBI telemetry.
Management and orchestration plane	Advanced with selected Optimal controls	Just-in-time privileged access, short-lived credentials, strong device posture, session recording, break-glass logging, separation of duties, and automated privilege revocation.

5G Domain	Minimum Target	Required Characteristics
Cloud-native hosting platform	Advanced	Signed images, software bills of materials, admission control, workload identity, namespace isolation, secrets governance, runtime detection, policy-as-code, and controlled registry access.
Interconnect and roaming edge	Advanced	Deny-by-default filtering, strict partner certificate governance, SEPP policy enforcement, protected attribute handling, anomaly detection, and separation from unrelated internal services.
O-RAN, RIC, SMO, and O-Cloud	Advanced for control functions; Initial to Advanced for staged RAN rollout	Interface-specific policy for A1, E2, O1, O2, and open fronthaul; signed xApps and rApps; bounded actuation authority; workload isolation; RIC and SMO telemetry; and rapid revocation paths.
Network slicing and MEC	Advanced for high-risk slices and shared edge platforms	Slice-aware policy, tenant and workload identity, explicit dependency maps, edge workload admission, data access scoping, and telemetry tied to slice or tenant context.
Visibility, analytics, and response	Advanced with selected Optimal controls	Centralized correlation across SBA, signaling, cloud, management, O-RAN, interconnect, and endpoint telemetry; behavioral baselining; automated containment; and cross-domain investigation workflows.

The distinction between Advanced and Optimal is operationally important because Advanced maturity should be treated as the minimum target for domains that can materially affect subscriber data, lawful intercept exposure, core control-plane behavior, roaming relationships, slice isolation, or large-scale service availability. Whereas, Optimal maturity should be pursued selectively where automation can be implemented safely and validated before it is allowed to change live service behavior. For example, automated revocation of an anomalous administrative session may be appropriate earlier than automated modification of radio control policy, because the latter can have direct service performance implications. The maturity target should therefore be risk-based and domain-specific rather than uniform across the entire operator environment.

The realism of this target depends on how an operator can leverage vendor product offerings and their functionality. Current telecom vendors generally provide several of the necessary building blocks, including certificate support, authenticated interfaces, role-based access controls, logging, cloud-native deployment options, orchestration controls,

and standards-based security functions. These capabilities, however, are not equivalent to a complete zero trust architecture when delivered in a default or minimally configured state. Thus, operators must require vendors to provide evidence of enforceable controls that can support zero-trust such as configuration, token validation behavior, certificate lifecycle controls, interface-level authorization, logging and event telemetry, segmentation support, software provenance, and management-plane privilege separation. Procurement language should distinguish between support for a control versus mandatory operational use of that control in the deployed configuration.

Public evidence of any mobile network operator having implemented a fully mature, end-to-end, Optimal-equivalent zero trust architecture across the complete 5G core, RAN, O-RAN, cloud platform, management plane, slicing environment, edge infrastructure, and interconnect is limited [17]. Therefore, operator and vendor announcements should be treated as evidence of direction, not proof of complete architectural maturity. A realistic approach is to establish measurable progress metrics within high-risk domains such as validated inventories, enforced mTLS and authorization, reduced standing privileges, segmented management paths, centralized telemetry, tested containment procedures, and independent evidence that the deployed configuration has closed specific trust gaps and products are capable of demonstrating effectively protection mechanisms to prevent adversarial attacks.

7 Summary

Implementing zero trust in 5G networks requires a more explicit and technically grounded execution path because the combination of 3GPP service-based architecture, cloud-native core platforms, edge computing, network slicing, and O-RAN disaggregation produces a trust problem that cannot be adequately addressed by perimeter-centric methods or by generic enterprise zero trust patterns alone. The most useful way forward is to preserve the rigor of established zero trust frameworks while translating them into telecom-specific domains, interfaces, and operational controls.

By adapting the CISA Zero Trust Maturity Model to 5G environments and grounding the roadmap in NIST, 3GPP, DHS/CISA, NSA/CISA, O-RAN, and GSMA interconnect concepts and security guide[16], network operators can develop a clear implementation path with concrete milestones and maturity metrics. Defense-in-depth remains essential, but it becomes more effective when we stop assuming that every adjacent layer maintains trustworthy conditions. In this sense, zero trust is not a replacement for telecom security engineering but rather, a blueprint to evolve the modern communication infrastructures into more coherent, measurable, secure and resilient ecosystems.

8 References

[1] National Institute of Standards and Technology, *Zero Trust Architecture*, NIST SP 800-207, Aug. 2020. URL: https://csrc.nist.gov/pubs/sp/800/207/final

[2] Cybersecurity and Infrastructure Security Agency, *Zero Trust Maturity Model Version 2.0*, Apr. 2023. URL: https://www.cisa.gov/zero-trust-maturity-model

[3] National Institute of Standards and Technology, *Implementing a Zero Trust Architecture*, NIST SP 1800-35, Jun. 2025. URL: https://csrc.nist.gov/pubs/sp/1800/35/final

[4] 3rd Generation Partnership Project, *System Architecture for the 5G System (5GS)*, 3GPP TS 23.501, current release. URL: https://portal.3gpp.org/desktopmodules/Specifications/SpecificationDetails.aspx?specificationId=3144

[5] 3rd Generation Partnership Project, *Security Architecture and Procedures for 5G System*, 3GPP TS 33.501, current release. URL: https://portal.3gpp.org/desktopmodules/Specifications/SpecificationDetails.aspx?specificationId=3169

[6] 3rd Generation Partnership Project, *Network Domain Security (NDS); IP Network Layer Security*, 3GPP TS 33.210, current release. URL: https://portal.3gpp.org/desktopmodules/Specifications/SpecificationDetails.aspx?specificationId=2279

[7] National Security Agency and Cybersecurity and Infrastructure Security Agency, *Security Guidance for 5G Cloud Infrastructures*, Parts I-IV, 2021-2022. URL: https://www.cisa.gov/resources-tools/resources/security-guidance-5g-cloud-infrastructures

[8] Enduring Security Framework, National Security Agency, Office of the Director of National Intelligence, and Cybersecurity and Infrastructure Security Agency, *Potential Threat Vectors to 5G Infrastructure*, May 2021. URL: https://www.cisa.gov/sites/default/files/publications/Potential_Threat_Vectors_to_5G_Infrastructure.pdf

[9] Enduring Security Framework, National Security Agency, and Cybersecurity and Infrastructure Security Agency, *Potential Threats to 5G Network Slicing*, Dec. 2022. URL: https://media.defense.gov/2022/Dec/13/2003132073/-1/-1/0/POTENTIAL%20THREATS%20TO%205G%20NETWORK%20SLICING_508C_FINAL.PDF

[10] National Security Agency and Cybersecurity and Infrastructure Security Agency, *5G Network Slicing: Security Considerations for Design, Deployment, and Maintenance*, Jul. 2023. URL: https://media.defense.gov/2023/Jul/17/2003260829/-1/-1/0/ESF%205G%20NETWORK%20SLICING-

SECURITY%20CONSIDERATIONS%20FOR%20DESIGN,%20DEPLOYMENT,%20AND%20MAINTENANCE_FINAL.PDF

[11] O-RAN ALLIANCE, *O-RAN Security Requirements and Controls Specifications*, current release. URL: https://www.o-ran.org/blog/o-ran-alliance-security-update-2025

[12] O-RAN ALLIANCE, *O-RAN Security Protocols Specifications*, current release. URL: https://www.o-ran.org/blog/o-ran-alliance-security-update-2025

[13] O-RAN ALLIANCE, *Zero Trust Architecture for Secure O-RAN*, O-RAN WG11 White Paper, May 2024. URL: https://mediastorage.o-ran.org/white-papers/O-RAN.WG11.ZTA%20for%20Secure%20O-RAN%20White%20Paper-2024-05.pdf

[14] O-RAN ALLIANCE, *O-RAN Study on Zero Trust Architecture for O-RAN*, O-RAN.WG11.TR.ZTA-R004-v04.00, Jun. 2026. URL: https://specifications.o-ran.org/download?id=1141

[15] FCC, Implications of Salt Typhoon Attack and FCC Response | Federal Communications Commission, Dec. 5, 2024. URL: https://www.fcc.gov/document/implications-salt-typhoon-attack-and-fcc-response

[16] GSMA, *FS.40 5G Security Guide*, Version 3.0, 16 Jul. 2024. URL: https://www.gsma.com/solutions-and-impact/technologies/security/wp-content/uploads/2024/07/FS.40-v3.0-002-19-July.pdf

[17] National Telecommunications and Information Administration, 5G Challenge Notice of Inquiry Docket No. 210105-0001 COMMENTS OF AT&T SERVICES, INC. URL: https://www.ntia.gov/files/ntia/publications/att_services_02102021.pdf

Contact Information

info@palindrometech.com

www.palindrometech.com

https://palindrometech.com/ai-cybersecurity

About Palindrome Technologies

Palindrome's organizational philosophy is built upon three fundamental principles

Assurance
Trust
Confidence

Founded in 2005, Palindrome Technologies Inc. is a leading applied information security research firm and analysis laboratory having expertise in emerging technologies, embedded systems, communication networks, software, and cloud platforms.

Prior forming Palindrome, the principals of the company worked for Bellcore (Bell Communications Research) in the Security & Fraud group where they supported security assurance efforts for telecommunication providers, product vendors and the US government.

Since its inception Palindrome has been providing a range of high-tech services related to securing emerging technologies, global enterprise organizations (I.e., healthcare, financial, energy, government) and carrier-grade networks.

Palindrome subject matter experts maintain internationally recognized ISO credentials and extensive working experience in securing AI/ML implementations to provide consultation and auditing capabilities to organizations seeking ISO/IEC 42001 Certification.

Palindrome is an accredited ISO/IEC 17025 testing laboratory as well as a FCC, GSMA, CTIA and IEEE designated Cybersecurity Testing Lab. Palindrome has been helping global enterprise organizations, service providers and product vendors with deploying and maintaining secure networks, services, and products. The Palindrome team is also known for its contributions to industry standards bodies (e.g., IEEE, GSMA, CTIA and ATIS), and branches of the US government such as FCC CSRIC VII, CSRIC IX and NIST.