

5G Network Security Design Principles:

Applying 5G Cybersecurity and Privacy Capabilities



Michael Bartock
Jeffrey Cichonski
Murugiah Souppaya
Information Technology Laboratory

Karen Scarfone
Scarfone Cybersecurity

Parisa Grayeli
Sanjeev Sharma
The MITRE Corporation

June 2025

Initial Public Draft

This publication is available free of charge from:
<https://doi.org/10.6028/NIST.CSWP.36E.ipd>

Abstract

This white paper describes the network infrastructure design principles that commercial and private 5G network operators are encouraged to use to improve cybersecurity and privacy. Such a network infrastructure isolates types of 5G network traffic from each other: data plane, signaling, and operation and maintenance (O&M) traffic. This white paper is part of a series called Applying 5G Cybersecurity and Privacy Capabilities, which covers 5G cybersecurity and privacy-supporting capabilities that were demonstrated on the NIST National Cybersecurity Center of Excellence (NCCoE) 5G security testbed as part of the [5G Cybersecurity project](#) at the NCCoE.

Audience

Technology, cybersecurity, and privacy professionals who are involved in using, managing, or providing 5G-enabled services and products. This includes commercial mobile network operators, potential private 5G network operators, and end-user organizations. Readers should already be familiar with the basics of mobile network architectures and components.

Keywords

3GPP, 5G, cybersecurity, privacy, virtual routing and forwarding (VRF)

Note to Reviewers

NIST is particularly interested in your feedback on the following questions:

1. How do you envision using this paper? What changes would you like to see to improve that use?
2. What additional information would you like this paper to provide?
3. What other 5G infrastructure cybersecurity and privacy capabilities are you most interested in learning more about?

Acknowledgments

We are grateful to the following individuals for their generous contributions.

AMI: Muthukkumaran Ramalingam, Stefano Righi

AT&T: Jitendra Patel, Bogdan Ungureanu

CableLabs: Tao Wan

Cisco: Matt Hyatt, Kori Rongey, Steve Vetter, Robin White

Dell Technologies: Dan Carroll

Intel: Steve Orrin

Keysight Technologies: Corey Piggott

MiTAC Computing Technology Corp.: Simon Hwang

The MITRE Corporation: John Kent, Theresa Suloway

NIST: Cheryl Pascoe, Adam Sedgewick, Kevin Stine

Nokia: Gary Atkinson, Rajasekhar Bodanki, Robert Cranston, Jorge Escobar, Don McBride

Palo Alto Networks: Aarin Buskirk, Bryan Wenger

T-Mobile: Todd Gibson

Overview

As specified by 3GPP standards, 5G systems use service-based architectures (SBAs) with a design that works well when implemented with cloud-native technologies leveraging microservices and container technology. A single 5G network function (NF) can be comprised of a multitude of containers running on many distributed servers. The 5G NFs communicate with each other over the network infrastructure, including carrier-grade routers and switches. The 5G Radio Access Network (RAN) components operate over a spread-out geographic area, while still requiring simultaneous connectivity to multiple types of 5G traffic.

Most network traffic in data centers and cloud environments flows over the same physical connections and is processed by the same network devices. Because physical separation is not feasible, methods for logically separating 5G traffic from other traffic and further separating types of 5G traffic from each other are needed to improve 5G cybersecurity and privacy.

What's the problem?

Data centers and cloud environments process and handle many types of traffic. Within the scope of 5G, the following types of 5G traffic are a starting point for logical separation:

- **Data Plane:** Transmitting user data, such as voice calls, video streaming, and internet browsing.
- **Signaling:** Setting up, maintaining, and tearing down communication sessions. It includes tasks like handovers, authentication, and resources allocation.
- **Operation and Maintenance (O&M):** Providing connectivity for the 5G network equipment, including software updates, fault detection, and performance optimization.

Each of these traffic types carries data with different sensitivity levels, and security and privacy implications if accessed by unauthorized or malicious users. O&M traffic provides access to devices that make up the 5G environment, allowing administrators important privileges and configuration capabilities, whereas the signaling traffic carries critical setup information, and the data plane carries user data. For example, the data plane segment is susceptible to distributed denial of service (DDoS) attacks on the N6 interface, the signaling segment can be impacted by signaling storms, and the O&M segment needs well-defined user and network access control.

A malicious actor can target the data plane, and if signaling and O&M traffic are not separated from the data plane, the malicious actor could target them as well. For example, the adversary can conduct privilege escalation and process injection for gaining administrative rights, attempt password cracking of valid user accounts on the nodes, exploit vulnerabilities in databases and file systems, and take advantage of improper configurations of routers and switches.¹

Another reason for separating the types of traffic is to prevent attackers from targeting one type to impact the performance of the others. For example, if an attacker overwhelms the data plane, and signaling and O&M traffic are not separated from the data plane, the attack can disrupt critical network functions and network management. Therefore, it is imperative to design the network infrastructure in a way that securely separates the different types of traffic.

The 5G standards defined by 3GPP do not specify cybersecurity and privacy protections for the underlying network infrastructure that support and operate the 5G system; these aspects are deemed implementation specific.² Mobile network operators make risk-based decisions on the countermeasures to mitigate attacks against their network.

¹ <https://fight.mitre.org/techniques/FGT1599/>

² <https://portal.3gpp.org/desktopmodules/Specifications/SpecificationDetails.aspx?specificationId=3169>

How can secure network design principles address the problem?

Network operators can use common network technologies to logically separate the 5G data plane, signaling, and O&M traffic from each other. Doing so can improve cybersecurity and privacy in multiple ways, including the following:

- **Reduce the risk of attacks spreading from one network segment to another.** This can limit the attackers' visibility and influence.
- **Make it easier to apply different cybersecurity requirements for each type of traffic.** For example, user data might need encryption for confidentiality, while signaling traffic requires integrity and authentication to prevent unauthorized access and control. O&M traffic needs to be protected to ensure that only authorized personnel can manage the network.

Logically separating the types of 5G traffic can also improve performance and manageability. It simplifies network management by allowing for tailored policies and monitoring for each type. Also, when issues arise, having a separate plane allows for easier identification and isolation of problems.

In addition to logical separation of traffic types, other technologies such as network slicing, policy-based network control, and traffic engineering can be applied to achieve additional security, performance, and manageability benefits.

The rest of this paper focuses on a method of achieving logical separation of traffic types in 5G networks: **virtual routing and forwarding (VRF)**. VRF is a technology used in networking to create multiple virtual instances of a routing table within a single physical router, thus avoiding the need for multiple physical sets of infrastructure (e.g., routers, fiber, power systems). VRF's key benefits include:

- **Traffic separation, isolation, and security:** VRF can ensure that different traffic types do not mix. Each VRF instance operates independently, maintaining its own routing table and forwarding decisions. By isolating routing tables, VRF enhances security. Even though the VRFs reside on the same physical network equipment, traffic from one VRF instance cannot be accessed by another, providing a secure environment for sensitive data.
- **Efficient use of resources:** VRF enables the use of a single physical router to support multiple virtual networks, reducing the need for additional hardware and simplifying network management.
- **Incremental growth:** New VRF instances can be added to the network as needed without requiring major changes to the existing infrastructure.

Figure 1 depicts how multiple VRFs are logically separated from each other while running within the same physical switch or router.




Figure 1. Logical View of VRFs

How can I use secure network design principles?

Architecting 5G network infrastructure is a complex process with many considerations. Defining requirements for performance and reliability early in the process is critical for having a functional network. There is an opportunity to also make security and privacy requirements foundational in initial planning phases. Including all of these requirements in the network design and procurement process helps to ensure that the equipment being deployed supports the technical mechanisms for logical separation.

Network engineers will need to appropriately implement, configure, and manage the various components. In addition to logical traffic separation (i.e., VRFs), other items that should be considered when configuring the network devices include, but are not limited to:

- **Security access controls**
 - Firewalling capabilities
 - Utilizing IP prefix lists
- **Cryptographic suites and algorithms supported**
- **Connectivity between devices**
 - Media and transceiver types
 - Correct usage of physical port types and speeds
 - Applicable routing protocols such as Border Gateway Protocol (BGP)

Implementing logical separation of 5G traffic types may also have functional and performance benefits such as the following:

- **Prioritization and Quality of Service (QoS):** Data plane requires high throughput and low latency to ensure a good user experience. Signaling is typically less bandwidth intensive but requires low latency and high reliability. O&M traffic is generally low in volume but is critical for maintaining network health and performance. Separating these planes allows for tailored QoS policies for each type of traffic.

- **Scalability:** Separation of traffic types allows network resources to be allocated more efficiently. For example, the data plane can be scaled independently to handle increased user data traffic without affecting the signaling or O&M traffic.
- **Reliability and Resilience:** Isolating the traffic types ensures that issues in one plane do not cascade and affect the others. For example, a surge in user data traffic should not overwhelm the signaling plane, which could lead to call drops.

Additional Technical Details

The rest of this white paper is intended for readers seeking more in-depth knowledge of NCCoE's network setup, including usage of VRF.

For background information on the NCCoE 5G Cybersecurity project, including the architecture and components of the 5G standalone network built within the demonstration lab environment, see NIST SP 1800-33 volume B, 5G Cybersecurity, Approach, Architecture, and Security Characteristics.³

Network Architecture

This section describes how the NCCoE set up its underlying network infrastructure to ensure the separation of data plane, signaling, and O&M traffic.

Figure 2 depicts the high-level architecture of the NCCoE's 5G network infrastructure. On the left side of the diagram is the user equipment (UE) (i.e., mobile devices using the 5G network) connecting to the 5G radio access network (RAN), connecting to the backhaul through a few routers, then connecting to the network infrastructure consisting of multiple routers. The network infrastructure connects the backhaul to the 5G core and management network and to the data network (DN). The network infrastructure is also connected to the cloud platform and the shared services. The data plane, signaling, and O&M VRFs have been set up in all the routers to separate these traffic types.

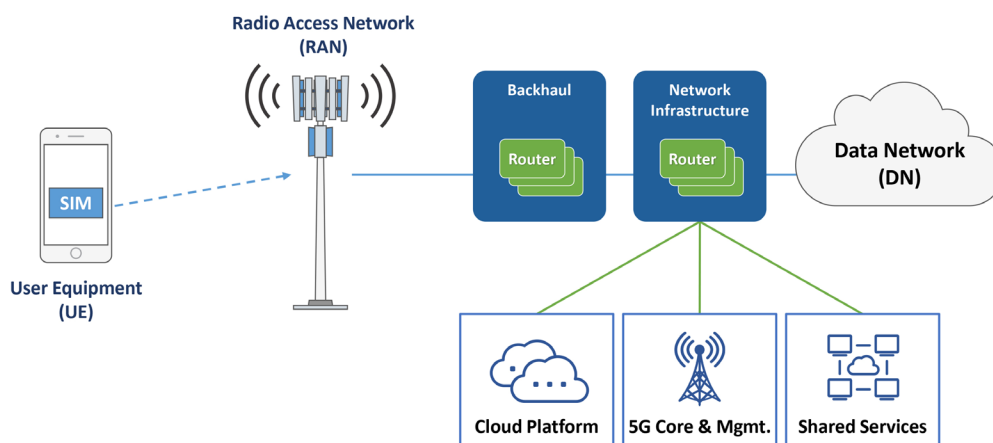


Figure 2. High-Level Network Infrastructure

³ <https://www.nccoe.nist.gov/sites/default/files/2022-04/nist-5g-sp1800-33b-preliminary-draft.pdf>

The network infrastructure supporting NCCoE's 5G system uses a spine-leaf architecture. This architecture is designed so that every leaf is connected to every spine to ensure that all leaf switches are no more than one hop away from one another. This provides multiple paths for traffic, supporting fault tolerance and minimizing latency. This architecture also provides scalability, where leaf switches can be added or removed to accommodate growth, while spine switches provide a scalable backbone.

The NCCoE architecture includes two spine switches with 40 GbE capabilities, two leaf switches with 100 GbE capabilities, and two leaf switches with 40 GbE capabilities. As depicted in Figure 3, each leaf switch connects to each spine switch to ensure that all leaf switches are no more than one hop away from one another. A scalable layer 3 routing protocol, Border Gateway Protocol (BGP), is used between the spine and leaf switches. Additionally, Equal Cost Multi-Path (ECMP) routing is used between the spine and leaf switches to load-balance traffic across the layer 3 network. The two leaf switch pairs, Leaf 1 and 2 as well as Leaf 3 and 4, use Virtual Link Trunking (VLT), allowing all connections to be active while also providing fault tolerance. The layer 3 and layer 2 boundary is at the leaf switches. The connections from the leaf switches to the spine switches are layer 3, whereas the connections from the leaf switches to the hosts are layer 2.

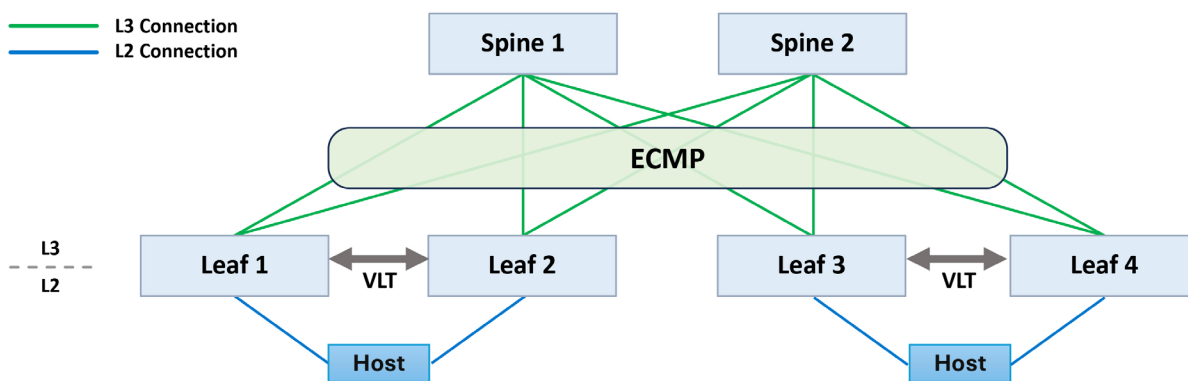


Figure 3. NCCoE's Lab Leaf and Spine Network Architecture

As depicted in Figure 4, the servers and storage are directly connected to the leaf switches. Leafs 1 and 2 are connected to Cloud Platform and 5G Core and Management hosts, using layer 2. Leafs 3 and 4 are connected to the Cloud Platform and Shared Services hosts, using layer 2.

The 5G RAN communicates with the 5G core through the cell site router, core aggregate router, and Leaf 1 and 2 switches. Leafs 1 and 2 use the External Border Gateway Protocol (EBGP) routing protocol to connect to the core aggregate router. Also, the Internal BGP (iBGP) routing protocol is used between the core aggregate router and the cell site router.

Link Aggregation Groups (LAGs) are used in the network to provide either increased link capacity or redundancy. For the backhaul connection in this project, a LAG is used between the cell site router and core aggregation router consisting of 2 x 1GbE links. The LAG is configured to use the default port threshold behavior, whereby all member ports must become inactive for the LAG to be declared down.

The VRFs are configured on all the spine and leaf switches for each of the 5G network types. The VRFs are also configured on Leafs 1 and 2 for the connectivity to the core aggregate router, and for the connectivity between the core aggregate router and the cell site router. This allows end-to-end separation of the data plane, signaling, and O&M traffic, ensuring there is no VRF leakage between these traffic types. Since the O&M and signaling communication is only required from the RAN to the 5G core, they do not traverse the entire network. The data plane goes through the whole network because it initiates in the RAN at the UE and must be able to get out to the data network.

The network testing device is connected from one end to the cell site router and on the other end to the User Plane Function (UPF) firewall to perform end-to-end testing. Leafs 3 and 4 are connected to the UPF firewall for connectivity to the internet, data network, and network testing devices.

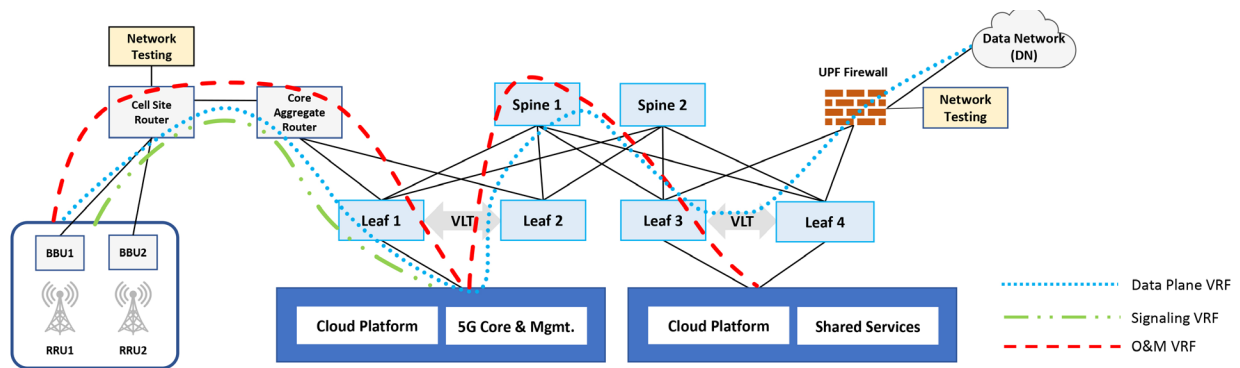


Figure 4. NCCoE's Lab Network Infrastructure

Network Configuration

Figure 5 shows portion of NCCoE's lab Leaf 1 configuration showcasing the data plane implementation as one VRF example. It includes the configuration of data plane VRF, IP prefix list, and VLANs, as well as the configuration of the interfaces to the spines and the core aggregate router, and the data plane BGP setup. Only the data plane related configuration is shown, and the rest of the configuration is omitted for abbreviation. The VLAN numbers were administrator defined and unique to each switch.

5G Network Security Design Principles: Applying 5G Cybersecurity and Privacy Capabilities

```
Leaf1#
!
ip vrf DataPlane
!
! Omitted for abbreviation
!
ip prefix-list DataPlane-vrf seq 10 permit 10.0.0.0/22
ip prefix-list DataPlane-vrf seq 20 permit 172.16.112.0/23
ip prefix-list DataPlane-vrf seq 30 permit 192.168.101.8/29
!
! Omitted for abbreviation
!
interface vlan113
 mode L3
 description "sub-interface to Spine 1 port 1/51 DataPlane-vrf"
 no shutdown
 ip vrf forwarding DataPlane
 ip address 192.168.1.17/31
!
interface vlan213
 mode L3
 description "sub-interface to Spine 2 port 1/51 DataPlane-vrf"
 no shutdown
 ip vrf forwarding DataPlane
 ip address 192.168.2.17/31
!
interface vlan313
 mode L3
 description "sub-interface to Nokia 7750 Port 1/1/10 DataPlane-vrf"
 no shutdown
 ip vrf forwarding DataPlane
 ip address 192.168.3.8/31
!
! Omitted for abbreviation
!
interface loopback2
 description "Router ID 5G-S4048-Leaf1 DataPlane vrf"
 no shutdown
 ip vrf forwarding DataPlane
 ip address 10.0.2.9/32
!
! Omitted for abbreviation
!
interface ethernet1/1/31:1
 description "to Spine 1 port 1/51"
 no shutdown
 switchport mode trunk
 switchport access vlan 1
 switchport trunk allowed vlan 111,113,115
 mtu 9216
 flowcontrol receive off
 spanning-tree disable
!
interface ethernet1/1/32:1
 description "to Spine 2 port 1/51"
 no shutdown
 switchport mode trunk
 switchport access vlan 1
 switchport trunk allowed vlan 211,213,215
 mtu 9216
 flowcontrol receive off
 spanning-tree disable
```

Setting up Data Plane VRF

Setting up Data Plane VRF prefix list

Setting up VLANs for connectivity from Leaf 1 to Spine 1, Spine 2 and Core Aggregate router for Data Plane VRF

Setting up Leaf 1's interfaces for Data Plane VRF loopback and for connectivity to Spine 1, and Spine 2

```
interface ethernet1/1/33
description "to Nokia 7750 Port 1/1/10"
no shutdown
switchport mode trunk
switchport access vlan 1
switchport trunk allowed vlan 311-313
mtu 9216
flowcontrol receive off
!
! Omitted for abbreviation
!
route-map DataPlane-vrf permit 30
match ip address prefix-list DataPlane-vrf
!
! Omitted for abbreviation
!
router bgp 65500
bestpath as-path multipath-relax
graceful-restart role receiver-only
maximum-paths ebgp 2
!
vrf DataPlane
bestpath as-path multipath-relax
graceful-restart role receiver-only
maximum-paths ebgp 2
!
address-family ipv4 unicast
redistribute connected route-map DataPlane-vrf
!
template DataPlane-vrf
advertisement-interval 1
fall-over
timers 3 9
!
neighbor 192.168.1.16
description "EBGP to Spine 1 DataPlane-vrf"
remote-as 65510
no shutdown
!
neighbor 192.168.2.16
description "EBGP to Spine 2 DataPlane-vrf"
remote-as 65511
no shutdown
!
neighbor 192.168.3.9
description "EBGP to Nokia 7750 DataPlane-vrf"
remote-as 65501
no shutdown
!
! Omitted for abbreviation
```

Setting up Leaf 1's interface for connectivity to Core Aggregate router

Setting up route-map for Data Plane VRF

Setting up BGP routing from Leaf t to Spine 1, Spine 2 and Core Aggregate router for Data Plane VRF

Figure 5. Sample Configuration of Leaf 1 with Focus on Data Plane VRF

Confirming Data Plane, Signaling, and O&M Traffic Separation

The NCCoE lab uses the Cisco Secure Network Analytics (SNA) network tool to monitor data plane, signaling, and O&M traffic. Custom reports can be configured to show VRF separation at a high level, as well as alerts if there is any leakage between VRFs. Figure 6 shows that SNA can display the three VRFs, and there is no leakage between them. If there were any leakage between the VRFs, one or more of the circles would turn red with alerts being shown. The VRF leakage could lead to unintended routing of traffic between isolated network segments, compromising security and operational disruptions.

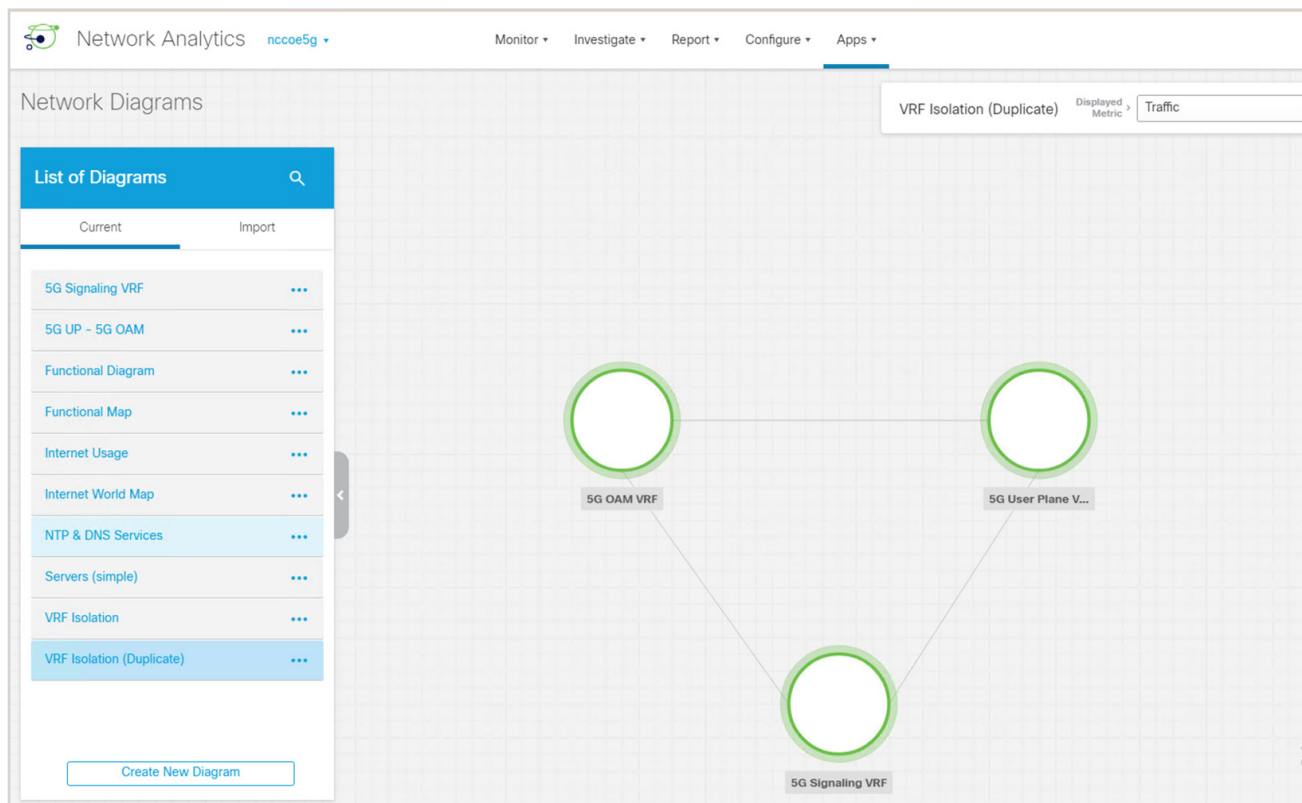



Figure 6. Cisco Secure Network Analytics (SNA) Displaying Three VRFs



Certain commercial equipment, instruments, software, or materials, commercial or non-commercial, are identified in this paper in order to specify the experimental procedure adequately. Such identification does not imply recommendation or endorsement of any product or service by NIST, nor does it imply that the materials or equipment identified are necessarily the best available for the purpose.

NIST Technical Series Policies

[Copyright, Use, and Licensing Statements](#)

[NIST Technical Series Publication Identifier Syntax](#)

Author ORCID iDs

Michael Bartock: 0000-0003-0875-4555

Jeffrey Cichonski: 0009-0006-1137-2549

Karen Scarfone: 0000-0001-6334-9486

Murugiah Souppaya: 0000-0002-8055-8527

How to Cite this NIST Technical Series Publication:

Bartock M. et al. (2025) 5G Network Security Design Principles: Applying 5G Cybersecurity and Privacy Capabilities. (National Institute of Standards and Technology, Gaithersburg, MD), NIST Cybersecurity White Paper (CSWP) NIST CSWP 36E ipd. <https://doi.org/10.6028/NIST.CSWP.36E.ipd>

Public Comment Period

June 17, 2025 – July 17, 2025

Submit Comments

5g-security@nist.gov

Or submit the web form at <https://www.nccoe.nist.gov/5g-cybersecurity>

National Institute of Standards and Technology

Attn: Applied Cybersecurity Division, Information Technology Laboratory

100 Bureau Drive (Mail Stop 2000)

Gaithersburg, MD 20899-2000

Additional Information

Additional information about this publication is available at

<https://www.nccoe.nist.gov/5g-cybersecurity>, including related content, potential updates, and document history.

All comments are subject to release under the Freedom of Information Act (FOIA).