

Dell Technologies 5G Core Validated Design with Oracle and VMware

June 2022

H19171

Reference Architecture Guide

Abstract

This reference architecture guide describes a Dell 5G solution incorporating Oracle Communications Core with VMware on-premises infrastructure and Dell storage and networking. The solution is designed to help communications service providers quickly deploy 5G services.

Dell Technologies Solutions



Copyright

The information in this publication is provided as is. Dell Inc. makes no representations or warranties of any kind with respect to the information in this publication, and specifically disclaims implied warranties of merchantability or fitness for a particular purpose.

Use, copying, and distribution of any software described in this publication requires an applicable software license.

Copyright © 2022 Dell Inc. or its subsidiaries. Published in the USA 06/22 Reference Architecture Guide H19171.

Dell Inc. believes the information in this document is accurate as of its publication date. The information is subject to change without notice.

Contents

Chapter 1	Introduction	5
	Introduction.....	6
	Solution overview	6
	We value your feedback	7
Chapter 2	Physical Design	8
	Overview.....	9
	Physical network design	9
	Hardware	10
	Network	18
Chapter 3	Cluster Design	24
	VMware Telco Cloud Automation	25
	Telco Cloud Automation architecture	25
	Sample workflow for infrastructure and CaaS automation	27
	Deployment architecture.....	28
	Services design	30
	Physical design.....	31
	Storage considerations for 5G Core	31
	Tanzu Kubernetes cluster design	32
	Tanzu Basic for 5G Core deployment model.....	32
	CNF design.....	33
	Kubernetes cluster deployment process	34
	Telco workload deployment.....	34
Chapter 4	Oracle 5G Core Design	36
	Oracle 5G Core	37
	NF architecture	40
	5G Core architecture	41
	NF deployment on the cloud infrastructure.....	43
	5GC NFs on VMware Telco Cloud Platform.....	45
	Observability tools	48
Chapter 5	Solution Validation	50
	Testing scope and setup	51
	Use cases.....	51

Chapter 6	Foundation of CSP Grade 5G Core	60
Network deployment options		61
Security.....		62
Conclusions		63
Chapter 7	References	65
Dell Technologies documentation		66
Oracle documentation		66
Appendix A	Hardware and Software Configuration	67
Overview.....		68

Chapter 1 Introduction

This chapter presents the following topics:

Introduction	6
Solution overview	6
We value your feedback	7

Introduction

The Dell Technologies 5G Core reference architecture designed with Oracle Communications and VMware is intended to help communications service providers (CSPs) to quickly deploy 5G services. This reference architecture features a 5G Core platform, a container management and orchestration platform, and telecommunications-grade infrastructure in a fully validated solution design.

The reference architecture provides:

- A solution that was developed by Dell Technologies, Oracle Communications, and VMware to help CSPs accelerate 5G Core network deployments while reducing the costs and risks associated with network transformation efforts.
- A trusted and validated industry-leading foundation that gives CSPs the flexibility and reliability that they require to move forward with confidence.
- An approach to delivering leading-edge telecommunications (telco) technology on a proven platform for a true competitive advantage.

The solution includes the following components:

- VMware Telco Cloud Automation, providing consistent multicloud operations from infrastructure to services.
- Oracle Communications 5G Core network software, a cloud-native 5G technology from the network through Business/Operation Support Systems (B/OSS) applications. Oracle 5G Core uses converged policy and charging to capitalize on network slicing and to support new business models.
- Dell PowerEdge R640, R650, and R750 servers, providing high availability (HA) in mobile core and harsher edge environments with customizable hardware acceleration options.
- Dell PowerSwitch S3048-ON and S5232F-ON networking switches, optimized for high-performance data center environments.

Solution overview

The 5G Core solution introduces innovative and disruptive networking paradigms to mobile networks. Its validated design brings tested and proven configurations that are designed to dynamically fit the needs of CSPs. By offering flexible design options and guidance on choosing the right partners and components, this solution can shorten deployment timelines, reducing (or, in some cases, eliminating) the time it takes to design, test, and integrate components from multiple partners.

The 5G future will not be built by a single vendor. It calls for a collaboration in innovation from the industry leaders in cloud, telco, containers, and other technologies. Dell Technologies, Oracle Communications, and VMware are taking a step forward toward that future with this design for a fully containerized, cloud-native 5G solution. Our goal is to provide the CSPs of today with a fully validated, accelerated path to the 5G services of tomorrow.

Key benefits

The Dell 5G Core validated design combines industry-leading technology from Oracle Communications, VMware, and Dell to offer customers the following benefits:

- Oracle Communications 5G Core takes advantage of automation to bring new services to market faster, and delivers converged policy and charging, signaling, routing, and network slice selection to capitalize on new business models.
- VMware Telco Cloud Platform deploys cloud-native and virtual network functions (NFs) consistently, at web-scale speed, and without disruption.
- Dell's telco-grade PowerEdge and PowerSwitch infrastructure supports high-density, high-performance 5G Core and edge workloads.
- Dell Technologies, Oracle Communications, and VMware provide consulting services for the solution.

Document purpose

This guide addresses the following topics:

- Physical design of the 5G solution
- VMware cluster design
- Oracle 5G Core design
- 5G Core validation in our labs
- Foundation of CSP Grade 5G Core

Audience

This document is intended for network and system architects and system administrators who must deploy an Oracle Communications 5G Core solution running on VMware Telco Cloud Platform and Dell PowerEdge servers and PowerSwitch switches. Some experience with containers and Kubernetes platforms is recommended.

We value your feedback

Dell Technologies and the authors of this document welcome your feedback on the solution and the solution documentation. Contact the Dell Technologies Solutions team by [email](#).

Note: This guide may contain language from third-party content that is not under Dell's control and is not consistent with Dell's current guidelines for Dell's content. When this content is updated by the relevant third parties, this guide will be revised accordingly.

Chapter 2 Physical Design

This chapter presents the following topics:

- Overview..... 9**
- Physical network design..... 9**
- Hardware 10**
- Network 18**

Overview

This chapter describes the physical components and cabling that are used in the reference architecture.

Physical network design

The following figure shows the distribution of physical components and cabling in the network design:

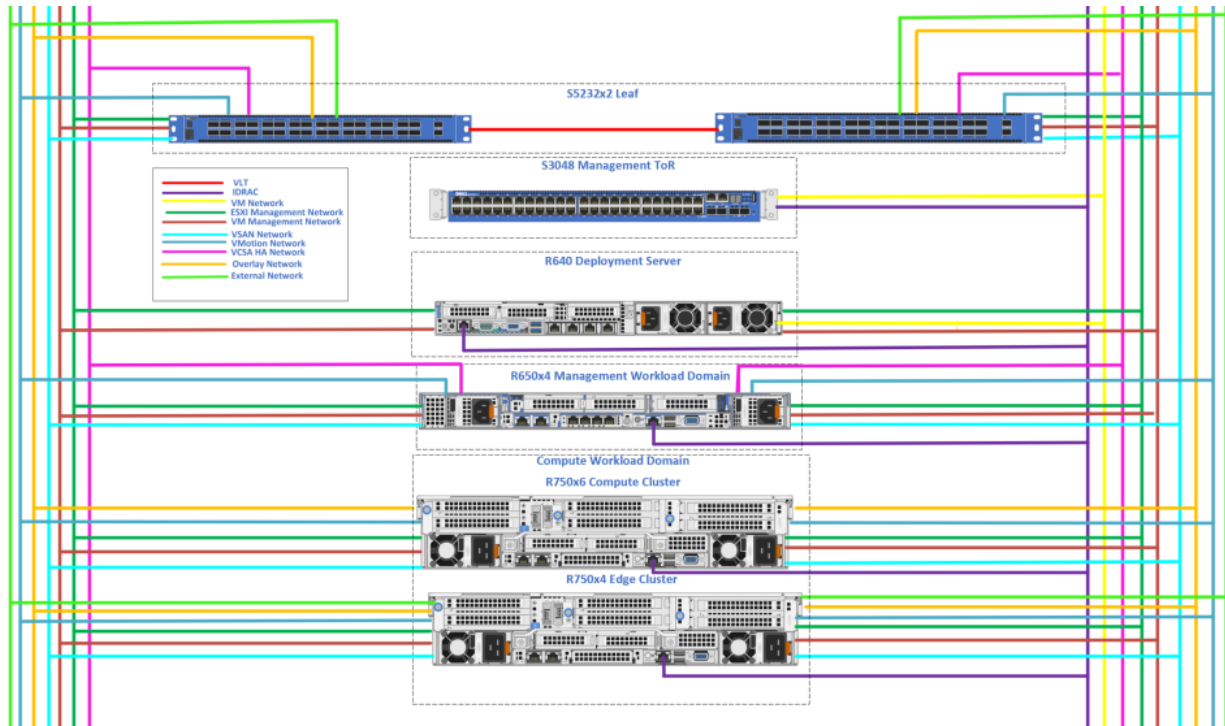


Figure 1. Network connectivity

One stamp (cluster of servers) depicts our 5G Core installation. The stamp hosts:

- Two PowerSwitch S5232F-ON leaf switches. The leaf switches provide connectivity between the endpoints and the data center. Leaf switches are configured as a Virtual Link Trunking (VLT) pair that enables all connections to be active while providing fault tolerance.
- Four servers dedicated to a VMware management workload domain that is running on PowerEdge R650 servers. All the VMware management functions, including Telco Cloud Automation and Harbor, will be running on this cluster.
- Six servers forming a VMware compute cluster and running on PowerEdge R750 servers. All Oracle NFs will run on this cluster.
- Four PowerEdge R750 servers dedicated for an edge cluster.
- One PowerSwitch S3048-ON switch for out-of-band (OOB) connectivity of all iDRAC server ports.

- One PowerEdge R640 server for deployment. This server is used to deploy the solution from a single point that connects to the OOB network. This point is accessible from the corporate network and the in-band network (the management network and leaf switches). The server has a Windows virtual machine (VM) that provides Remote Desktop Protocol (RDP) access to the management workload domain and compute cluster.

Hardware

This section describes the hardware that comprises the Dell Telco Cloud Platform.

Servers

PowerEdge R750 rack server

The PowerEdge R750 rack server is an enterprise server that is powered by third-generation Intel® Xeon® Scalable processors to deliver outstanding performance for the most demanding workloads. The PowerEdge R750, a two-socket/2U rack-mounted server, is ideal for a data center with high performance computing (HPC) workloads.

The following figure shows a front view of the PowerEdge R750 server:



Figure 2. PowerEdge R750 server (front view)

The following figure shows a rear view of the R750 server:




Figure 1. PowerEdge R750 server (rear view)

The following tables show the recommended hardware configuration for the PowerEdge R750 server in the solution:

Table 1. Recommended hardware configuration for the PowerEdge R750 server (compute)

Hardware	Description
Quantity	6 servers
Compute chassis	2U R750 2.5-in. chassis with up to 24 SAS/SATA drives
Intended use	Oracle infrastructure

Hardware	Description
Riser configuration	Riser Config 2, full length, 4x16, 2x8 slots, DW GPU capable
CPU	Two Intel Xeon Gold 6348 2.6 GHz 28 cores 56 threads per CPU 56 cores 112 threads total
Embedded NIC 1	Broadcom Gigabit Ethernet BCM5720
Integrated NIC 1	Intel® Ethernet 25G 2P E810-XXV OCP
NIC slot 4	Intel Ethernet 25G 2P E810-XXV Adapter
NIC slot 5	Intel Ethernet 25G 2P E810-XXV Adapter
NIC slot 7	Mellanox ConnectX-5 EN 25 GbE Dual-port SFP28 Adapter
Standard memory	8 x 32 GB DDR4 (256 GB) per server
Storage	8 x HDD 900 GB 15KRPMs SAS drives, 2 x SSDs x 960 GB cache
Storage controller	Dell HBA355i No RAID
IDSDM card reader	Present with dual 1 6GB micro_SDHC/SDXC Card for ESXi operating system installation
Storage (aux)	BOSS-S2 controller cards
Power supply	Dual, hot plug, redundant power supply (1+1), 1400 W, mixed mode
System management	iDRAC data center
Security	TPM 2.0-NTC (optional)
Fans	6x STD
Operating system	VMware ESXI 7.0 U1
BIOS	1.2.4
Lifecycle controller	5.10.00.00
Storage controller firmware	15.15.15.00

Table 2. Recommended hardware configuration for the PowerEdge R750 (edge)

Hardware	Description
Quantity	4 servers
Edge chassis	2U R750 2.5-in. chassis with up to 24 SAS/SATA drives
Intended use	Oracle infrastructure
Riser config	Riser Config 2, full length, 4x16, 2x8 slots, DW GPU-capable

Hardware	Description
CPU	Two Intel Xeon Gold 6348 2.6 GHz 28 cores, 56 threads per CPU 56 cores, 112 threads total
Embedded NIC 1	Broadcom Gigabit Ethernet BCM5720
Integrated NIC 1	Intel Ethernet 25G 2P E810-XXV OCP
NIC slot 5	Intel Ethernet 25G 2P E810-XXV Adapter
Standard memory	8 x 32 GB DDR4 (256 GB) per server
Storage	8 x HDD 900 GB 15KRPMs SAS drives, 2 x SSDs x 960 GB cache
Storage controller	Dell HBA355i no RAID
IDSDM card reader	Present with dual 16 GB micro_SDHC/SDXC Card for ESXi operating system installation
Power supply	Dual, hot plug, power supply redundant (1+1), 1400 W, mixed mode
System management	iDRAC data center
Security	TPM 2.0-NTC (optional)
Fans	6 x STD
Operating system	VMware ESXI 7.0 U1
BIOS	1.2.4
Lifecycle controller	5.10.00.00
Storage controller firmware	15.15.15.00

PowerEdge R650 rack server

The PowerEdge R650 rack server, an enterprise server that is powered by third-generation Intel Xeon Scalable processors, is designed to optimize workload performance and data center density. The PowerEdge R650 is a 1U rack-mounted server.

The following figure shows a front view of the PowerEdge R650 server:



Figure 2. PowerEdge R650 server (front view)

Table 3. Recommended hardware configuration for the PowerEdge R650 server

Hardware	Description
Quantity	4 servers
Management chassis	1U R650 2.5-in. chassis with up to 10 hard drives (SAS/SATA) including a maximum of 4 universal drives, 3 PCIe Slots
Intended use	Oracle infrastructure
Riser config	Riser config 0, 2 CPU, half length, low profile, 3 x16 slots, SW GPU-capable
CPU	Two Intel Xeon Gold 6348 2.6 GHz 28 cores 56 threads per CPU 56 cores 112 threads total
Embedded NIC 1	Broadcom Gigabit Ethernet BCM5720
Integrated NIC 1	Intel Ethernet 25G 2P E810-XXV OCP
NIC Slot 2	Intel Ethernet 25G 2P E810-XXV adapter
Standard memory	8 x 32 GB DDR4 (256 GB) per server
Storage	6 x HDD 900 GB 15KRPMs SAS drives, one SSD x 960 GB cache
Storage controller	Dell HBA355i no RAID
IDSDM card reader	Present with dual 16 GB micro SDHC/SDXC card for ESXi operating system installation
Power supply	Dual, hot plug, power supply redundant (1+1), 1400 W, mixed mode
System management	iDRAC data center
Security	TPM 2.0 -NTC (optional)
Fans	4 x HP
Operating system	VMware ESXI 7.0 U1
BIOS	1.2.4
Lifecycle controller	4.40.29.00
Storage controller firmware	15.15.15.00

The following figure shows a rear view of the R650 server:




Figure 3. PowerEdge R650 server (rear view)

PowerEdge R640 rack server

The PowerEdge R640 is a scalable computing and storage server in a 1U, two-socket platform with an ideal mix of performance, cost, and density for most data centers. The R640 server can address demanding workloads such as virtualization, dense private cloud, HPC, and software-defined storage.

The following figure shows a front view of the PowerEdge R640 server:



Figure 4. PowerEdge R640 server (front view)

Table 4. Recommended hardware configuration for the PowerEdge R640 server

Hardware	Description
Quantity	1 server
Deployment chassis	1U R640 2.5 chassis with up to 8 HD, 3 PCIe
Intended use	Oracle deployment
Riser config	Riser Config 2, 3x16 LP
CPU	<ul style="list-style-type: none"> Two Intel Xeon Gold 6240 2.6 GHz 18 cores, 36 threads per CPU 36 cores, 72 threads total
Integrated NIC 1	Intel Gigabit 4P I350-t rNDC
NIC slot 1	Mellanox ConnectX-4 LX 25 GbE SFP adapter
NIC slot 2	Mellanox ConnectX-4 LX 25 GbE SFP adapter
Standard memory	12 x 16 GB DDR4 (192 GB) per server
PCIe slot 1	MT27710 Family [ConnectX-4 Lx]
PCIe slot 2	MT27710 Family [ConnectX-4 Lx]
Storage	900 GB, HDD 15 K SAS, 12 Gb, 512 n, 2.5, HP 960G SSD SAS, MU, 12, 2.5, HP, PX05SV
Storage controller	PERC H740P Mini
IDSDM card reader	Present with dual 16 GB micro_SDHC/SDXC Card for ESXi operating system installation
Power supply	Dual, redundant, hot-plug PS, 750 W
System management	iDRAC data center
Fans	8 x STD
Operating system	VMware ESXI 7.0 U1

Hardware	Description
BIOS	2.11.2
Lifecycle controller	5.00.00.00
Storage controller firmware	51.14.0-3900

The following figure shows a rear view of the PowerEdge R640 server:




Figure 5. PowerEdge R640 server (typical configuration)

Network interface cards

Broadcom Gigabit Ethernet BCM5720

The Broadcom 5720 Dual Port 1Gb Network Interface Card (NIC) connects your desktop and server to your network. Remote management support enables you to maximize management resources. This NIC has a compact design and is ideal for high-performance network applications.

The following figure shows the Broadcom 5720 Dual Port 1Gb NIC:



Figure 6. Broadcom 5720 Dual Port 1Gb NIC

Intel Ethernet PCI-Express NIC dual port server adapter

The Ethernet PCI-Express NIC dual port server adapter connects your server to your network. The dual port server adapter is a reliable, standards-based solution featuring a compact space-saving design. It is ideal for high-performance network applications.

The following figure shows the adapter:



Figure 7. Intel Ethernet 25G 2P E810-XXV OCP network adapter

Mellanox ConnectX-5 dual port 10/25GbE SFP28 adapter

The Mellanox ConnectX-5 dual port 10/25gbe sfp28 adapter connects your server to your network. The adapter has been tested and validated on Dell systems and is supported by Dell Technical Support when used with a Dell system.

The following figure shows the adapter:



Figure 8. Mellanox ConnectX-5 dual port 10/25gbe sfp28 adapter

Intel Gigabit 4P I350-t rNDC

The Intel Gigabit 4P I350-t rNDC Ethernet PCI-Express NIC in a low-profile form factor connects your server to your network. The Gigabit Ethernet quad port server adapter

has been tested and validated on Dell systems and is supported by Dell Technical Support when used with a Dell system.

The following figure shows the adapter:



Figure 9. Intel Gigabit 4P I350-t rNDC

Mellanox ConnectX-4 LX 25GbE SFP adapter

Mellanox ConnectX-4 LX dual-port 25GbE SFP28 low-profile gigabit Ethernet NICs deliver high-bandwidth, industry-leading connectivity for performance-driven server and storage applications in enterprise data centers, Web 2.0, HPC, and embedded environments. Mellanox Ethernet adapters provide dedicated adapter resources that guarantee isolation and protection for VMs in the server.

By enabling a single physical NIC to appear as multiple virtual NICs, Mellanox Ethernet adapters give data center managers better server utilization for local area network (LAN) and storage-attached network (SAN) unification, while reducing power and cable complexity and capital expenditure. New large-scale cloud environments require the implementation of overlay network protocols to overcome the issues of security, isolation, and virtual LAN (VLAN) limitations in the cloud.

The following figure shows the Mellanox Ethernet adapter:



Figure 10. Mellanox ConnectX-4 LX 25 GbE SFP adapter

Network

The following section of the guide describes the network components and design concepts of the 5G installation.

Switches

PowerSwitch S3048-ON

The PowerSwitch S3048-ON 1000BASE-T top-of-rack (ToR) switch is the industry's first 1GbE enterprise switching platform to deliver both an industry-hardened operating system and support for open networking that accommodates third-party operating systems. This open networking platform is built for high-performance software-defined data centers, allowing users to run traditional workloads and deploy new workloads. By using the S3048-ON switch, you can run operating system options that are optimized for diverse deployment needs on a common hardware platform and architecture.

The following figure shows the PowerSwitch S3048-ON switch:



Figure 11. PowerSwitch S3048-ON switch

PowerSwitch S5232F-ON

PowerSwitch S5232F-ON 25/100GbE fixed switches provide state-of-the-art high-density 25/100 GbE ports and a broad range of functionality to meet the growing demands of today's telco data center environment. The compact, single-RU, S5232F-ON switch provides performance and flexibility for a variety of network designs. In high-density deployments, you can use the S5232F-ON switch with breakout cables to provide up to 128 10 GbE or 25 GbE ports.

The minimum software release that is required for PTP support is OS10 Release 10.5.2.6, which supports IEEE-1588v2 with the G.8275.1 telco profile and one-step timing.

Note: The PowerSwitch S5232F-ON switch does not support SyncE timing on SFP+ ports 33 and 34. SyncE is supported only on certain optics. For a list of SyncE-compatible devices, see Appendix III of the [ITU G.8262](#) standard.

The following figure shows the PowerSwitch S5232F-ON switch:



Figure 12. PowerSwitch S5232F-ON

Virtual network configuration

The vSphere Distributed Switch (VDS) is configured on the hosts in each tier, providing a similar network configuration across multiple hosts. Each ESXi host has multiple port groups that are configured and allotted to specific VLANs. The VDS has a separate uplink connection for the physical data center network, separating uplink traffic from other network traffic. These uplinks are mapped with a pair of physical NICs on each host to connect the virtual switches to the physical switch.

The management workload domain consists of infrastructure and VM networks, as shown in the following figure:

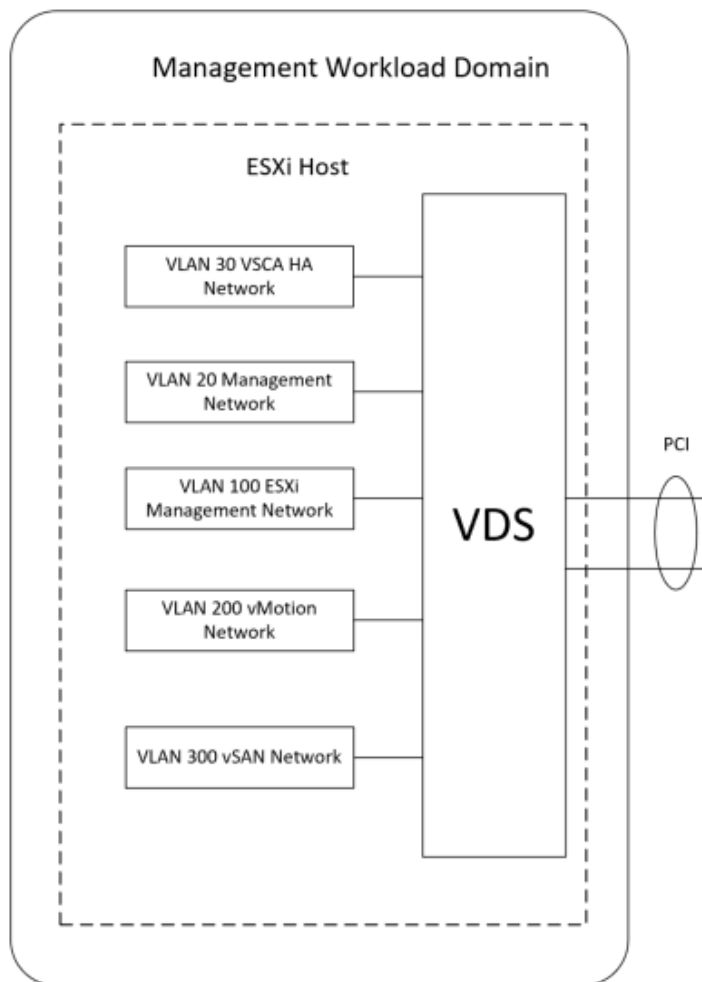


Figure 13. Management workload domain VDS and VLANs

The edge cluster consists of infrastructure, management, and workload networks, as shown in the following figure:

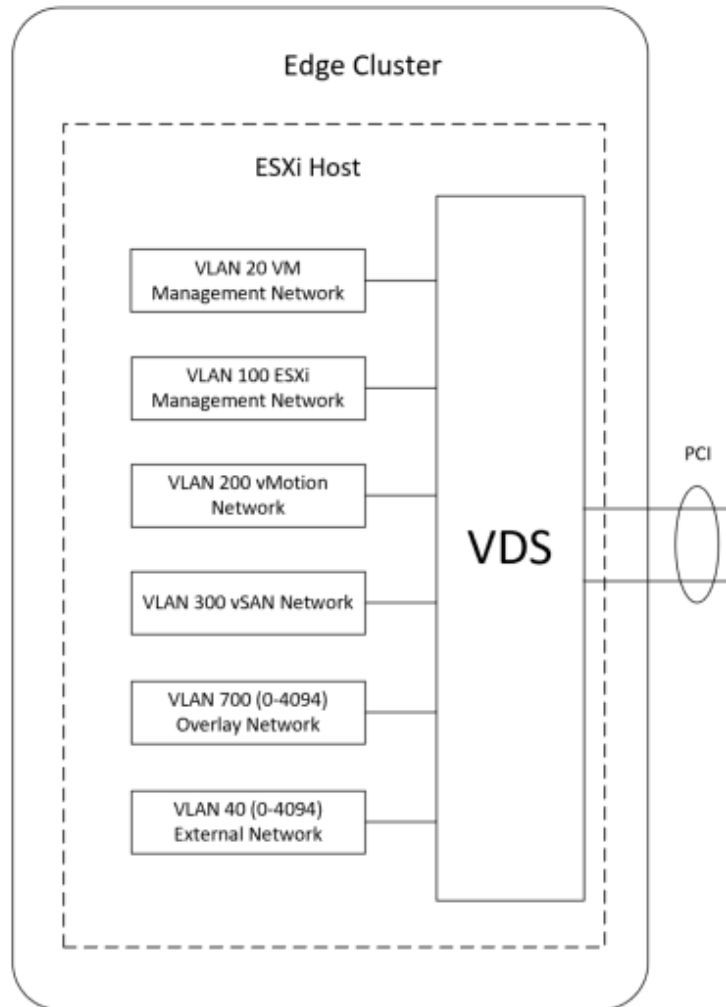


Figure 14. Edge cluster VDS and VLANs

The compute cluster consists of networks that are required to deploy tenants, as shown in the following figure:

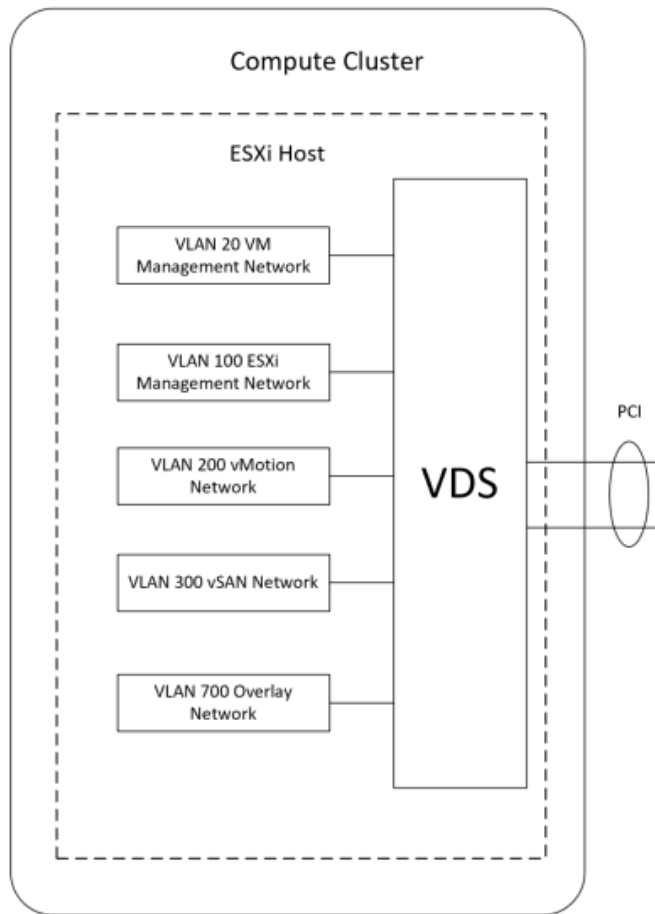


Figure 15. Compute cluster VDS and VLANs

Networks by function

This section provides a description of the networks that are used to support this deployment.

VLT

Leaf switches are configured as a VLT pair.

iDRAC

The integrated Dell Remote Access Controller (iDRAC) network is used for configuring host BIOS parameters, installing host software, and providing connectivity to the TCA and vCenter user interfaces. The iDRAC is a piece of hardware that sits on the server motherboard and enables systems administrators to update and manage Dell systems, even when the server is turned off. The iDRAC provides both a web interface and a command-line interface through which administrators can perform remote management tasks.

The networks that are used to deploy 5G Core are listed in the following table and described below:

Table 5. 5G Core deployed networks

Network	VDS/VSS group	Interface type	VLANs	Subnets	Mask	Gateway
VM network	vSwitch0	VLAN	0	Default	Default	Default
ESXI management network	Infra_Network_VDS	VLAN	100	192.168.100.100	/24	192.168.100.254
VM management network	VM_Network_VDS	VLAN	20	192.168.20.100	/24	192.168.20.254
VSAN network	Infra_Network_VDS	VLAN	300	192.168.3.100	/24	192.168.3.254
VMotion network	Infra_Network_VDS	VLAN	200	192.168.2.100	/24	192.168.2.254
VCSA A network	VM_Network_VDS	VLAN	30	192.168.30.100	/24	192.168.30.254
Overlay network	Edge_VM_Network_VDS	VLAN	700	192.168.7.100	/24	192.168.7.254
External network	Edge_VM_Network_VDS	VLAN	40	172.16.60.100	/24	172.16.60.254

VM network

The VM network enables VMs to communicate with one another, which means that VMs running on an ESXi host can connect to the virtual and physical network.

ESXI management network

This network is used for ESXi host management traffic. It enables login into individual ESXi UIs to view or change configurations.

Note: VMware ESXi is an enterprise-class, type-1 hypervisor that VMware has developed for deploying and serving virtual computers.

VM management network

The VM management network is where the VM connects to vCenter Server and ESXi hosts.

vSAN network

This network is used for vSAN shared storage traffic. VMware vSAN is a software-defined storage product that is used in collaboration with the VMware ESXi hypervisor. Software-defined storage products provision and manage storage based on policies, regardless of the underlying hardware.

vMotion network

This network is used for vSphere vMotion traffic. vMotion enables the live migration of running VMs from one physical server to another with zero downtime, continuous service availability, and complete transaction integrity. vMotion is a key enabling technology for creating a dynamic, automated, and self-optimizing data center.

VCSA HA network

VMware vCenter HA protects the vCenter Server Appliance (VCSA) against host and hardware failures. The active-passive architecture of the solution helps reduce downtime when you patch VCSA. The VCSA HA feature works as a cluster of three VMs containing active, passive, and witness nodes.

Overlay network

GENEVE (Generic Network Virtualization Encapsulation) provides the overlay capability in VMware NSX-T to create isolated, multitenant broadcast domains across data center fabrics. GENEVE enables you to create elastic, logical networks that span physical network boundaries, providing extensibility while still using the offload capabilities of NICs for performance improvement. Using the GENEVE overlay, VMware NSX-T isolates the network into a pool of capacity and separates the consumption of these services from the underlying physical infrastructure. You can organize the pool of network capacity in logical networks that are directly attached to specific applications. For more information, see [What is GENEVE?](#)

External network

The edge node in VMware NSX-T uses the external network for north-south peering with external devices.

Note: NSX-T data center is the software-defined networking component for Telco Cloud Platform. NSX-T enables you to create, delete, and manage software-based virtual networks.

Chapter 3 Cluster Design

This chapter presents the following topics:

VMware Telco Cloud Automation	25
Telco Cloud Automation architecture	25
Sample workflow for infrastructure and CaaS automation	27
Deployment architecture	28
Services design	30
Physical design	31
Storage considerations for 5G Core	31
Tanzu Kubernetes cluster design	32
Tanzu Basic for 5G Core deployment model	32
CNF design.....	33
Kubernetes cluster deployment process	34
Telco workload deployment	34

VMware Telco Cloud Automation

VMware Telco Cloud Automation is a unified orchestrator that onboards and orchestrates workloads seamlessly from VMs and container-based infrastructures. Telco Cloud Automation distributes workloads from the core to the edge and from private to public clouds for unified orchestration.

VMware Telco Cloud Platform is a common platform spanning core and RAN functionality. The platform self-tunes automatically depending on the workload that is deployed through it. All VNFs/CNFs from 5G Core to RAN are deployed using the same automation platform, operational tools, and Container-as-a-Service (CaaS) layer based on [VMware Tanzu](#).

The following figure shows the VMware Telco Cloud Platform architecture:

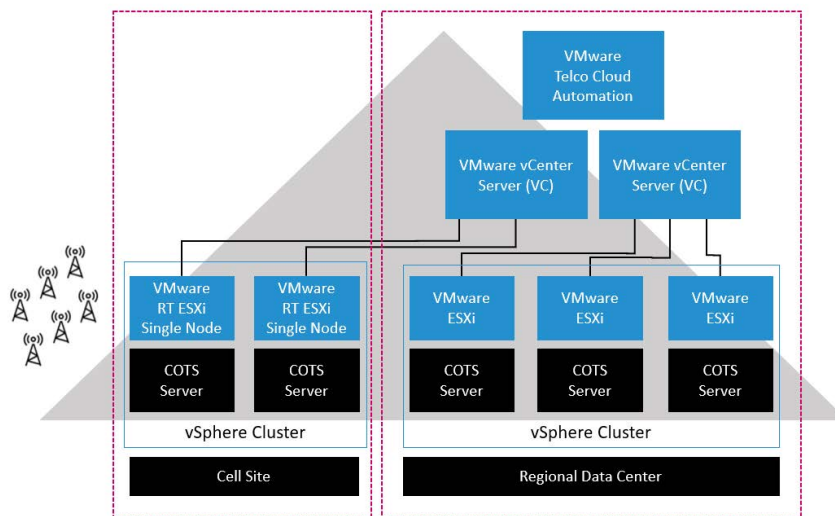


Figure 16. VMware Telco Cloud Platform architecture

Virtual platform infrastructure

The virtual infrastructure, the foundation of the platform, contains software-defined infrastructure, software-defined networking, and software-defined storage.

Compute workload domain

With Telco Cloud Platform core architecture, you can deploy and extend core applications into multiple data centers. The compute workload domain that you deploy in data centers using Telco Cloud Automation can contain multiple vSphere hosts, each running the user workloads to enable the 5G Core solution. These hosts can start with one ESXi host and then scale based on the resource and availability requirements of the solution that is being deployed.

Telco Cloud Automation architecture

VMware Telco Cloud Automation consists of two components: Telco Cloud Manager and a Telco Cloud Automation control plane. Telco Cloud Manager provides telco operators with NFV MANO capabilities and enables the automation of deployment and configuration

of the NFs and network services. The Telco Cloud Automation control plane provides the infrastructure for placing workloads across the clouds.

VMware Telco Cloud Automation implements the solution's high-level architecture through logical building blocks and core components, as shown in the following figure:

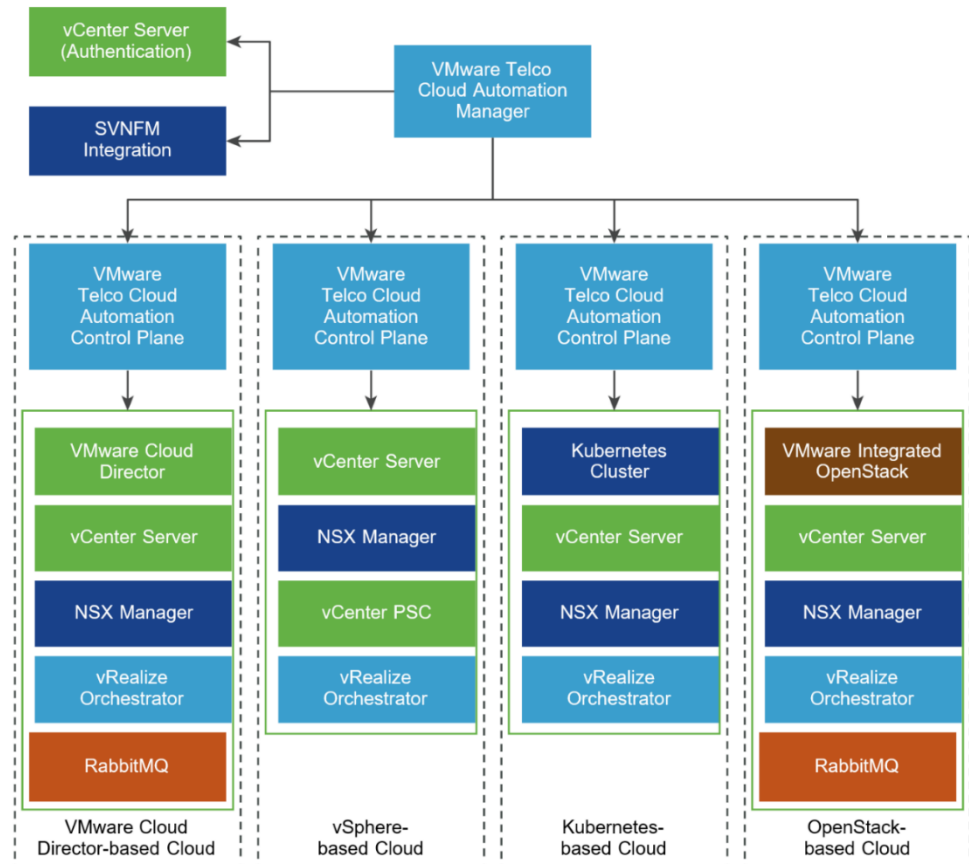


Figure 17. VMware Telco Cloud Automation architecture

Network services

The network services that are shown in Figure 17 are:

- **vCenter Server:** Used for authenticating and signing into VMware Telco Cloud Automation.
- **SOL 003 SVNFM:** Any SOL 003 SVNFM that can be registered with VMware Telco Cloud Automation.
- **VMware Telco Cloud Automation Control Plane (TCA-CP):** Deployed on the VIM and paired with VMware Telco Cloud Automation Manager.
- **VMware Telco Cloud Automation Manager:** Connects with TCA-CP to communicate with the VIMs. The VIMs are cloud platforms such as vCloud Director, vSphere, Kubernetes cluster, or VMware Integrated OpenStack.
- **vRealize Orchestrator:** Registered with TCA-CP and used to run NFV workflows. Customers can register for each VIM or for the entire network of VIMs.

The following figure shows the Telco Cloud Automation integration points:

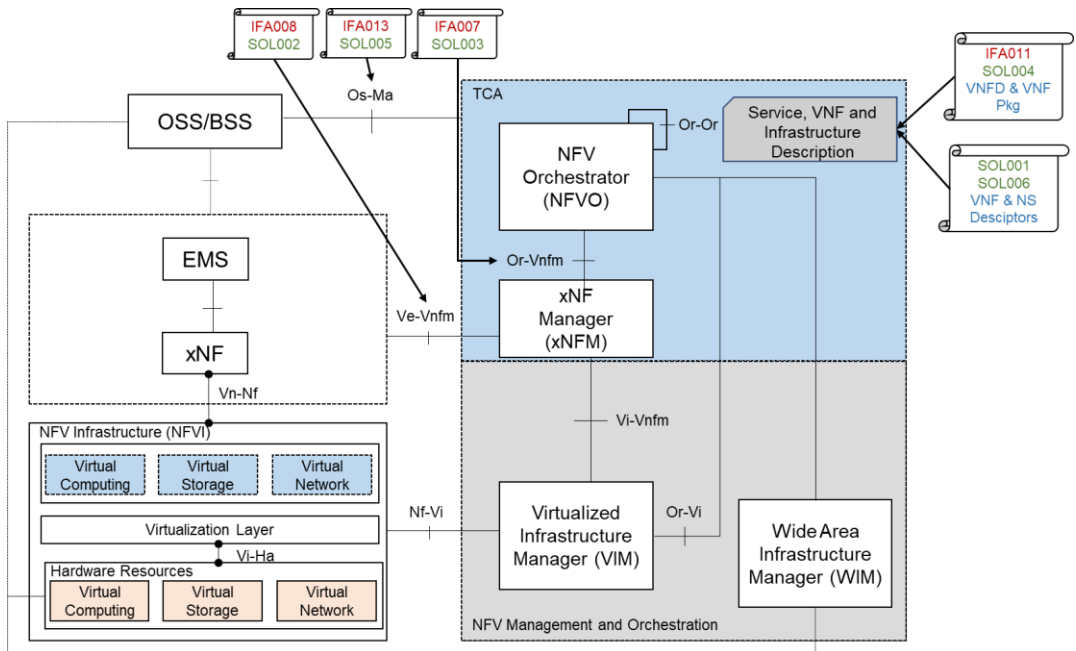


Figure 18. Telco Cloud Automation interoperability

Interoperability points

- SOL 001 ([ETSI GS NFV SOL 001](#)) refers to VNF Descriptor (VNFD), a recipe from the VNF provider that describes all the components of its VNF and how they are interconnected. SOL 001 covers VNFD and Network Service Descriptor (NSD) in accordance with [ETSI GS NFV-IFA 011](#). All NFs designed with Telco Cloud Automation are compliant with SOL 001 version 2.5.1.
- VNFD is packaged with manifest files, software images, and additional files in a VNF package known as CSAR (Cloud Service Archive), a zip file in a format that conforms with YAML specification version 1.2 in accordance with ETSI GS NFV SOL 004.
- Other interfaces include ETSI GS NFV SOL 003 for integration of external generic NFV orchestrators, ETSI NFV IFA008/SOL002 for third-party EMS integration, Or-VNfm ETSI NFV IFA007/SOL003 for VNF Manager (VNFM) integration, and Os-Ma ETSI NFV IFA013/SOL005 for OSS/BSS integration. For more information, see [Network Functions Virtualization](#) from ETSI.

Sample workflow for infrastructure and CaaS automation

Infrastructure automation enables operators to deploy virtualized infrastructure and provision new sites easily through automation and templates, reducing the need for complex manual site operations. The sample workflow for an automated site consists of the following steps:

1. Before starting the deployment, validate the prerequisites for each site: the host configuration, the physical switch and domain configurations, and the network design, including time synchronization and license availability. Infrastructure automation validates these prerequisites to ensure easy and fast deployment.

2. Install ESXi on each server and configure appropriate networking. Confirm IP reachability from the site to the data.
3. Bootstrap the Telco Cloud Automation virtual appliance to enable automated configurations from the UI.
4. Create site blueprints in the Telco Cloud Automation UI, including blueprints for:
 - Domains
 - DNS
 - NTP
 - Proxy servers
 - IP address ranges
 - Gateways
 - NIC configurations
 - Images
 - Licenses
 - VLANs
 - MTU size settings for jumbo frames,
 - Physical site location coordinates
5. Add hosts to each site.
6. Bootstrap the VMware telco cloud infrastructure (the software-defined data center cloud or SDDC).
7. Launch the Telco Cloud Automation control plane in each cloud location. The Telco Cloud Automation appliance in the central location provides its services at each cloud location that is using the VMware TCA-CP.
8. Deploy the Tanzu Kubernetes clusters to enable the management and worker cluster.
9. Instantiate the cloud-native network functions and the related network services. Customize the Tanzu worker clusters in accordance with the CSAR files for each workload.
10. Integrate the instantiated workloads with the appropriate OSS and BSS system to enable further LCM operations.

Deployment architecture

This section describes the high-level physical and virtual infrastructure, networking, and storage elements in the 5G Core reference architecture.

VMware Telco Cloud Platform for Core is a cloud-native core solution that is designed for running core functions. It provides the core modernization path, evolving from legacy core to virtualized core. It transforms the core into a 5G multiservice hub minicloud, enabling CSPs to monetize their core investments. The platform is designed to meet the performance and latency requirements that are inherent to core workloads.

The following figure shows the Telco Cloud Automation infrastructure domains:

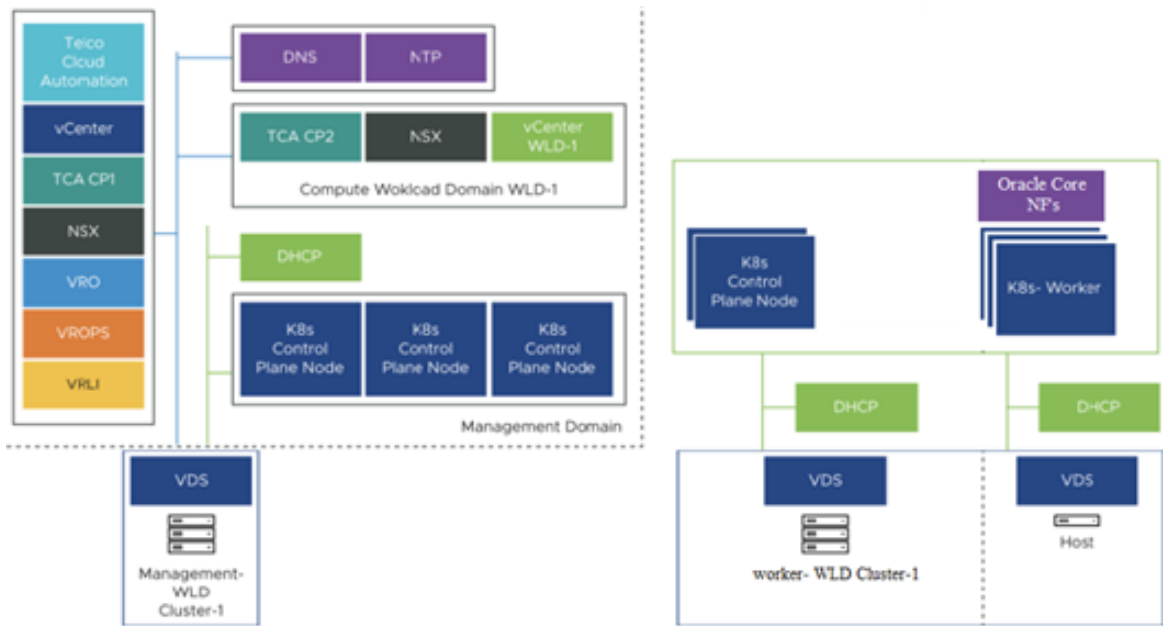


Figure 19. Telco Cloud Automation infrastructure domains

Features and benefits

The VMware Telco Cloud Platform for Core architecture:

- Enables CSPs to run virtualized core control-plane and user-plane functionality.
- Simplifies CSP operations consistently across distributed sites with centralized cloud-first automation, while reducing operating expenses.
- Provides operational consistency by removing business uncertainties and reduces the ballooning costs that are associated with 5G deployment.
- Enables CSPs to accelerate innovation speed, deploy 5G services quickly, and scale the services as customer demands increase.

VMware Telco Cloud Platform for Core is powered by field-proven virtualization, carrier-grade Container-as-a-Service (CaaS), and multilayer automation that is consistent with its 5G Core and edge offerings. The end-to-end consistency that is achieved by the coherent underpinning platform across 5G networks enables CSPs to provide 5G services that are customized for different enterprise and consumer markets while providing unparalleled operational efficiency.

The unique Telco Cloud Platform features that were created to deploy 5G Core NFs efficiently are:

- Telco Cloud Automation running on Data Center (DC) and providing orchestration and management services for telco clouds for Core.
- A dedicated DHCP server which is available to support the DHCP service offering for Kubernetes clusters.

- DC consisting of four R750 servers that primarily host and manage the 5G Core network services.
- Shared services such as Harbor and MetallLB running on the management cluster.
- Dedicated vCenter that is deployed on the DC cluster to manage the SDDC management and operational management components: vCenter Server, NSX, vRealize Operations, vRealize Log Insight, and so on.

Services design

Various external services are required for the initial deployment of the Telco Cloud Platform Core architecture components and the Tanzu Kubernetes grid cluster.

VMware Telco Cloud Platform 5G Core service requirements

The required external services and dependencies are listed in the following table and described below:

Table 6. Telco Cloud Platform services requirements

Service	Purpose
DNS	Provides name resolution for the various components in the solution.
DHCP	Provides automated IP address allocation for the Tanzu Kubernetes grid cluster. Ensure that this service is available locally to each site.
NTP	Synchronizes time between the various telco core management components at the central data center or RDC.

DNS

During Telco Cloud Platform core deployment, you must provide the DNS domain information to configure the various components. DNS resolution must be available for all the components in the Telco Cloud Platform Core solution. Required DNS records include the servers, VMs, and virtual IPs used. Ensure that both forward and reverse DNS resolution is functional for each component before deploying Telco Cloud Platform Core management components or creating any workload domains.

DHCP

Telco Cloud Platform Virtual Core uses DHCP to automatically configure the Tanzu Kubernetes cluster with an IPv4 address at the data center. The scope of DHCP must be large enough to accommodate all initial and future Kubernetes workloads to be used in the Telco Cloud Platform Core solution.

Note: While deploying the Tanzu Kubernetes Grid control plane, dedicate a static IP address for the Kubernetes API endpoint.

NTP

All management components must be synchronized against a standard time using the NTP. Telco Cloud Platform Core components such as vCenter Single Sign-On (SSO) are sensitive to a time drift between distributed components. The synchronized time between

the various components also helps with troubleshooting efforts. Each ESXi node in the Telco Cloud Platform Core solution must be time-synchronized.

To meet the NTP source requirements, ensure that:

- The IP addresses of either one NTP server or three or more NTP servers are provided during the initial deployment. Do not use two NTP reference servers in any deployment.
- The NTP sources are reachable by all the components in the Telco Cloud Platform Core solution.
- Time skew is less than five minutes between NTP sources.

Physical design

Physical ESXi hosts

The physical specifications of the ESXi hosts allow for successful deployment and operation of the design. The physical design specifications of the ESXi host determine the characteristics of the ESXi hosts that you use to deploy the Telco Cloud Platform Core solution—for example, consistent PCI card slot placement, especially for network controllers, is essential for accurate physical to virtual I/O resources. By using identical configurations, you can balance the VM storage components across storage and compute resources.

Four PowerEdge R750 servers are used in this reference architecture to meet the stringent requirements of Core NFs. The following Oracle Core NFs were onboarded and validated in these servers: NRF, PCF, BSF, NSSF, and SCP.

Storage considerations for 5G Core

In Kubernetes, a volume is a directory on a disk that is accessible to the containers inside a pod. While Kubernetes supports many types of volumes, we recommend a vSAN storage type for the Telco Cloud Platform Core deployment. The vSphere cluster at the data center can have three or more ESXi hosts to support the vSAN storage.

Note: If vSAN storage is used in the data center, follow the vSAN storage policies.

vSAN storage policies

vSAN storage policies define storage requirements for your storage class. Cloud Native Persistent Storage or Volume (PV) inherits the performance and availability characteristics that are made available by the vSAN storage policy. These policies determine how the storage objects are provisioned and allocated in the data store to guarantee the required level of service. “Kubernetes Storage Class” is a way for Kubernetes administrators to describe the “classes” of storage that are available for a Tanzu Kubernetes cluster. Different storage classes map to different vSAN storage policies.

Tanzu Kubernetes cluster design

Resource pools The Tanzu Kubernetes clusters are deployed in the compute workload domains. Telco Cloud Platform Core consumes resources from the compute workload domain, and resource pools provide guaranteed resource availability to workloads. Administrators can:

- Add more resources as capacity grows.
- Map each Kubernetes cluster to a resource pool.
- Dedicate a resource pool to a Kubernetes cluster or share the pool across multiple clusters.

In a core deployment design, a Kubernetes control-plane node and worker nodes can be placed on a vSphere cluster to support CNF workloads. The vSphere cluster can be managed by a vCenter in the compute workload domain.

Tanzu Basic for 5G Core deployment model

This section describes the Tanzu Basic for 5G Core deployment architecture and the placement of its components in the Telco Cloud Platform 5G Core design.

Tanzu Kubernetes management cluster

Tanzu Kubernetes management cluster functions as the primary management and operational center for the Tanzu Basic for Core instance. In this management cluster, the cluster API runs to create Tanzu Kubernetes clusters, and you configure the shared and in-cluster services that the clusters use. Shared and in-cluster services are services that run in the Tanzu Kubernetes grid instance to provide authentication and authorization of Tanzu Kubernetes clusters, logging, and ingress control.

Note: The management cluster is designed to operate the platform and manage the life cycle of Tanzu Kubernetes clusters. Do not use the management cluster as a general-purpose compute environment for end-user workloads.

Tanzu Kubernetes workload cluster

The Tanzu Kubernetes workload cluster is deployed from the Tanzu Kubernetes management cluster. Tanzu Kubernetes clusters can run different versions of Kubernetes, depending on the Cloud Native Network Functions (CNF) workload requirements. Tanzu Kubernetes clusters support multiple types of cloud-native infrastructures (CNIs) for pod-to-pod networking, with Antrea as the default CNI and the vSphere CSI provider for storage by default. When deployed through Telco Cloud Automation, VMware NodeConfig Operator is bundled into every workload cluster to handle the node operating system configuration, performance tuning, and operating system upgrades that are required for various telco CNF workload types for 5G Core. Other networks such as macvlans can be defined in cluster templates and will be mapped to actual networks when the cluster is deployed. Single Root I/O Virtualization (SR-IOV) networks are attached to pods while the NF is instantiated.

The following figure depicts the Tanzu Kubernetes management cluster:

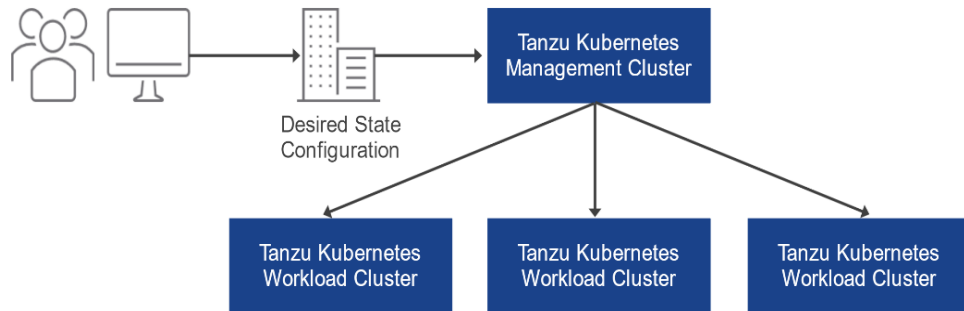


Figure 20. Tanzu Kubernetes management cluster

CNF design

This section describes the CNF requirements and how CNFs can be onboarded and instantiated in Telco Cloud Platform for 5G Core.

Helm charts

Helm is the default package manager in Kubernetes and is widely used by CNF vendors to simplify container packaging. With Helm charts, dependencies between CNFs are handled in formats that are agreed on by the upstream community, enabling telco operators to consume CNF packages in a declarative and easy-to-operate manner. With proper version management, Helm charts also simplify workload updates and inventory control.

The Helm repository is a required component in the Telco Cloud Platform 5G Core. Production CNF Helm charts must be stored centrally and be accessible by the Tanzu Kubernetes clusters. To reduce the number of management end points, the Helm repository must work seamlessly with container images. A container registry must be capable of supporting both container image and Helm charts. You can use Telco Cloud Automation to change the values being used by Helm charts. For more information, see the [VMware Telco Cloud Automation User Guide](#).

CSAR design

NF Helm charts are uploaded as a catalog offering wrapped around the ETSI-compliant TOSCA YAML (CSAR) descriptor file. The descriptor file includes the structure and composition of the NF and supporting artifacts such as the Helm charts version, provider, and set of preinstantiation jobs. Core NFs have sets of prerequisite configurations on the underlying Kubernetes cluster. Those requirements are also defined in the NF CSAR.

The CSAR extension supports the following features:

- SR-IOV interface configuration and addition along with DPDK binding
- Nonuniform memory access (NUMA) alignment of vCPUs and virtual functions
- Latency sensitivity
- Custom operating system package installations

- Full [GRUB](#) configuration

You can update CSAR files to reflect changes in CNF requirements or deployment models. CNF developers can update the CSAR package directly in the Telco Cloud Automation designer or use an external continuous integration and continuous deployment (CI/CD) process to maintain and build newer versions of the CSAR package.

Kubernetes cluster deployment process

The following figure shows the Kubernetes cluster deployment process:

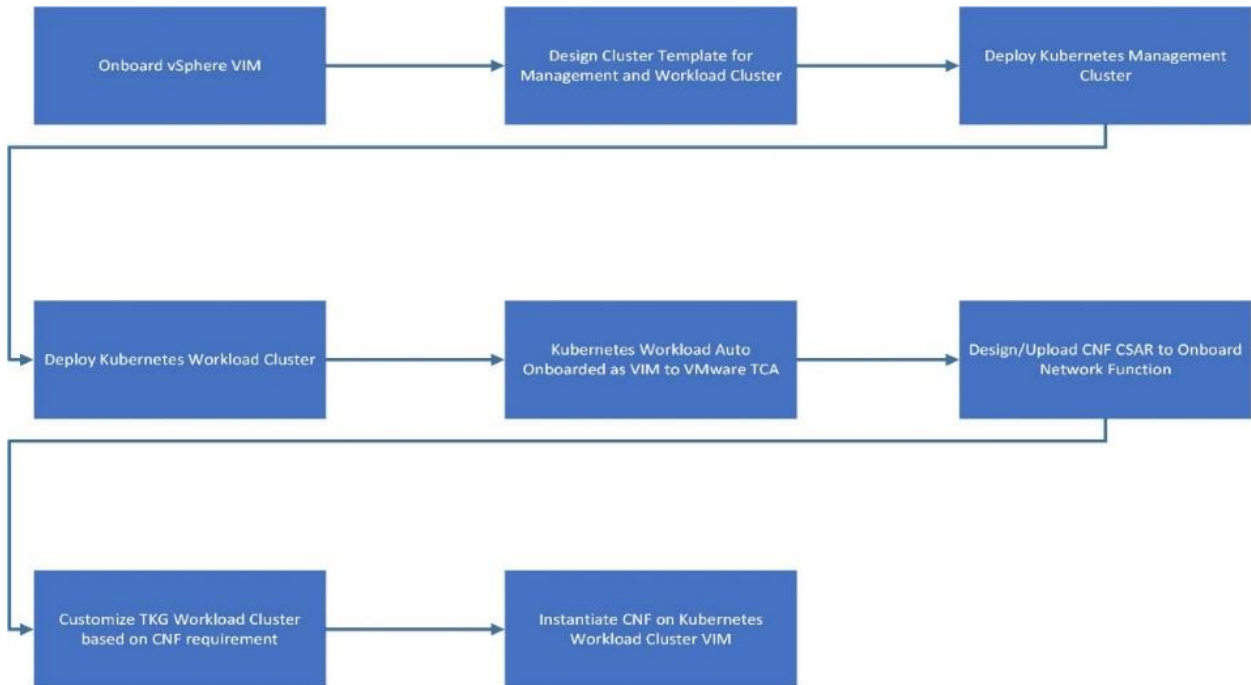


Figure 21. Kubernetes cluster deployment process

Telco workload deployment

VMware Telco Cloud Automation deploys the 5G Core workloads. At a high level, the deployment process consists of the following steps:

1. Make VMware Telco Cloud Automation aware of the infrastructure that can host CNF workloads by associating the VMware TCA-CP to a telco cloud infrastructure such as vSphere.
2. Create the Kubernetes cluster on the infrastructure by using VMware Telco Cloud Automation, which will be used to deploy the CNF workload.
3. Customize the control-plane node that manages the deployment of the workload based on the CNF requirements.
4. Modify the operating system on the Kubernetes control-plane node to make it compatible with the CNF deployment.
5. Install any add-ons that must be available on the control-plane nodes.

6. Customize the worker nodes on which the CNF is deployed based on the CNF requirements.
7. Deploy Platform as a Service (PaaS) components on the worker nodes that complement the CNF deployment.
8. Configure role-based authentication control policies for the Kubernetes cluster.
9. Deploy CNF on the worker nodes.
10. Push configuration information to the workloads.

For more information, see [Dell Technologies Reference Architecture for VMware Telco Cloud Platform 5G Edition 2.0 | Dell Technologies Info Hub](#).

Chapter 4 Oracle 5G Core Design

This chapter presents the following topics:

Oracle 5G Core	37
NF architecture	40
5G Core architecture	41
NF deployment on the cloud infrastructure	43
5GC NFs on VMware Telco Cloud Platform	45
Observability tools	48

Oracle 5G Core

Introduction

The wide range of use cases in 5G requires extra effort in making the 5G network more flexible, agile, scalable, and secure. The network core is where it all starts. The core of a telco network is where all the important decisions are made related to signaling, policy, charging, routing, and more. The fifth generation of mobile networks has radically changed the way elements in the core communicate with each other. The 5G Core control-plane functions follow a Service-Based Architecture (SBA), where network elements advertise and provide services that can be consumed by other elements in the core through application programming interfaces (APIs).

Oracle Communications NFs

Oracle Communications provides a full set of 5G Core NFs, with a special focus on:

- Routing and selection:** Establishing a robust, scalable, secure, and optimized service-aware routing and selection framework that includes services such as registration and discovery, automated resource control, network visibility, traffic management, core security authentication and authorization, slice selection, and binding support.
- Policy and charging:** Providing a unique, intuitive policy design and run-time experience to enable CSPs to quickly deploy new policies and services, while ensuring the reliability of existing services through a fully automated test framework. The solution is flexible enough to manage different domain-specific policies and granular enough to manage individual services.
- Network exposure:** Exposing network functions and 5G services in a secure and reliable way to both trusted and untrusted entities.

The following figure shows the Oracle 5G Core portfolio:

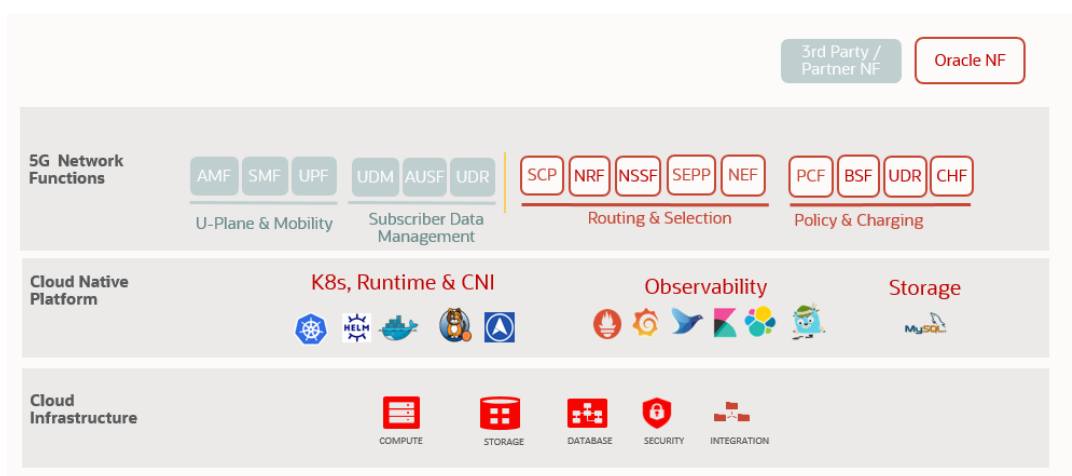


Figure 22. Oracle 5G Core portfolio

Oracle’s strategy is to provide core network signaling functions that leverage Oracle expertise and strength. The following components are responsible for:

- Signaling and routing:** NRF, SCP, NSSF, BSF, and SEPP
- Policy:** PCF, UDR

- **Charging:** CHF
- **API exposure:** NEF

The policy framework is one of the key components of the 5G network, empowering CSPs to better differentiate and customize tailored offerings for a wide range of consumer and enterprise use cases. At the core of the Oracle Communications solution policy is a flexible and intuitive policy design experience that uses new tools and architecture for design, run-time, debugging, and testing. This design enables CSPs to generate and test operator policies from scratch and deploy them into their production environment in minutes. The evolution to 5G requires a coherent and unified policy across the network that can easily influence edge computing routing and integration with network data analytics for smarter and more dynamic policy decisions. The solution must be agile and secure enough to manage different domain-specific policies for more differentiated and personalized customer offerings. It must also possess a policy framework that is granular enough to manage individual services and be capable of managing diverse services across network slices. Therefore, policy plays an even more important role in 5G than in any of the preceding generations of network technologies.

In addition to the need to generate new revenue streams, the top priorities for operators throughout the different mobile technology generations have been reducing network maintenance costs, handling traffic overload, managing signaling peaks, and securing the network from any kind of attack or breach. With the introduction of cloud-native technologies and the addition of another “G” to the spectrum, addressing these challenges is even more critical and complex in 5G. Oracle’s cloud-native 5G Core NFs enable CSPs to build a scalable, flexible, and secure signaling and routing framework as well as a routing and selection framework to manage the signaling complexity of 5G networks, ensuring that the network is robust enough to adapt to the evolving needs of your business.

The following section of this guide provides a brief description of the Oracle NFs. For more information, see the links in [Oracle documentation](#).

Policy Control Function

The PCF enables service providers to control the 5G network by implementing complex policies decisions based on network, subscriber, and service information. This information includes session management (SM) policies, access and mobility management (AM) policies, and policy authorization (PA) services. Use of these policies provides tailored offerings for a wide range of use cases—from Enhanced Mobile Broadband through Ultra-Reliable and Low Latency Communication to Massive IoT.

Unified Data Repository

The unified data repository (UDR) hosts the structured data for Unified Data Management (UDM), Oracle Communications' PCF, and NEF, as defined by 3GPP. The UDR architecture provides a highly available, distributed, and flexible data storage platform using service-based cloud-native design principles. The data repository is based on Oracle MySQL cluster technology, which supports both SQL and NoSQL database access and leverages Oracle expertise in delivering carrier-grade database solutions that are optimized for real-time critical requirements.

Binding Support Function

Different NFs in a 5G network coordinate with PCF to provide call, messaging, data, and other support services to users. To support all subscribers and user equipment (UE), multiple and separately addressable instances of the PCF are deployed in the network. Therefore, different NFs need binding support to identify the right PCF instance for session correlation and policy revalidation. BSF provides this binding support in the 5G Core network.

Network Repository Function

NRF works as a centralized repository for all the 5G Core NFs in an operator's network. The repository maintains updated records of the services that are provided by each of the elements in the 5G Core that can be instantiated, scaled, and terminated without manual intervention. This decoupling between the service provider and the service consumer increases the flexibility, scalability, and efficiency of the new 5G Core network, and helps operators manage their 5G network effectively by providing automated resource control in the core.

Service Communications Proxy

The new SBA brings unprecedented benefits for operators. However, this architecture is not fully equipped to deal with some of the major challenges that come with increased signaling traffic, such as:

- Routing and optimization
- Traffic management
- Robustness, scalability, and security
- Network visibility
- Core security—authorization and authentication

The Service Communications Proxy (SCP) addresses this problem by creating a secure 5G Core signaling architecture that provides routing control, resiliency, and observability into the 5G Core network. Oracle SCP is aligned with the latest 3GPP standards and can also be configured to work in prestandard mode to interwork with NFs that do not support Rel-16 headers.

Network Slice Selection Function

Along with increased speed, bandwidth, and latency, 5G and cloud-native technologies enable new business models and offer customized dedicated network slices for different services across various industries. This requires a high degree of deployment flexibility and efficient network resource utilization for operators so that they can help enterprises to efficiently launch innovative services in a cost-effective way. The network slice selection function (NSSF) provides network slice selection capabilities by enabling service providers to configure end-to-end dedicated logical network instances that are optimized for specific functional requirements of different applications and services.

Security Edge Protection Proxy

Security edge protection proxy (SEPP) is a nontransparent proxy sitting at the edge of the PLMN network and enabling secured inter-PLMN NF communication. Oracle Communications SEPP supports both standard and nonstandard security frameworks.

Network Exposure Function

Network exposure function (NEF) securely exposes the 5G network capabilities to Application Functions. Exposure functions act as a centralized point for service exposure and play a key role in authorizing all access requests originating from outside the 3GPP network to enable Cellular IoT, non-IoT, edge computing, and API gateway use cases for operators.

Charging Function

The Charging Function (CHF) is based on Oracle Communications Billing and Revenue Management (BRM). It acts as an Account Balance Management Function (ABMF) with the BRM Elastic Charging Engine (ECE) as a Converged Charging Server, a well-integrated and proven solution in mobile broadband deployments.

NF architecture

The Oracle Communications cloud-native 5G Core is built on principles that are defined by the Cloud Native Computing Foundation (CNCF), with seamless integration in open-source orchestration and automation frameworks and a range of popular cloud services that are sponsored by the CNCF. Oracle's 5G solutions are developed using DevOps principles and are designed to support zero-manual-touch management, in which the life cycle of NFs can be governed by automated CI/CD workflows.

All Oracle 5G core NFs are developed as cloud-native applications composed of a collection of microservices that run in a cloud-native environment and separate processing or business logic from state management. Microservices can be grouped into three layers:

- **Connectivity:** Components interfacing with external entities. An API gateway is used to interface with external traffic to the NF. These are stateless sets of components.
- **Business logic:** An application layer running the business logic, policy engine, and various services that can be enabled based on deployment requirements. These are stateless sets of components.
- **Data management:** Data layer that is responsible for storing various types of persistent data.

The following figure shows the Oracle CNF architecture:

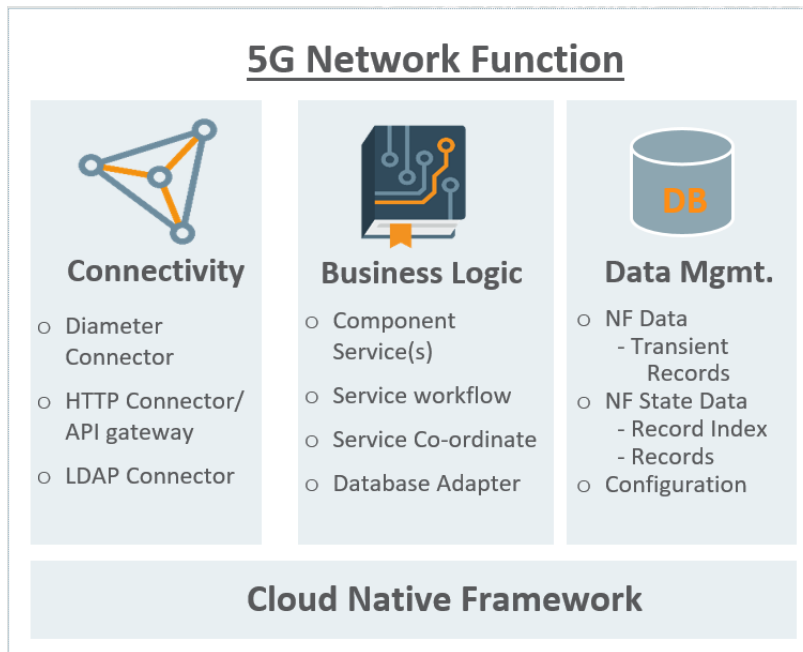


Figure 23. Oracle CNF architecture

Each layer includes multiple microservices, which are deployed as Kubernetes pods in replica sets. The data management layer is implemented as an independent cloud-native database tier (cnDB) based on MySQL Cluster Carrier Grade Edition (CGE). Each NF has a schema and a set of tables in the database tier. The same database tier can be used for multiple NFs.

Building a telecom network using authentic cloud-native solutions empowers CSPs to replicate the success of IT web-scale models, which have been proven to be agile, cost-effective, and customer-oriented. From an operational perspective, CSPs can benefit from increased automation and more efficiency gains through the adoption of agile processes. This includes continuous delivery of software, continuous deployment that is enabled by automated testing, methods to enable canary releases, end-to-end service orchestration, and, eventually, one-click procurement of slices that are deployed as a service for specific use cases and market segments.

5G Core architecture

The 5G Core solution that Dell Technologies, VMware, and Oracle have validated is intended to simplify the deployment of a 5G core network for CSPs, offering flexible design choices and reducing the time that it takes to design, test, and integrate components from multiple partners.

The architecture includes a basic set of 5G core NFs: PCF, BSF, NRF, NSSF, and SCP. These NFs provide the policy, signaling, and routing frameworks of the 5G SBA, which means that other network functions from Oracle and other vendors can be added to support a variety of additional 5G use cases.

The 5G core NFs are deployed on a common cloud infrastructure, including a Tanzu Kubernetes cluster. A common database tier NF provides the data management layer for

configuration and state data of all the NFs. Common tools are used to operate the solution.

The following figure shows the NFs that were used to validate the 5G Core solution:

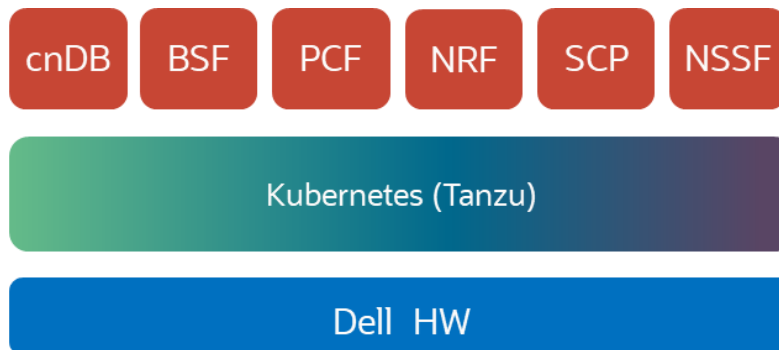


Figure 24. NFs selected for the initial validation of the 5G Core solution

Logical network architecture

In accordance with 3GPP standards, each NF exposes its functionality through a Service Based Interface (SBI), which employs a well-defined REST interface using HTTP/2. All the NFs expose their SBI interfaces on a common service network.

The NFs register their profile to the NRF, specify the services that are supported, and query the NRF to discover the NFs that are available for the service required. Also, an NF may subscribe to the NRF to be notified about status changes of other NFs.

A cURL-based tool is used to simulate other network functions that are not yet included in the solution, but are required in a complete 5G Core deployment such as AMF and SMF. The cURL simulator makes it possible to generate custom HTTP2 messages, simulating the request from an NF such as AMF or SMF according to standardized 3GPP procedures. The simulated AMF or SMF behaves as a consumer NF sending a request to a producer NF such as PCF, BSF, NSSF.

The following figure shows the logical network architecture of the 5G Core solution:

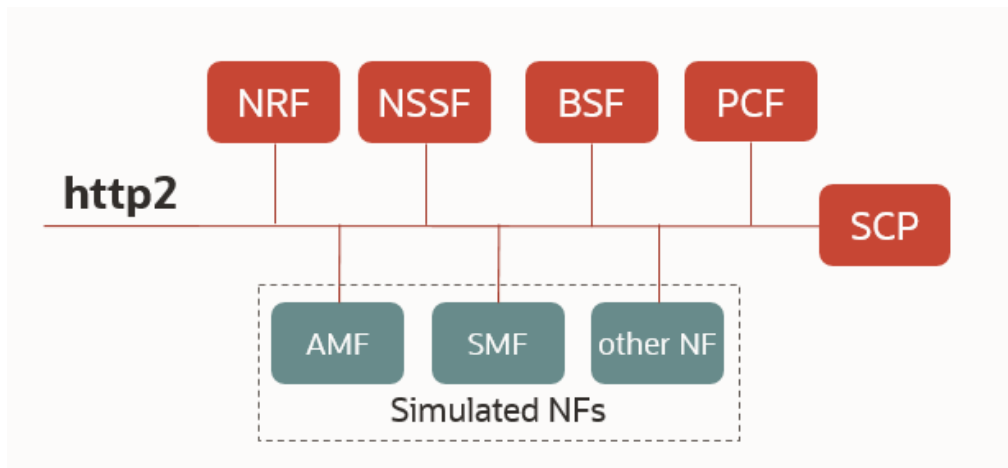


Figure 25. Logical network architecture of the 5G Core solution

The SCP may be used to route messages between consumer and producer NFs (also known as “indirect communications”), optimizing traffic routing with additional capabilities such as load-balancing and alternate routing. This requires the consumer NF to send a request to the SCP, which then routes the request to the target NF. The SCP supports indirect communications in accordance with 3GPP Release-16 Model C and D. In Model C, the consumer NF communicates directly with the NRF to discover the target producer NFs, and then uses SCP to route the requests; in Model D, the discovery of the target is offloaded to the SCP also. Model C has been tested with all the NFs.

NF deployment on the cloud infrastructure

The 5G NFs are cloud-native functions (CNFs) composed of multiple microservices that are deployed as pods in a Kubernetes cluster. Some microservices are common to many NFs:

- **Ingress and Egress Gateway:** Acting as a gateway for all ingress and egress http2 traffic with other NFs in the networks.
- **Configuration management service (the name may differ among NFs):** Providing the interface for NF configuration.
- **App-info and Perf-Info:** Monitoring the NF status and capacity.

Other microservices are specific to the business logic of each NF—for example, the NRF architecture includes dedicated microservices that manage procedures for registration, discovery, or subscriptions/notifications according to 3GPP standards. Similarly, the PCF architecture includes dedicated microservices that run the procedures for AM or SM Policy, connections to CHF or UDR, and so on.

The following figure shows the NF architecture with different microservices:

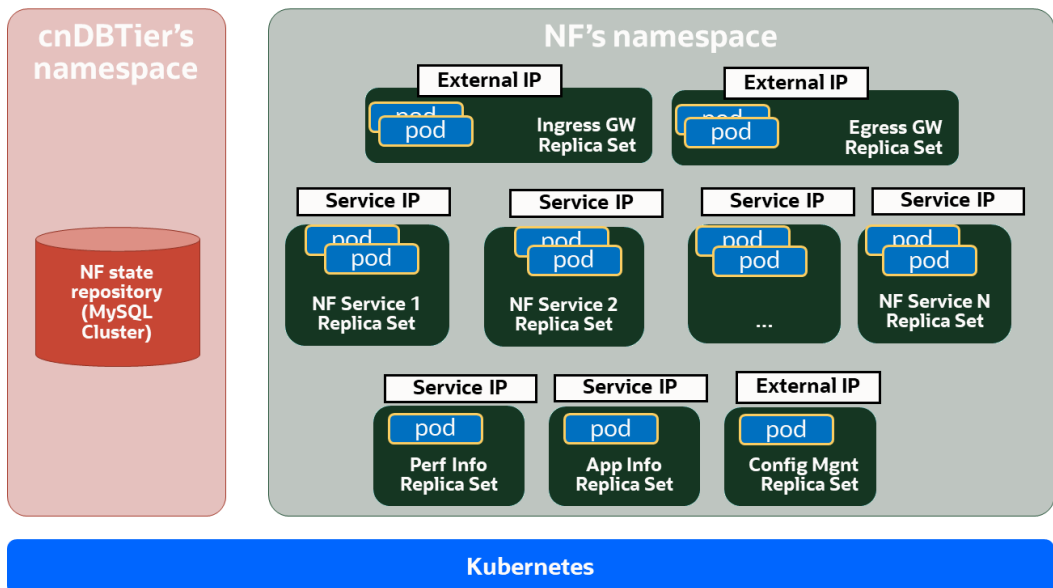


Figure 26. Generic NF architecture with different microservices

Each NF uses a dedicated namespace. Kubernetes pods implementing the different microservices of the NF are deployed in replica sets. The deployment settings for each NF

can be customized in a dedicated YAML file, `custom_value.yaml`, which lets you enable or disable specific features, set resources such as the minimum and maximum number of replicas in each set, and allocate cloud resources for execution (request and limit CPU and memory), `ServiceType` (clusterIP, LoadBalancer), and many others. API gateways and configuration management microservices get an external IP so that they can be accessed from other NFs or from management clients.

All NFs are stateless and store various types of persistent data on a common data management layer, the `cnDBTier`, which is deployed as a separated cloud-native function. The `cnDBTier` is built on MySQL Cluster Carrier Grade Edition and provides a reliable and scalable data layer which can replicate data locally and across different geographical sites.

The MySQL cluster includes a cluster manager that exposes management services for the database management API (deployment, registration, and so on), life cycle (creation, schemas, keys, and so on), and operations (provisioning, export, audit, and so on). MySQL Server is the API node that processes database transactions from the NFs. Data is stored in node groups of one to four data nodes with equivalent characteristics (compute, memory, storage) and is synchronously replicated across the data nodes in each node group. Replication nodes replicate updates to the node groups, regardless of the source.

The following figure shows the `cnDBTier` architecture:

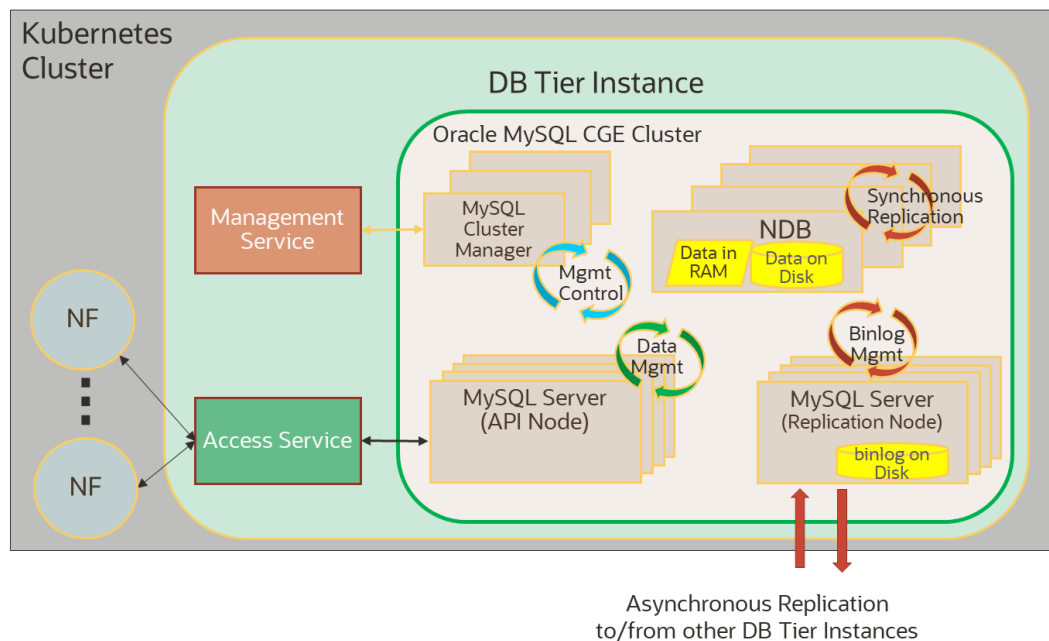


Figure 27. cnDBTier architecture

The `cnDBTier` is independent of the NFs using it. Each NF has a schema and a set of tables in the data tier. The `cnDBTier` is deployed first, so that other NFs can use it.

NF deployment and life cycle management

The following figure shows the general process for NF deployment:

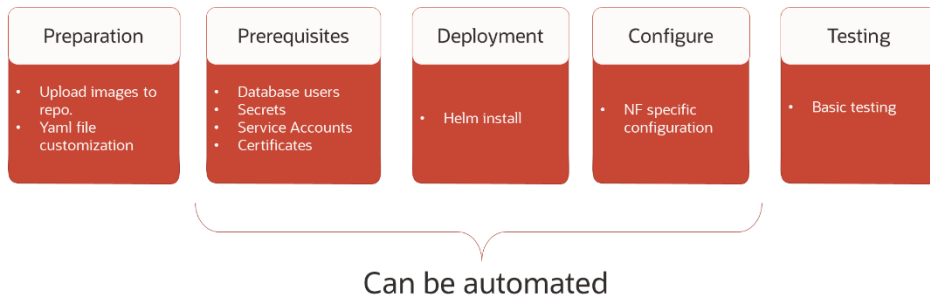


Figure 28. NF life cycle management

The software packages that are provided for each NF include software images, Helm charts, and YAML files. Before the deployment, at the preparation stage, some YAML files are customized for the specific environment. Preinstallation scripts are provided to perform several tasks, such as creating a dedicated namespace for each NF and creating database tables and secrets (the prerequisites stage). Then, NF microservices are deployed on the Kubernetes cluster using Helm. When the cluster is available, it can be configured using REST API and tested.

As shown in Figure 28, automation tools can be used to automate some of the LCM steps.

5GC NFs on VMware Telco Cloud Platform

The 5G Core solution is built on the cloud infrastructure that is described in [Chapter 3 Cluster Design](#). This includes VMware Telco Cloud Platform on Dell telco-grade infrastructure. The solution leverages the Telco Cloud Automation tool that is included in the Telco Cloud Platform suite for the LCM of Oracle’s cloud-native NFs. Other Oracle tools and Rest APIs are used for configuration management, while multiple open-source observability tools are used to collect and analyze log files, metrics, and trace files. Software images and Helm charts are uploaded on a common Harbor repository.

A CNF descriptor is created through the Telco Cloud Automation UI, which provides a wizard to set inputs such as preinstallation scripts to be run, the location of the Harbor repository for Helm charts, and binary files and custom YAML files to be used for NF instantiation.

At the end of the process, a CSAR archive is created for the NF. The archive is automatically onboarded in the Telco Cloud Automation catalog and made available to be instantiated on a Kubernetes cluster through the same Telco Cloud Automation interface, as shown in the following figure:

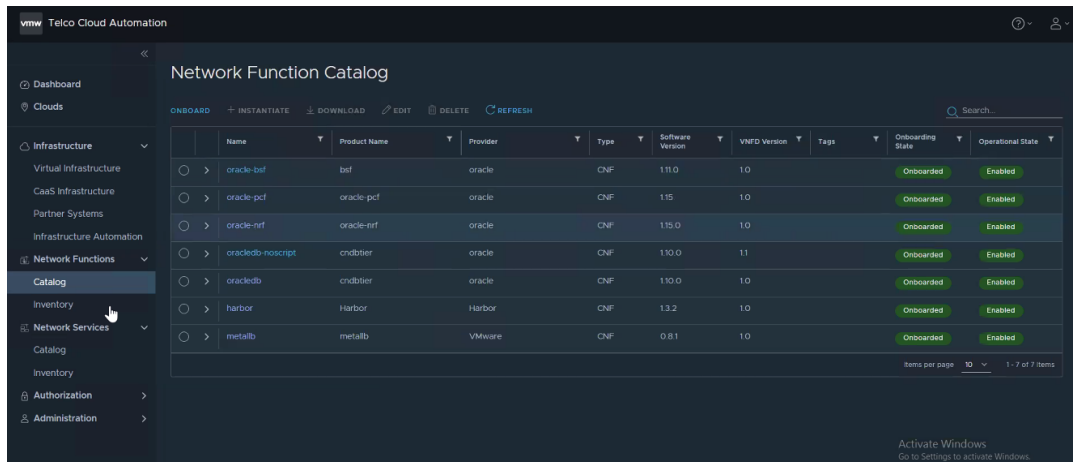


Figure 29. VMware Telco Cloud Automation NFs Catalog view

Configuration management

Oracle 5G NFs provide multiple options for configuration management. All NFs provide a REST API interface to configure the services and manageable objects according to a defined interface specification. Any REST client can be used for the purpose.

Configuration tasks in each NF are managed by a dedicated microservice, which exposes an external IP address through a load balancer. The PCF also provides a UI through the same IP, which allows the same NF configuration as is available through the REST API. This UI is normally used for configuration of the PCF, which requires a large number of options and possible settings compared to other NFs.

The following figure shows the configuration options:

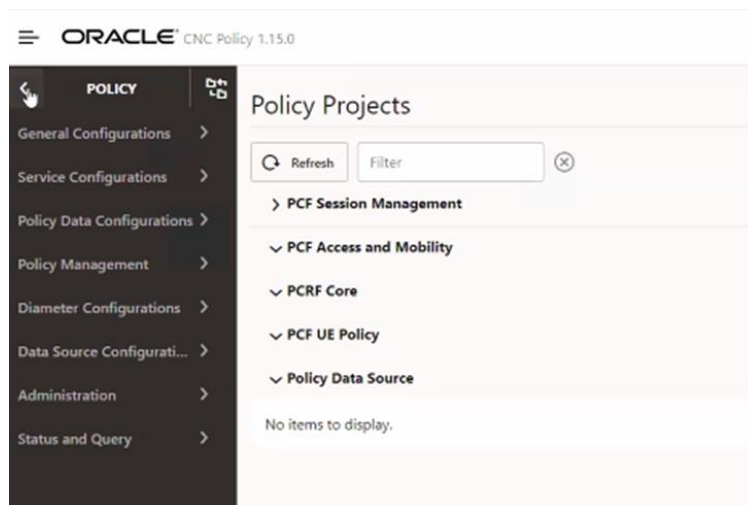


Figure 30. PCF UI for configuration

Use the UI to define all PCF global parameters with the same granularity as through the REST API.

From the same PCF console, you can access the Policy Design Studio application, which is used to design operator-defined policies. The Policy Design Studio enables users to build powerful logic using intuitive building blocks. Operators can use various building blocks to design logic such as conditions (for example, IF statements checking

parameters in an incoming request) or actions (for example, installing a PCC rule when the condition is met). Multiple blocks can be mixed to handle common use cases conveniently.

The following figure shows the design options that are available through the PCF UI:

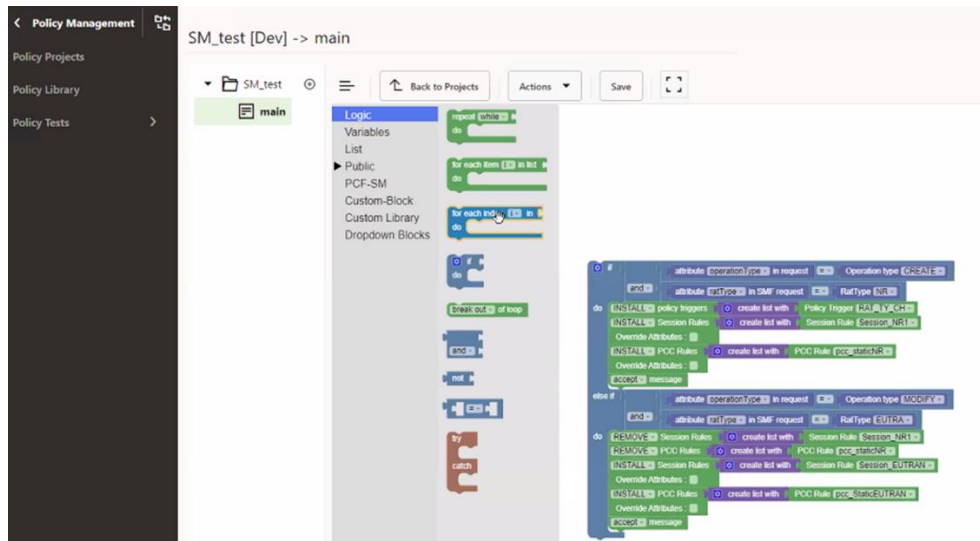


Figure 31. Policy design options in the PCF UI

Another option for Oracle NF configuration is to use the Cloud Native Core Console (CNCC), a single screen for configuring and managing multiple NFs. CNCC is deployed as an additional application on the same Kubernetes cluster. It includes a UI or API portal for NF configuration, and an Identity and Access Management (CNCC IAM) module. The IAM module acts as a local identity provider and broker for external identity providers and manages required authentication and authorization procedures such as creating and assigning roles to users.

The following figure shows the Oracle CNCC Welcome page:

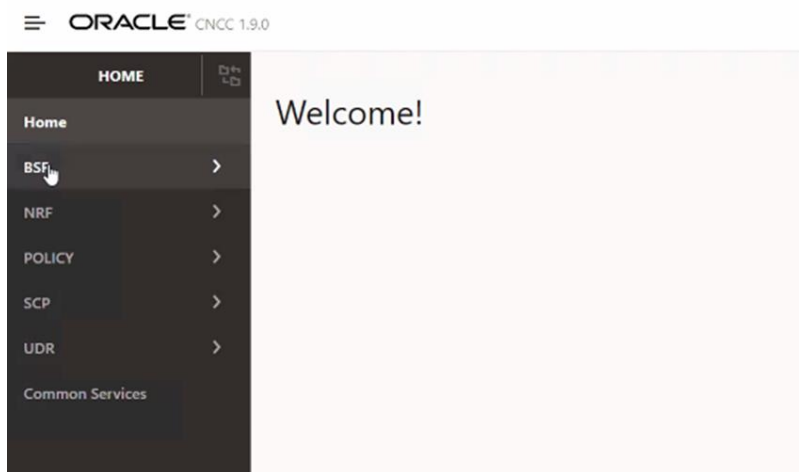


Figure 32. Welcome screen of Oracle’s Cloud Native Core console (CNCC)

Currently, CNCC supports configuration of PCF, UDR, BSF, NRF, SEPP, and SCP.

Observability tools

Open-source observability tools are used to collect and visualize metrics, log files, and trace files that are generated by 5G NFs during validation of the solution.

- **Prometheus** collects and stores the metrics that are generated by all 5G NFs.
- **Grafana** provides viewing and filtering options for metrics from the Prometheus database, as shown in the following figure:



Figure 33. Metrics viewed through the Grafana dashboard

- **Fluentd** collects logs from all the microservices and stores them in the ElasticSearch database.
- **Kibana** is used to filter and views the logs available in ElasticSearch, as shown in the following figure:

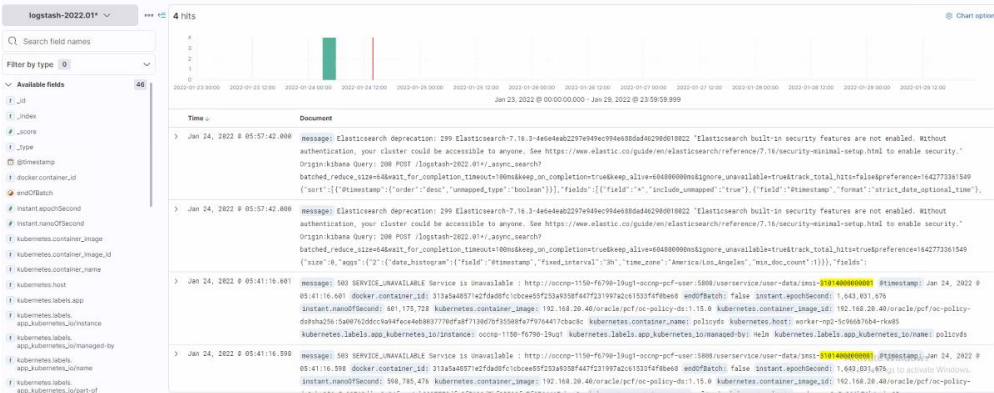


Figure 34. Example of log file visualization through the Kibana UI

- A **Jaeger** agent in each Kubernetes worker node collects trace files from all NFs and stores them in the ElasticSearch database, which is viewable through the Jaeger UI. You can filter the files for specific messages and interfaces.

The following figure shows an example of trace visualization through a Jaeger UI:

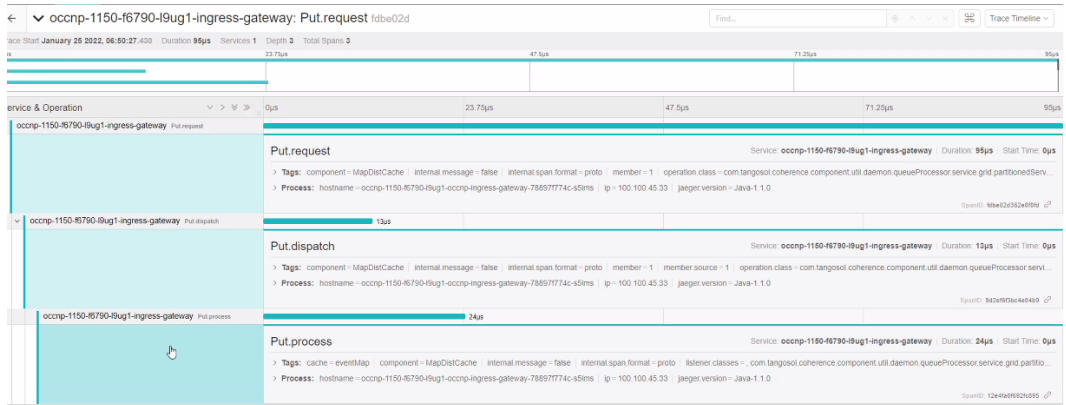


Figure 35. Example of Jaeger UI for trace visualization

The open-source observability tools that are used to monitor the solution run on the same Tanzu Kubernetes cluster as the NF. The latest software version that is available from public repositories is installed using Helm commands.

Note: In future releases, it will be possible to automate the installation procedures using Telco Cloud Automation.

The following figure shows the NF and operational tools that we used to validate the 5G Core solution:

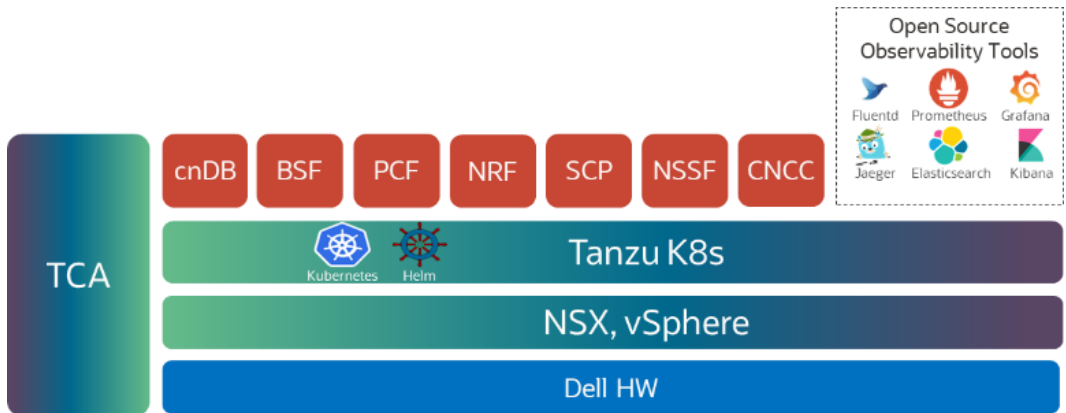


Figure 36. NFs and operational tools used for validation

Chapter 5 Solution Validation

This chapter presents the following topics:

Testing scope and setup	51
Use cases	51

Testing scope and setup

We tested the 5G Core solution in the laboratory to verify the deployment of Oracle 5G Core NFs on Dell and VMware cloud infrastructure and the correct functioning of the NFs in accordance with 3GPP procedures.

We used a cURL-based simulator to simulate other 5G NFs (mainly AMF and SMF), originating requests corresponding to different use cases.

The cURL simulator makes it possible to create HTTP2 command-line requests in which a method (GET, POST, HEAD, and so on), some request headers, and sometimes a request body are specified. The NF receiving the request responds with a status line indicating whether things went well, response headers, and possibly a response body.

Observability tools verify the correct functioning of the NF:

- Jaeger checks the response from each NF and validates the correct behavior according to 3GPP standard procedures.
- Prometheus and Grafana dashboard collect and visualize metrics from the NFs.
- Elasticsearch and Kibana collect and visualize log files that are generated by the NFs.

Note: Performance testing and HA testing are out of scope in the current 5G Core solution. They will be available in future releases of the solution.

Use cases

This section provides a description of the use cases that we used to validate the correct behavior of the NF after deployment.

NRF use cases

This section describes the objective, procedure, and expected results for NRF use cases.

Objective: Registration of an NF profile in NRF

The following figure shows the call flow:

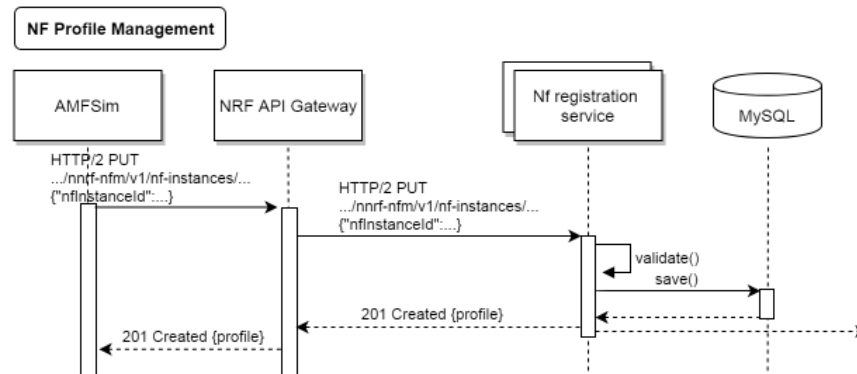


Figure 37. NF profile registration in NRF call flow

Procedure: Simulating an NF as SMF or AMF, an HTTP2 PUT request is sent to NRF with the NF profile to be registered.

Expected result: A 201 answer, indicating that registration is successful. Also, when registered profiles are checked in NRF using the REST API to display detailed profile information, the response returns the same information as was sent in the request.

NF profile discovery

Objective: Discovery of NF profiles registered by NFs in NRF.

The following figure shows the call flow:

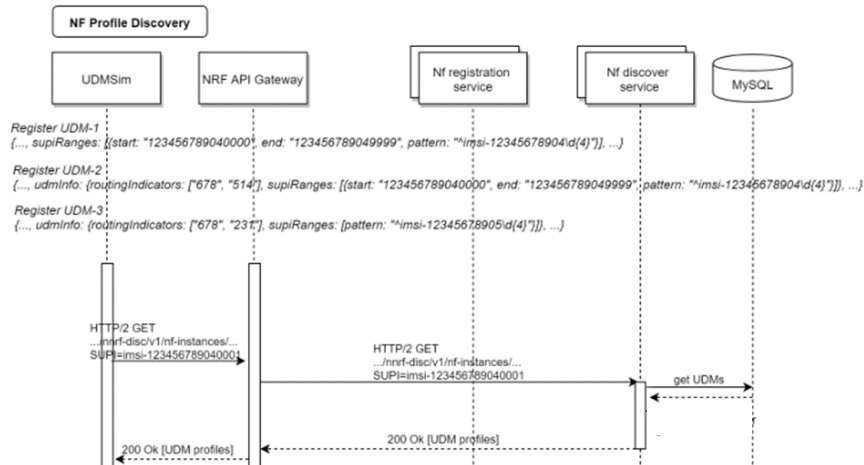


Figure 38. NF profiles in NRF discovery call flow

Procedure: One or more NF profiles is registered, including a real (nonsimulated) NF such as PCF, and an HTTP2 GET request requesting a registered NF profile is sent from a simulated NF to NRF.

Expected result: A 200 response, including the profile information of the NF that was registered in NRF.

PCF/BSF use cases

This section describes the objective, procedure, and expected results for PCF/BSF use cases.

SM session establishment

Objective: Create a session in PCF simulating an SMF. PCF will select a policy based on information in the request and the policy configuration in PCF.

The following figure shows the call flow:

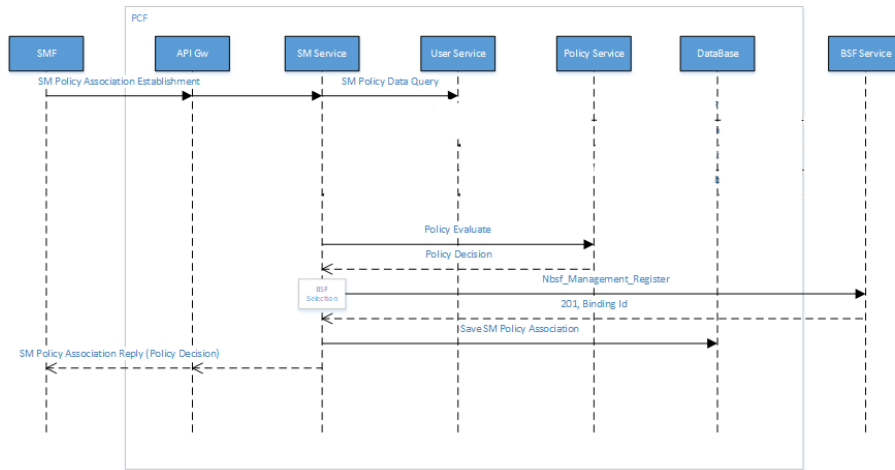


Figure 39. SM session establishment call flow

Procedure: A session establishment request is sent by the SMF simulator to PCF. Policies in PCF are selected based on the session information in the request and the conditions that are defined in the project. One rule is selected (it can be static or dynamic) and sent back to SMF.

Expected result: The selected rule in the response is based on information in the request and on conditions and actions that were configured in policy projects in PCF. The session should also be registered in BSF with PCF fqdn.

SM session update

Objective: Modify the SM session that you created with new parameters.

The following figure shows the call flow:

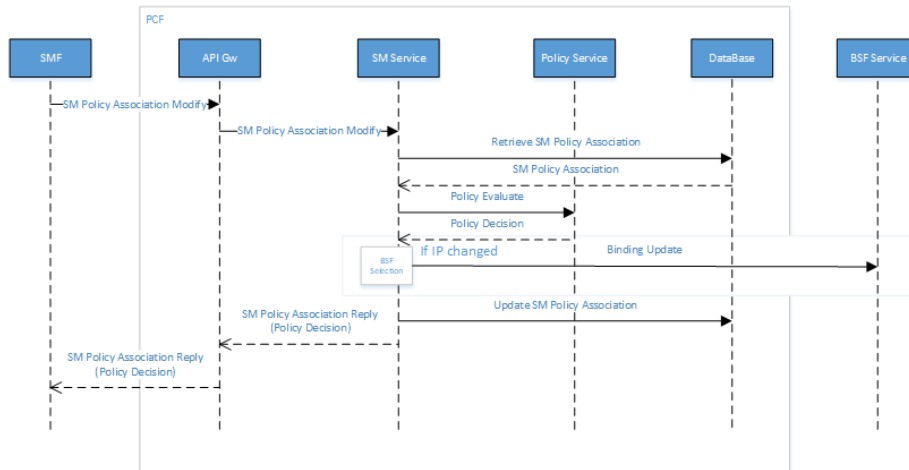


Figure 40. SM session update call flow

Procedure: An SM policy association modify request is sent by the SMF simulator to PCF with modified session information. SMF checks policies to select a new rule if applicable for the new information.

Expected result: Correct rule information is sent in response from PCF based on the information received and the conditions and action that are configured in the policy. Also, if there is an IP change, binding information is updated in BSF.

SM session termination

Objective: Terminate the SM session that you previously created and updated.

The following figure shows the call flow:

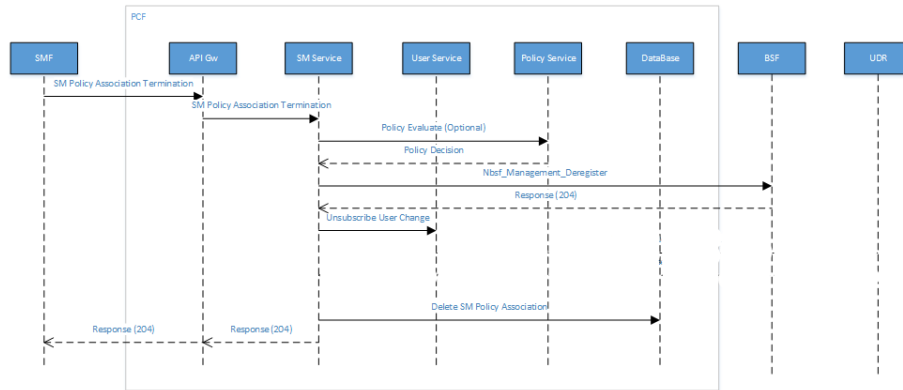


Figure 41. SM session termination call flow

Procedure: A session termination request with the session ID of the session that you previously created is sent from the SMF simulator to PCF.

Expected result: A 204 response from PCF. Confirm that the session was deleted in the UI of PCF and BSF.

AM session establishment

Objective: Create an AM session in PCF from the AMF simulator and select a rule based on information in the request.

The following figure shows the call flow:

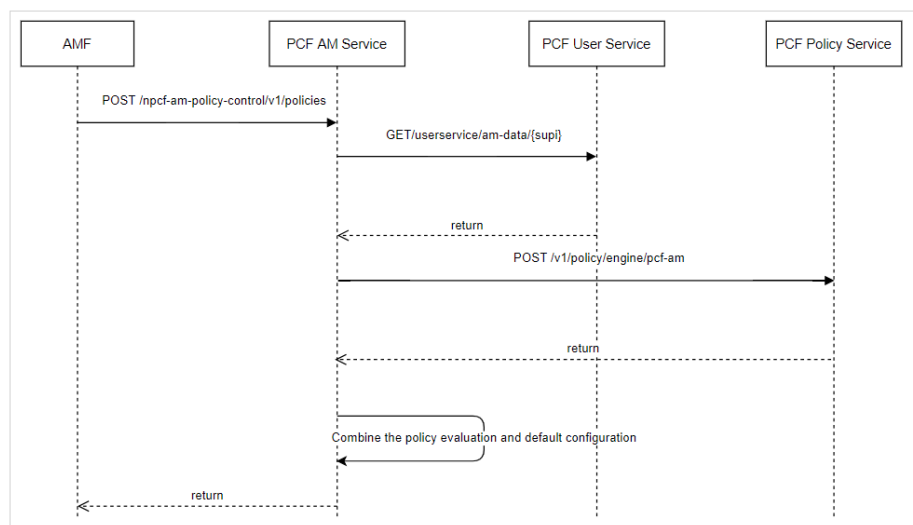


Figure 42. AM session establishment call flow

Procedure: A request to PCF including access network information is sent from the AMF simulator. The AM policy project has been configured so that PCF selects a rule according to the information received and returns the rule in the response message.

Expected result: The expected rule is received in the response message from PCF based on the access network information in the request. You can also verify the session creation in the PCF session viewer in the UI.

NSSF use cases

This section describes the objective, procedure, and expected results for NSSF use cases.

Not permitted S-NSSAI sent from UE

Objective: To test that the validation of the S-NSSAI that is requested by the UE was received from AMF.

The following figure shows the call flow:



Figure 43. S-NSSAI test call flow

Procedure: No authorized S-NSSAIs are included in the request from AMF (a simulated AMF is used in the first phase).

Expected result: Because no S-NSSAIs are allowed in accordance with the policy configuration, NSSF returns a “403 Forbidden” response.

Initial registration to the 5GS without AMF redirection

Objective: To simulate an initial registration of UE where the selected target AMF is the same as the initial AMF.

The following figure shows the call flow:

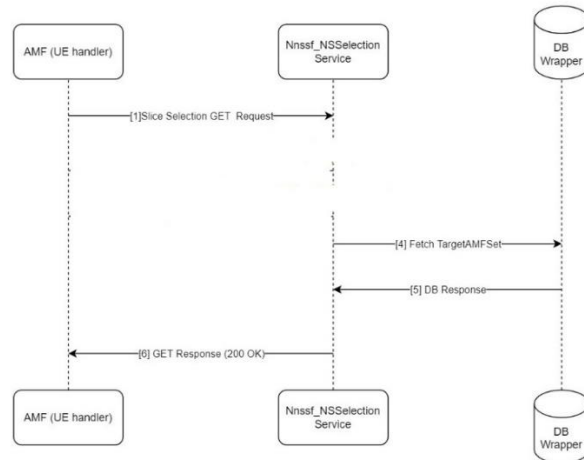


Figure 44. Initial UE registration

Procedure: An initial slice selection request is sent by the AMF simulator to NSSF. In accordance with the policy configured in NSSF, the target AMF selected is the same as the original AMF.

Expected result: A 200 OK response from NSSF containing the selected AMF information (here, the original AMF information).

Permitted S-NSSAI sent from UE

Objective: To test the NSSF feature to validate the S-NSSAI that is requested by UE and received from AMF.

The following figure shows the call flow:

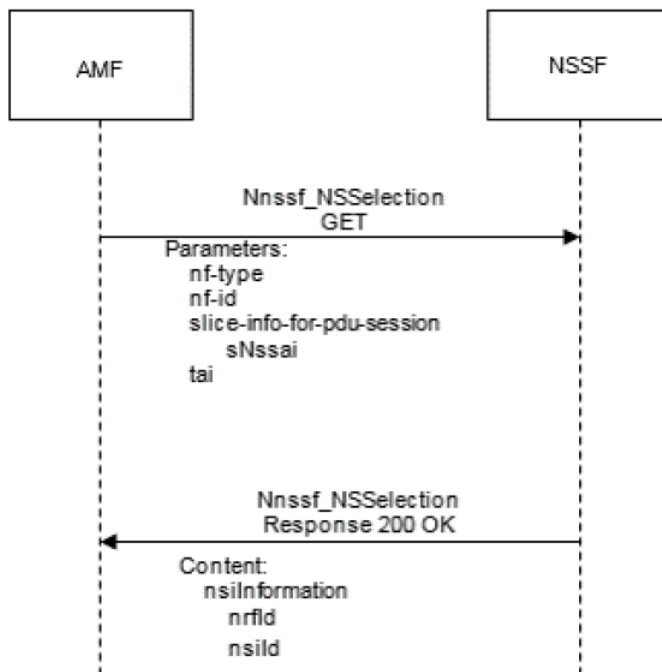


Figure 45. S-NSSAI from UE test call flow

Procedure: The AMF queries the NSSF with the specific S-NSSAI, the NF type of the NF service consumer, the requester ID, the PLMN ID of the SUPI, and the location information.

Expected result: The NSSF determines and returns the appropriate NRF to be used to select NF services within the selected network slice instance. The NSSF may also return an NSI ID identifying the network slice instance to use for this S-NSSAI.

Subscribe Service Operation

Objective: To test the NSSF availability service feature that validates the subscribe operation used by an NF service consumer (AMF) to get notifications for any change in NSSAI availability information.

The following figure shows the call flow:



Figure 46. Subscribe service test call flow

Procedure: AMF sends a to NSSF with a notification URL and a list of tracking area IDs (TAIs) as the JSON body.

Expected result: NSSF stores the subscription request and responds with the list of allowed S-NSSAIs per TAI in the request. NSSF also returns a subscription-id and an expiry date. The expiry date reflects the duration up to which NSSF sends notifications for any change in the grant status of S-NSSAIs for subscribed TAIs.

SCP use cases

This section describes the objective, the procedure, and the expected results for SCP use cases.

SCP indirect communication - 3GPP Rel-16 Model C

Objective: To route a request from a service consumer NF (such as a simulated SMF) to a service producer NF (such as a PCF) indirectly through SCP using the 3gpp-Sbi-Target-Apiroot custom http header.

The following figure shows the call flow:

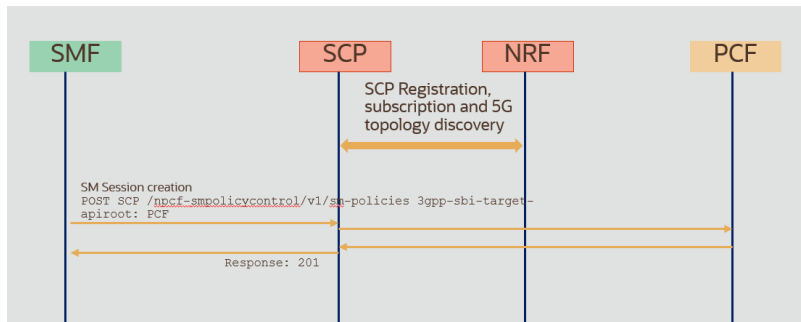


Figure 47. SCP indirect communication call flow

Procedure: After a successful registering of SCP and PCF to NRF, a request is sent from the SMF simulator to SCP that includes the `3gpp-sbi-target` API root that defines the target as PCF `fqdn`. Because SCP previously registered to NRF and discovered the network topology, it will use information in `3gpp-sbi-target` to route the request to PCF. PCF will send the response through SCP.

Expected result: Request and response are routed successfully through SCP. Confirm the successful routing by checking the response and the log and trace files in SCP and PCF.

Load balancing

Objective: Test the signaling load-balancing feature between two service producer NFs (such as PCF) originating from the service consumer (SMF simulator).

The following figure shows the call flow:

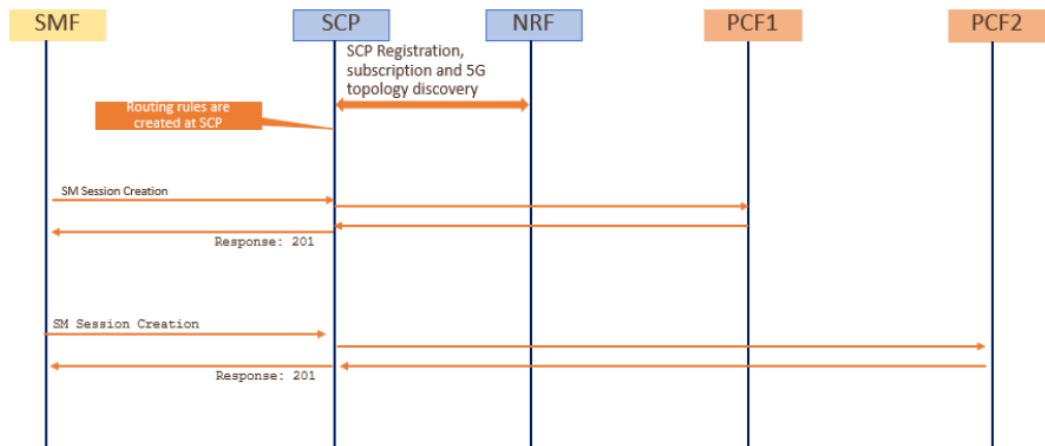


Figure 48. Load-balancing test call flow

Procedure: Signaling traffic is generated using the SMF simulator and sent to SCP. Routing rules that have been configured in SCP are applied, and requests that are received are forwarded to both PCFs according to the rules (percentage).

Expected result: Confirmation using log files or observability tools that traffic is balanced to both PCFs as defined in configured routing rules in SCP.

Alternate routing producer failure

Objective: To test the SCP feature of selecting an alternate route to a secondary producer in case a higher priority producer fails.

The following figure shows the call flow:

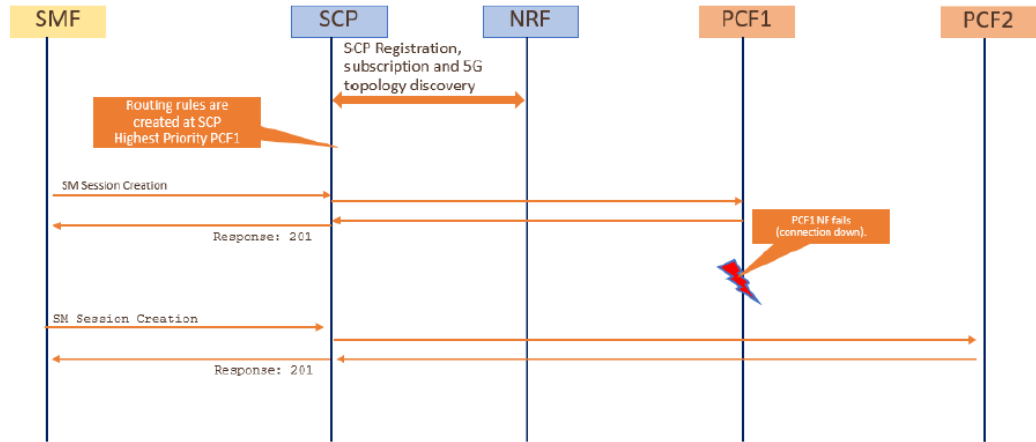


Figure 49. Alternate routing call flow

Procedure: When routing rules are configured in SCP, PCF1 has the highest priority. In normal conditions, all traffic is to be routed towards PCF1. A failure in PCF1 is simulated so that SCP selects the next lower priority NF producer of the same type (in this case, PCF2) to route traffic.

Expected result: If two PCFs are running, signaling traffic is routed to PCF1. After simulation of PCF1 failure, traffic is rerouted to PCF2. Verify the rerouting by using log files and observability tools.

Chapter 6 Foundation of CSP Grade 5G Core

This chapter presents the following topics:

Network deployment options	61
Security	62
Conclusions	63

Network deployment options

The NFs that we selected for the validated 5G Core solution provide a routing and policy framework to build a multivendor 5G network with the resiliency and scalability that is needed for 5G use cases.

Multivendor interoperability

This Dell Validated Design solution can integrate additional NFs from third-party vendors (AMF, SMF, UPF, AUSF, UDM) and from Oracle (SEPP, NEF) to provide a complete 5G Core solution and support new use cases.

Oracle 5G NFs provide standard interfaces that are aligned with [Rel-15](#) and [Rel-16](#) procedures. For our design validation, we deployed the NFs in live networks and integrated them with other NFs from all major vendors in the market.

We tested multiple use case and end-to-end procedures using NF simulators such as AMF or SMF to generate HTTP requests according to 3GPP standards.

Scalability and geo-redundancy

The solution can scale within a single site or in geo-redundant scenarios across two or more sites.

A single network site may include multiple instances of the same NF type, such as AMF, SMF, and PCF. The SCP can optimize the traffic in the network, ensuring load distribution and resiliency.

All Oracle NFs support geo-redundant deployments across two or more sites. NFs are made of stateless microservices that manage the SBI interfaces and the application logic of each NF, while session states are stored in a separated DBTier layer that is built on MySQL. The DBTier uses MySQL CGE to provide disaster recovery capabilities across two, three, or four sites.

The following figure shows an example of a two-site deployment with a geo-redundant DBTier for Oracle NFs:




Figure 50. Two-site deployment with geo-redundant DBTier for Oracle NFs

The Oracle NFs in each site use a common DBTier, which replicates data on the DBTier cluster in the other site. If an NF in one site (or even the whole site) becomes unavailable, the corresponding NF in the other site can take over the subscriber sessions of the failed NF because the session data is replicated on the DBTier in the other site. The NFs must be properly dimensioned to manage extra traffic in case the other site becomes unavailable.

Security

Because 5G touches every aspect of our lives with its broad set of use cases, the potential security threat is likely to increase. Operators must invest heavily to secure their 5G networks before they can consider use cases supporting business-critical and mission-critical industry-vertical applications such as healthcare and banking.

In the 5G SBA, all NFs can communicate with each other through an API that is based on HTTP and JSON. This communication capability requires a careful design to protect the data in the network (to avoid eavesdropping of the traffic) or to authorize communications among NFs (to avoid that a malicious actor registers a fake producer NF to the NRF and starts receiving traffic from other consumer NFs).

The solution supports standard security mechanisms that are defined by 3GPP:

- Token-based authorization to control the access of consumer NFs to the services offered by producer NFs.
- Mutual authentication and encryption of the communication between NFs or between NFs and SCP, based on TLS/HTTPs.

Token-based authorization with NRF

NRF supports 3GPP 29.510-based verification for access-token authorization requests for specific NF producers based on the allowed NF type and PLMN that are present in the NF profiles. An extension to this requirement is to include screening for access-token requests based on the NF type.

NRF performs the required authorization, and, if successful, will issue the token with the requested claims. The NRF provides an option to the user to tailor the authorization of the producer-consumer NF types along with the producer NF's services. The operator configures the mapping of the requester NF type, target NF type, and allowed services of the target NF. An access-token request is received based on the configuration and is further processed only if the authorization is successful. Allowed services can be configured as a single wildcard (*), which denotes that all the target NFs services are allowed for the consumer NF. The operator can also configure the HTTP status code and error description to be used in the error response that is sent by the NRF when the access token request is rejected.

Private keys are used by NRF to sign the access token that is generated. The private key is available only in the NRF. Public certificates are used by producer NFs to validate the access token that is generated by NRF. Accordingly, public certificates are available with producer NFs. NRF does not need the public certificate while signing the access token. The expiry time of the certificate is required to set appropriate validity time in the AccessTokenClaim.

Private keys and public certificates

Private keys are used by NRF to sign the access token that is generated. The private key is available only in the NRF.

Public certificates are used by producer NFs to validate the access token that the NRF generates. Accordingly, public certificates are available with producer NFs. NRF does not need the public certificate while signing the access token. The expiry time of the certificate is required to set an appropriate validity time in the AccessTokenClaim.

Two types of signing algorithms are supported by the NRF:

- **ES256:** ECDSA digital signature with SHA-256 hash algorithm
- **RS256:** RSA digital signature with SHA-256 hash algorithm

Support for HTTPS

HTTPS, an extension of HTTP, can be used to establish a secure connection between NFs in a 5G network, as defined in [3GPP TS 33.501](#). The HTTPS protocol uses Transport Layer Security (TLS) to encrypt the communication protocols between NFs, providing confidentiality and integrity protection to 5G Service Based Interface (SBI) messages.

To enable HTTPS on SBI messages, two NFs start the TLS handshake to agree an algorithm and keys to send messages securely to each other. After the handshake is completed, all communications between the client and the server are encrypted. This includes the full URL, data (plain text or binary), cookies, and other headers. The domain or host to which the client requested a connection is not encrypted because, when the connection is initiated, an HTTP request is made to the target server to create the secure connection. When HTTPS is established, the full URL is used.

This initialization must occur only once for each unique connection. In this respect, HTTP/2 has a distinct advantage over HTTP/1.1 because it multiplexes connections instead of opening multiple connections.

In indirect communications, the consumer and producer NFs communicate through the SCP, which accepts secured ingress connection requests from a consumer NF and establishes a secured egress connection with a producer NF.

Both HTTPS/TLS and Access Token/O Auth require a way to manage the keys and certificates that are required on all NFs. This is managed by a manual procedure in which each NF generates the private/public keys and can generate a Certificate Signature Request (CSR) to a Certificate Authority (CA). The CA can provide certificates to be installed on the database.

Conclusions

This guide describes a high-level reference architecture that has been developed by Dell Technologies and its partners Oracle and VMware. This guide can be used as the basis for future commercial deployments of 5G Core networks. The reference architecture provides a reusable model to help CSPs accelerate the adoption of this new technology with the advantages that are needed for next-generation 5G services. These advantages include:

- A 5G Core solution that provides technology and architectural flexibility to maximize network performance.
- Lowered Total Cost of Ownership through automation.
- A fully cloud-native solution providing simplification of deployment, automation of manual tasks, and network resiliency.
- A 5G Core platform that enables increased acceleration of new services to market and that can cost-effectively scale to meet the demand for those services.

This 5G Core reference model can be a blueprint for CSPs to deploy a 5G Core network that achieves these benefits.

Chapter 7 References

This chapter presents the following topics:

Dell Technologies documentation	66
Oracle documentation.....	66

Dell Technologies documentation

The following Dell Technologies documentation provides additional information about this solution. Access to these documents depends on your login credentials. If you do not have access to a document, contact your Dell Technologies representative.

- [Dell Technologies Reference Architecture for VMware Telco Cloud Platform 5G Edition 2.0](#)

Oracle documentation

The following Oracle web pages provide information about Oracle 5G Core, including solution briefs, datasheets, videos, and customer references.

- [Oracle: 5G Core Network Evolution](#)
- [Oracle: The Critical Role of Policy in the 5G Ecosystem](#)
- [Oracle: A more flexible and agile 5G core with Network Repository Function](#)
- [Oracle: A Secure, Cloud Native Signaling Solution for the New 5G Core](#)

The following Oracle web page provides links to product documentation such as user guides and installation guides:

- [My Oracle Support \(MOS\)](#)

Appendix A Hardware and Software Configuration

This appendix presents the following topics:

Overview **68**

Overview

The tables in this appendix show the key recommended parts for each node.

Note: When orders are placed, the Dell Technologies ordering center adds new SKUs and substitutes those SKUs that are shown in the tables with current local SKUs.

Hardware configuration

The following table shows the Bill of Materials (BOM) for the PowerEdge R650 server for the management cluster:

Table 7. PowerEdge R650 (4 units) BOM for the management cluster

Description	SKU	Quantity
PowerEdge R650 server	210-AYJZ	4
10 x 2.5 front storage	379-BEID	4
SAS/SATA/NVMe-capable backplane	379-BDSW	4
No rear storage	379-BDTE	4
No trusted platform module	461-AADZ	4
2.5-in. chassis with up to 10 hard drives (SAS/SATA) including a maximum of 4 universal drives, 3 PCIe slots, 2 CPU	321-BGHI	4
Intel Xeon Gold 6348 2.6G, 28C/56T, 11.2 GT/s, 42M cache, turbo, HT (235W) DDR4 3200	338-CBCI	4
Additional processor selected	379-BDCO	4
Heatsink for 2 CPU configuration (CPU more than 165W)	412-AAVM	4
Performance optimized	370-AAIP	4
3200 MT/s RDIMMs	370-AEVR	4
No RAID	780-BCDI	4
Dell HBA355i controller front	405-AAXY	4
Front PERC mechanical parts, rear load	750-ACFQ	4
Performance BIOS settings	384-BBBL	4
UEFI BIOS boot mode with GPT partition	800-BBDM	4

Description	SKU	Quantity
4 high-performance fans for 2 CPU (more than 165 W)	750-ADIH	4
Dual, hot plug, redundant power supply (1+1), 1400 W, mixed mode	450-AIQZ	4
Riser config 0, 2 CPU, half length, low profile, 3 x16 slots, SW GPU capable	330-BBRP	4
PowerEdge R650 motherboard	329-BFGQ	4
iDRAC9, Enterprise 15G	385-BBQV	4
Intel E810 XXV dual port 10/25 GbE SFP28, OCP NIC 3.0	540-BCXW	4
No bezel	350-BBBW	4
Luggage tray x8 and x10 chassis, R650	350-BCEI	4
BOSS blank	403-BCID	4
No quick sync	350-BBXM	4
iDRAC, legacy password	379-BCSG	4
iDRAC group manager, enabled	379-BCQV	4
No operating system	611-BBBF	4
No media required	605-BBFN	4
Redundant SD cards enabled	385-BBCF	4
16 GB micro-SDHC/SDXC card	385-BBOK	4
16 GB micro-SDHC/SDXC card	385-BBOK	4
ReadyRails sliding rails without cable management arm or strain relief bar	770-BECD	4
No systems documentation, no OpenManage DVD kit	631-AACK	4
PowerEdge R650 shipping	340-CUQR	4
R650 ship 4x3.5, 10x2.5, 8x2.5 NVMe	340-CUQN	4

Description	SKU	Quantity
R650 Dell/EMC label (BIS) for 2.5-in. chassis	343-BBQY	4
PowerEdge R650 no CCC marking, no CE marking	389-DYIC	4
IDSDM card reader	385-BBOV	4
Basic next business day 36 months	709-BBFM	4
ProSupport and next business day onsite service initial, 36 months	865-BBMY	4
Basic deployment Dell server R series 1U/2U	854-554	4
32 GB RDIMM, 3200 MT/s, dual rank 16 Gb base	370-AGDS	32
900 GB 15K RPM SAS 12 Gbps 512n 2.5 in hot plug hard drive	400-ASGV	24
960 GB SSD SATA mixed use 6 Gbps 2.5 in hot plug	400-AZVM	4
Jumper cord C13/C14, 0.6 M, 250 V, 13 A (North American, Guam, North Marianas, Philippines, Samoa)	492-BBDH	8
Intel E810 XXV dual-port 10/25 GbE SFP28 adapter, PCIe low profile	540-BCXV	4
Intel E810-XXV 25 GbE SFP28 2P OCP 3.0	540-BCXW	4

The following table shows the PowerEdge R750 BOM for the resource cluster:

Table 8. PowerEdge R750 (6 units) BOM for the resource cluster

Description	SKU	Quantity
PowerEdge R750 server	210-AYCG	6
2.5 chassis	379-BDTF	6
SAS/SATA backplane	379-BDSS	6
No rear storage	379-BDTE	6
No GPU enablement	379-BDSR	6
No trusted platform module	461-AADZ	6
2.5-in. chassis with up to 24 SAS/SATA Drives	321-BGFC	6
Intel Xeon Gold 6348 2.6G, 28C/56T, 11.2 GT/s, 42M Cache, Turbo, HT (235 W) DDR4 3200	338-CBCI	6
Additional processor selected	278-BDCO	6
Heatsink for 2 CPU with GPU configuration	412-AAVC	6
Performance optimized	370-AAIP	6
3200 MT/s RDIMMs	370-AEVR	6
No RAID	780-BCDI	6
Dell HBA355i controller front	405-AAXY	6
Front PERC mechanical parts, for 2.5-in. x24 SAS/SATA chassis	750-ADED	6
Performance BIOS settings	384-BBBL	6
UEFI BIOS Boot Mode with GPT partition	800-BBDM	6
Standard fan x6 V3	750-ADGK	6
Dual, hot plug, power supply redundant (1+1), 1400 W, mixed mode	450-AJHG	6
Riser config 2, full length, 4x16, 2x8 slots, DW GPU capable	330-BBRW	6
R750 motherboard	329-BFGT	6
iDRAC9, Enterprise 15G	385-BBQV	6
Intel E810 XXV dual port 10/25 GbE SFP28, OCP NIC 3.0	540-BCXW	6

Description	SKU	Quantity
No bezel	350-BBBW	6
Dell EMC luggage tag	350-BCED	6
Assembly BOSS blank	329-BERC	6
No quick sync	350-BBYX	6
iDRAC, legacy password	379-BCSG	6
iDRAC group manager, enabled	379-BCQV	6
No operating system	611-BBBF	6
No media required	605-BBFN	6
Redundant SD cards enabled	385-BBCF	6
16 GB micro-SDHC/SDXC card	385-BBOK	6
16 GB micro-SDHC/SDXC card	385-BBOK	6
ReadyRails sliding rails	770-BBBQ	6
No systems documentation, no OpenManage DVD kit	631-AACK	6
PowerEdge R750 shipping	340-CULS	6
PowerEdge R750 shipping material	481-BBFG	6
PE R750 No CCC or CE marking	389-DYHD	6
Dell/EMC label (BIS) for 2.5-in. chassis	389-DYHF	6
IDSDM card reader	385-BBOV	6
Basic next business day 36 months	709-BBFM	6
ProSupport and next business day onsite service initial, 36 months	865-BBMY	6
Basic deployment Dell Technologies server R Series 1U/2U	854-554	6
32 GB RDIMM, 3200 MT/s, dual rank 16 Gb BASE	370-AGDS	48
960 GB SSD SATA mixed use 6 Gbps 2.5 in Hot Plug	400-AZVM	24
900 GB 15 K RPM SAS 12 Gbps 512n 2.5 in hot plug hard drive	400-ASGV	48

Description	SKU	Quantity
Jumper cord C13/C14, 0.6 M, 250 V, 13A (North American, Guam, North Marianas, Philippines, Samoa)	492-BBDH	12
Intel E810 XXV dual port 10/25 GbE SFP28 adapter, PCIe full height	540-BCYK	6

The following table shows the PowerEdge R650 BOM for the edge cluster:

Table 9. PowerEdge R750 (4 units) BOM for the edge cluster

Description	SKU	Quantity
PowerEdge R750 server	210-AYCG	4
2.5 Chassis	379-BDTF	4
SAS/SATA backplane	379-BDSS	4
No rear storage	379-BDTE	4
No GPU enablement	379-BDSR	4
No trusted platform module	461-AADZ	4
2.5-in. chassis with up to 24 SAS/SATA drives	321-BGFC	4
Intel Xeon Gold 6348 2.6G, 28C/56T, 11.2 GT/s, 42M cache, turbo, HT (235 W) DDR4 3200	338-CBCI	4
Additional processor selected	278-BDCO	4
Heatsink for 2 CPU with GPU configuration	412-AAVC	4
Performance optimized	370-AAIP	4
3200 MT/s RDIMMs	370-AEVR	4
No RAID	780-BCDI	4
Dell HBA355i controller front	405-AAXY	4
Front PERC mechanical parts, for 2.5-in. x24 SAS/SATA chassis	750-ADED	4
Performance BIOS settings	384-BBBL	4
UEFI BIOS boot mode with GPT partition	800-BBDM	4
Standard fan x6 V3	750-ADGK	4

Description	SKU	Quantity
Dual, hot plug, power supply redundant (1+1), 1400 W, mixed mode	450-AJHG	4
Riser config 2, full length, 4x16, 2x8 slots, DW GPU capable	330-BBRW	4
R750 motherboard	329-BFGT	4
iDRAC9, Enterprise 15G	385-BBQV	4
Intel E810 XXV dual port 10/25 GbE SFP28, OCP NIC 3.0	540-BCXW	4
No bezel	350-BBBW	4
Dell EMC luggage tag	350-BCED	4
Assembly BOSS blank	329-BERC	4
No quick sync	350-BBYX	4
iDRAC, legacy password	379-BCSG	4
iDRAC group manager, enabled	379-BCQV	4
No operating system	611-BBBF	4
No media required	605-BBFN	4
Redundant SD cards enabled	385-BBCF	4
16 GB micro-SDHC/SDXC card	385-BBOK	4
ReadyRails sliding rails	770-BBBQ	4
No systems documentation, no OpenManage DVD kit	631-AACK	4
PowerEdge R750 Shipping	340-CULS	4
PowerEdge R750 shipping material	481-BBFG	4
PE R750 no CCC or CE marking	389-DYHD	4
Dell/EMC label (BIS) for 2.5-in. chassis	389-DYHF	4
IDSDM card reader	385-BBOV	4
Basic next business day 36 months	709-BBFM	4

Description	SKU	Quantity
ProSupport and next business day onsite service initial, 36 months	865-BBMY	4
Basic deployment Dell server R Series 1U/2U	854-554	4
32 GB RDIMM, 3200 MT/s, dual rank 16 Gb BASE	370-AGDS	32
960 GB SSD SATA mixed use 6 Gbps 2.5-in. hot plug	400-AZVM	16
900 GB 15 K RPM SAS 12 Gbps 512n 2.5-in. hot plug hard drive	400-ASGV	32
Jumper cord C13/C14, 0.6 M, 250 V, 13A (North American, Guam, North Marianas, Philippines, Samoa)	492-BBDH	8
Intel E810 XXV dual port 10/25 GbE SFP28 adapter, PCIe full height	540-BCYK	4

The following table shows the BOM for the PowerEdge R640 for the deployment cluster:

Table 10. PowerEdge R640 (1 unit) BOM for the deployment cluster

Description	SKU	Quantity
PowerEdge R640 server	210-AKWU	1
No trusted platform module	461-AADZ	1
PowerEdge R640 shipping	340-BKNE	1
DIMM blanks for system with 2 processors	370-ABWE	1
Standard 1U heatsink	412-AAIQ	1
Performance optimized	370-AAIP	1
No operating system	619-ABVR	1
No media required	421-5736	1
iDrac9, Enterprise	385-BBKT	1
OpenManage Enterprise server ConfigMgmt	528-BBWT	1
iDRAC group manager, enabled	379-BCQV	1
8 standard fans for R640	384-BBQJ	1

Description	SKU	Quantity
Dual, redundant, hot-plug PS, 750 W	450-ADWS	1
No systems docs, no OM DVD kit	631-AACK	1
US order	332-1286	1
iDRAC, legacy password	379-BCSG	1
No quick sync	350-BBKB	1
Slide RdyRL, no CMA	770-BBBC	1
ONSITE INSTL DECLINED	900-9997	1
900 GB, HDD 15 K SAS, 12 Gb, 512 n, 2.5, HP	400-ASGV	7
IDSDM and Combo card reader	385-BBLE	1
C13–C14, PDU, 12 A, 6.5 ft, 2 m, NA	492-BBDI	2
UEFI BIOS boot mode with GPT partition	800-BBDM	1
Riser config 2, 3x 16 LP	330-BBGN	1
Performance BIOS settings	384-BBBL	1
X550 QP 10 GbE, Base-T, rNDC	540-BBUY	1
Redundant SD cards enabled	385-BBCF	1
16 GB micro-SDHC/SDXC card	385-BBKG	1
No RAID	780-BCDI	1
HBA330 12 Gbps Cntrlr Mncrd	405-AAJU	1
960G SSD SAS, MU, 12, 2.5, HP, PX05SV	400-ASFI	1
2.5-in. chassis with up to 8 HD, 3 PCIe	321-BCQJ	1
PE R640 x8 Drive Shp Mtl	343-BBEV	1
No internal optical drive x4, x8 chassis	429-ABBF	1
No bezel	350-BBBW	1
Dell EMC luggage tag	350-BBJS	1
X710 DP, 10 Gb DA/SFP+, CvNwAd, LP	555-BCKN	1

Description	SKU	Quantity
Gold 6240 2.6 G, 24.75 M, 150 W	1338-BSGN	1
PowerEdge R640 MLK motherboard	329-BEIJ	1
Gold 6240 2.6 G, 24.75 M, 150 W	338-BSGN	1
Additional processor selected	379-BDCO	1
2933 MT/s RDIMMs	370-AEPP	1
32 GB RDIMM, 2933 MT/s, dual rank	370-AEQH	6

The following table shows the BOM for the Dell Networking S5232-ON switch for leaf switches:

Table 11. Dell Networking S5232-ON (2 units) BOM for leaf switches

Description	SKU	Quantity
OS10 Enterprise, S5232F-ON operating system version: 10.5.3.4.108	634-BRUO	2
Dell EMC S52XX-ON user guide	343-BBLP	2
US order	332-1286	2
S5232F-ON, PSU to I/O air, 2x PSU, OS10	210-APHN	2
HW WRTY-SVC NW S5232F-ON	818-5110	2
RTD PARTS NW S5232F-ON 1YR	818-5111	2
SW WARRANTY NW S5232F-ON 90D	818-5112	2
CBL, Q28–Q28, 100 G, 1 M direct attach	470-ABOV	2
Cbl, 100 GbE QSFP28 Pssv, 1 Meter	470-ABOR	18
Cbl, 100 GbE QSFP28 Pssv, 2 Meter (optional)	470-ABOS	18
Info 3rd Party O/S warranted by vendor	997-6306	6
ONSITE INSTL DECLINED	900-9997	6
250 V, 12 A, 2 MTR, C13/C14	450-AASX	12

The following table shows the BOM for the Dell Networking S3048-ON switch for management:

Table 12. Dell Networking S3048-ON (1 unit) switch for management BOM

Description	SKU	Quantity
Dell Networking S3048-ON operating system version: 10.5.3.4.108	210-AEDP	1
Operating system	528-BBSY	1
System documentation	634-BCXR	1
Power cords	450-AASX	1
Software features	634-BDXE	1
Hardware support services	802-7389, 802-7390, 802-7391, 995-9859, 997-6306	1
Deployment services	900-9997	1

Software configuration

The following tables show the software components of the solution:

Table 13. VMware Telco Cloud Platform Core 2.0 software components

Component	Version
Telco Cloud Platform	Core 2.0
ESXi version	Dell EMC customized VMware ESXi 7.0 U1 A01
VMware Tanzu Kubernetes Grid	1.2
vRealize Orchestrator Appliance	8.2
NSX-T	NSX-T Data Center Advance Edition 3.1
Telco Cloud Automation	1.8
vCenter	7.0 U1a

Table 14. Oracle CNC and 5G NF software components

Component	Version
Oracle Cloud Native Core (CNC)	2.5.0.0
NRF	1.15.0
PCF	1.15.0
BSF	1.11.0
NSSF	1.8.0
SCP	1.15.0