

Oracle® Communications

Cloud Native Core, Network Repository Function Benchmarking Guide



Release 25.2.201

G52536-01

April 2026

The Oracle logo, consisting of a solid red square with the word "ORACLE" in white, uppercase, sans-serif font centered within it.

ORACLE®

Copyright © 2023, 2026, Oracle and/or its affiliates.

This software and related documentation are provided under a license agreement containing restrictions on use and disclosure and are protected by intellectual property laws. Except as expressly permitted in your license agreement or allowed by law, you may not use, copy, reproduce, translate, broadcast, modify, license, transmit, distribute, exhibit, perform, publish, or display any part, in any form, or by any means. Reverse engineering, disassembly, or decompilation of this software, unless required by law for interoperability, is prohibited.

The information contained herein is subject to change without notice and is not warranted to be error-free. If you find any errors, please report them to us in writing.

If this is software, software documentation, data (as defined in the Federal Acquisition Regulation), or related documentation that is delivered to the U.S. Government or anyone licensing it on behalf of the U.S. Government, then the following notice is applicable:

U.S. GOVERNMENT END USERS: Oracle programs (including any operating system, integrated software, any programs embedded, installed, or activated on delivered hardware, and modifications of such programs) and Oracle computer documentation or other Oracle data delivered to or accessed by U.S. Government end users are "commercial computer software," "commercial computer software documentation," or "limited rights data" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, the use, reproduction, duplication, release, display, disclosure, modification, preparation of derivative works, and/or adaptation of i) Oracle programs (including any operating system, integrated software, any programs embedded, installed, or activated on delivered hardware, and modifications of such programs), ii) Oracle computer documentation and/or iii) other Oracle data, is subject to the rights and limitations specified in the license contained in the applicable contract. The terms governing the U.S. Government's use of Oracle cloud services are defined by the applicable contract for such services. No other rights are granted to the U.S. Government.

This software or hardware is developed for general use in a variety of information management applications. It is not developed or intended for use in any inherently dangerous applications, including applications that may create a risk of personal injury. If you use this software or hardware in dangerous applications, then you shall be responsible to take all appropriate fail-safe, backup, redundancy, and other measures to ensure its safe use. Oracle Corporation and its affiliates disclaim any liability for any damages caused by use of this software or hardware in dangerous applications.

Oracle®, Java, MySQL, and NetSuite are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

Intel and Intel Inside are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. AMD, Epyc, and the AMD logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark of The Open Group.

This software or hardware and documentation may provide access to or information about content, products, and services from third parties. Oracle Corporation and its affiliates are not responsible for and expressly disclaim all warranties of any kind with respect to third-party content, products, and services unless otherwise set forth in an applicable agreement between you and Oracle. Oracle Corporation and its affiliates will not be responsible for any loss, costs, or damages incurred due to your access to or use of third-party content, products, or services, except as set forth in an applicable agreement between you and Oracle.

Contents

1	Introduction	
1.1	Purpose and Scope	1
1.2	References	1
2	Deployment Environment	
2.1	Deployed Components	1
2.1.1	Hardware Details	1
2.1.2	System Software	1
2.1.3	Observability Services	1
2.1.4	CNE Common Services	2
3	Test Topology	
3.1	Single Set Test Topology	1
3.2	Multiset Test Topology (NRF Growth)	1
4	Benchmark Testing	
4.1	Overview	1
4.2	NRF Features	1
4.3	Software Test Constraints	4
4.4	Service Operations	5
4.5	53K TPS per NRF Set/106K TPS for Two NRF Sets	6
4.5.1	NRF Services Resource Requirements	6
4.5.2	cnDBTier Resource Requirements per NRF Deployment	8
4.5.3	State Data Quantum at NRF	11
4.5.3.1	NF Count Per NF Types	11
4.5.3.2	NF Subscription Count	12
4.5.4	Call-Mix	12
4.5.4.1	Discovery TPS per NfType and Number of Profiles in Response	16
4.5.5	Test Observations	16
4.5.6	NRF Resource Utilization	17
4.5.7	cnDBTier Services Resource Utilization	17

Preface

- [Documentation Accessibility](#)
- [Diversity and Inclusion](#)
- [Conventions](#)

Documentation Accessibility

For information about Oracle's commitment to accessibility, visit the Oracle Accessibility Program website at <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=docacc>.

Access to Oracle Support

Oracle customer access to and use of Oracle support services will be pursuant to the terms and conditions specified in their Oracle order for the applicable services.

Diversity and Inclusion

Oracle is fully committed to diversity and inclusion. Oracle respects and values having a diverse workforce that increases thought leadership and innovation. As part of our initiative to build a more inclusive culture that positively impacts our employees, customers, and partners, we are working to remove insensitive terms from our products and documentation. We are also mindful of the necessity to maintain compatibility with our customers' existing technologies and the need to ensure continuity of service as Oracle's offerings and industry standards evolve. Because of these technical constraints, our effort to remove insensitive terms is ongoing and will take time and external cooperation.

Conventions

The following text conventions are used in this document:

Convention	Meaning
boldface	Boldface type indicates graphical user interface elements associated with an action, or terms defined in text or the glossary.
<i>italic</i>	Italic type indicates book titles, emphasis, or placeholder variables for which you supply particular values.
monospace	Monospace type indicates commands within a paragraph, URLs, code in examples, text that appears on the screen, or text that you enter.

My Oracle Support

My Oracle Support (<https://support.oracle.com>) is your initial point of contact for all product support and training needs. A representative at Customer Access Support can assist you with My Oracle Support registration.

Call the Customer Access Support main number at 1-800-223-1711 (toll-free in the US), or call the Oracle Support hotline for your local country from the list at <http://www.oracle.com/us/support/contact/index.html>. When calling, make the selections in the sequence shown below on the Support telephone menu:

- For Technical issues such as creating a new Service Request (SR), select **1**.
- For Non-technical issues such as registration or assistance with My Oracle Support, select **2**.
- For Hardware, Networking and Solaris Operating System Support, select **3**.

You are connected to a live agent who can assist you with My Oracle Support registration and opening a support ticket.

My Oracle Support is available 24 hours a day, 7 days a week, 365 days a year.

Acronyms

The following table provides information about the acronyms and the terminology used in the document.

Table Acronyms

Term	Definition
ASM	Aspen Service Mesh
B	Bytes
CA	Certificate Authority
CNE	Cloud Native Environment
CPU	Central Processing Unit
KB	Kilo Bytes
LB	Load Balancer
NF	Network Function
NRF	Oracle Communications Cloud Native Core, Network Repository Function
OSO	Operations Services Overlay
RAM	Random Access Memory
RTT	Round Trip Time
SBA	Service Based Architecture
SUT	System Under Test
Set	A set comprises of multiple Georedundant NRF sites
TLS	Transport Layer Security
XFCC	x-forwarded-client-cert

What's New in This Guide

This section lists the documentation updates for release 25.2.2xx.

Release 25.2.201- G52536-01, April 2026

Generic Updates

- Added the System Under Test (SUT) acronym in the [Acronyms](#) section.
- Updated the cnDBTier version in the [System Software](#) section.
- Updated the NRF version test constraint, Average HTTP Request Packet Size, Average HTTP Response Packet Size, and Average size of NF Profiles in the [Software Test Constraints](#) section.
- Updated the CNE common services versions in the [Observability Services](#) section.
- Updated the CNE services names in the [CNE Common Services](#) section.
- Updated the images for singleset and multiset topology in the [Single Set Test Topology](#) and [Multiset Test Topology \(NRF Growth\)](#) sections.

Benchmark Testing Updates

- Updated the [NRF Features](#) section.
- Updated the following sections in the Observations for 51.6K TPS chapter:
 - [Service Operations](#)
 - [cnDBTier Resource Requirements per NRF Deployment](#)
 - [NF Subscription Count](#)
 - [Call-Mix](#)
 - [NF Count Per NF Types](#)
 - [Discovery TPS per NfType and Number of Profiles in Response](#)
 - [Test Observations](#)
 - [NRF Resource Utilization](#)
 - [cnDBTier Services Resource Utilization](#)
 - [Latency Observations](#)
 - Removed the "NRF Configuration" section as it has redundant information and added a note for the following features to use the recommended values as suggested in the configuration in the [NRF Features](#) section:
 - * Overload Control
 - * Pod Protection Support for NRF Subscription Microservice
 - * Ingress Gateway Pod Protection Using Rate Limiting
 - * Egress Gateway Pod Protection Using Rate Limiting

1

Introduction

Oracle Communications Cloud Native Core Network Repository Function (NRF) is one of the main components of the 5G Service Based Architecture (SBA). NRF maintains an updated repository of all the Network Functions (NFs) available in the operator's network along with the services provided by each NFs in the 5G core. NRF instantiates, scales, and terminates the services with minimal or no manual intervention.

For more information about NRF architecture, see *Oracle Communications Cloud Native Core, Network Repository Function User Guide*.

1.1 Purpose and Scope

This document describes the measurements that are used to perform the performance evaluation of NRF, microservices, and deployment environment setup software such as Aspen Service Mesh (ASM), Operations Services Overlay (OSO), Cloud Native Environment (CNE), and so on.

This document provides the following information:

- Benchmark NRF performance and capacity
- Benchmark data from our labs
- Key metrics used to manage NRF performance and capacity

Note

The performance and capacity of the NRF system may vary based on the call model, Feature or Interface configuration, and underlying CNE and hardware environment.

1.2 References

Following are the reference documents:

- *Oracle Communications Cloud Native Core, Cloud Native Environment Installation, Upgrade, and Fault Recovery Guide*
- *Oracle Communications Cloud Native Core, Operations Services Overlay Installation and Upgrade Guide*
- *Oracle Communications Cloud Native Core, cnDBTier User Guide*
- *Oracle Communications Cloud Native Core, Network Repository Function User Guide*
- *Oracle Communications Cloud Native Core, Network Repository Function REST Specification Guide*
- *Oracle Communications Cloud Native Core, Network Repository Function Troubleshooting Guide*
- *Oracle Communications Cloud Native Core, Network Repository Function Network Impact Report*

- *Oracle Communications Cloud Native Configuration Console User Guide*
- *Oracle Communications 5G Automated Test Suite Guide*
- *Oracle Communication Cloud Native Core, Data Collector User Guide*

2

Deployment Environment

This section provides information about the cloud native infrastructure used for Oracle Communications Cloud Native Core, Network Repository Function (NRF) benchmarking.

2.1 Deployed Components

This section provides details about the deployed components.

2.1.1 Hardware Details

This section describes the hardware details.

Table 2-1 Hardware Details

Nodes	Type
Master Nodes	ORACLE SERVER X8-2
Top of Rack Switch	Cisco 93108TC-EX
Worker Nodes	ORACLE SERVER X8-2

2.1.2 System Software

This section describes the system software details.

Table 2-2 System Software

System Software	Details
Operating System (+Kernel Version)	5.15.0-105.125.6.2.2.el8uek.x86_64 Oracle Linux Server 8.8
Hypervisor	BareMetal Server
CNE	23.4.1
OSO	NA
Kubernetes	1.27.x
ASM	1.21.6
cnDBTier	25.2.200

2.1.3 Observability Services

This section describes the observability services used for NRF benchmarking.

Table 2-3 Observability Services

Service Name	Version
Fluentd OpenSearch	1.16.2
Grafana	9.5.3
Jaeger	1.52.0
Kyverno	1.9.0
MetalLB	0.14.4
OpenSearch	2.11.0
OpenSearch Dashboard	2.11.0
Prometheus	2.51.1

2.1.4 CNE Common Services

The following are the CNE common services observed in the deployment:

- alertmanager-occne-kube-prom-stack-kube-alertmanager
- occne-alertmanager-snmp-notifier
- occne-bastion-controller
- occne-fluentd
- occne-kube-prom-stack-grafana
- occne-kube-prom-stack-kube-operator
- occne-kube-prom-stack-kube-state-metrics
- occne-kube-prom-stack-prometheus-node
- occne-metallb
- occne-metallb-controller
- occne-metrics-server
- occne-opensearch-cluster-client
- occne-opensearch-cluster-master
- occne-opensearch-dashboards
- occne-promxy
- occne-promxy-apigw-nginx
- occne-tracer-jaeger
- occne-tracer-jaeger-collector
- occne-tracer-jaeger-query
- prometheus-occne-kube-prom-stack-kube-prometheus

3

Test Topology

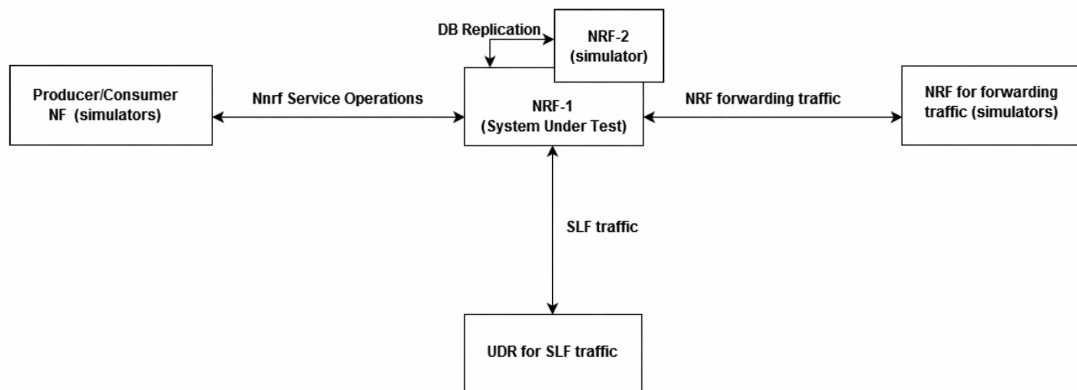
This section describes the topology used for testing Oracle Communications Cloud Native Core, Network Repository Function (NRF).

3.1 Single Set Test Topology

The following image represents the test topology consisting of the following components:

- Two-site NRF
- UDR for SLF traffic
- Producer/Consumer NF for Nnrf service operations (simulators)
- NRF for forwarding to a different segment (simulators)

Figure 3-1 Singleset Topology



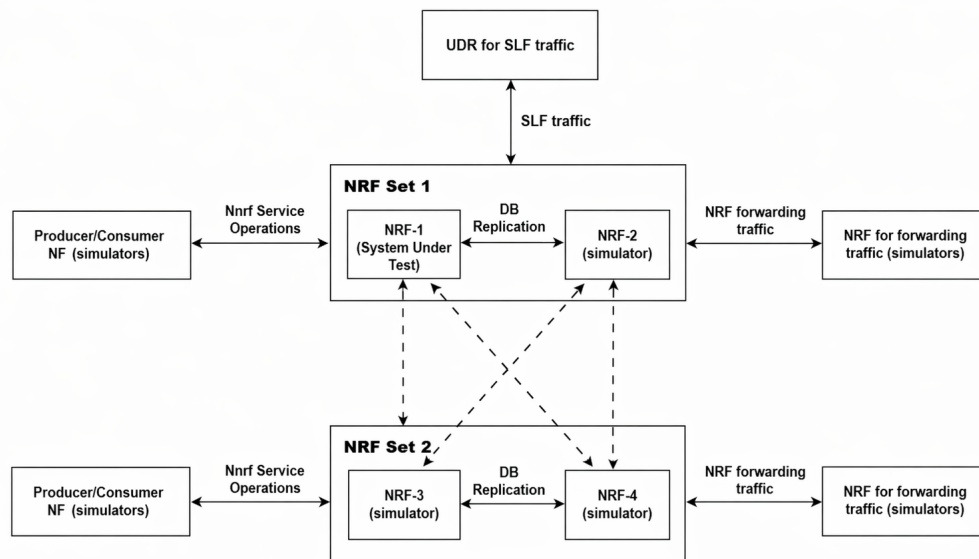
1. Producer NF sends NfRegister, NfHeartbeat, NfUpdate, and NfDeregister requests to NRF.
2. Consumer NF sends NfStatusSubscribe, NfStatusSubscribe (PATCH), NfStatusUnsubscribe, NfDiscover, NfAccessToken, NfProfileRetrieval, and NfListRetrieval requests to NRF.
3. NRF sends request to UDR for processing SLF based discovery query.
4. NRF sends response to all service requests it receives.
5. NRF triggers NfStatusNotify to all the Consumer NFs who have created subscription.
6. NRF forwards NfDiscovery requests to the NRF simulator.

3.2 Multiset Test Topology (NRF Growth)

The following image represents the test topology consisting of the following components:

- Two NRF Sets with one NRF Set having two NRF sites and the other set with one NRF site.
- UDR for SLF traffic
- Producer/Consumer NF for Nnrf service operations (simulators)
- NRF for forwarding to a different segment (simulators)

Figure 3-2 Multiset Topology



1. Producer NF sends NfRegister, NfHeartbeat, NfUpdate, and NfDeregister requests to NRF.
2. Consumer NF sends NfStatusSubscribe, NfStatusSubscribe (PATCH), NfStatusUnsubscribe, NfDiscover, NfAccessToken, NfProfileRetrieval, and NfListRetrieval requests to NRF.
3. NRF sends request to UDR for processing SLF based discovery query.
4. NRF sends response to all service requests it receives.
5. NRF triggers NfStatusNotify to all the Consumer NFs who have created subscription.
6. NRF forwards NfDiscovery requests to the NRF simulator.
7. NRF Set 1 retrieves the remote set state data from NRF Set 2 and vice versa.

4

Benchmark Testing

This section describes the environment used for benchmarking Oracle Communications Cloud Native Core, Network Repository Function (NRF). Benchmarking is performed with the settings described in this section.

The default values or recommendations for any required software or resource are available from third-party vendors. The operators may choose different values.

4.1 Overview

To qualify the test run, the following elements are considered:

- Pod restart
- CPU and Memory utilization
- Error rate
- Ingress and egress traffic rate
- Success rate
- Message request and response processing time
- Infrastructure resource requirements and utilization

4.2 NRF Features

The following table lists the features that are enabled during benchmarking.

Table 4-1 NRF Features

NRF Features	Feature Status	Notes
Subscriber Location Function	ENABLED	SLF is enabled for 3 NfTypes (AUSF, UDR, and UDM) in slfLookupConfig.
SLF Selection from Registered NF Profiles	ENABLED	Dynamic SLF feature is enabled.
Access Token Caching during SLF Query	ENABLED	Access Token Caching during SLF Query is enabled with 24 hour validity period based caching. For more information about the accessTokenCacheEnabled under slfOptions, see <i>Oracle Communications Cloud Native Core, Network Repository Function REST Specification Guide</i> .
Overload Control	ENABLED	This feature must be configured with the recommended values as mentioned in the <i>Oracle Communications Cloud Native Core, Network Repository Function REST Specification Guide</i> . Note: No feature specific failures are induced in the performance test.

Table 4-1 (Cont.) NRF Features

NRF Features	Feature Status	Notes
Preferred Locality Feature Set	ENABLED	Extended Preferred locality is configured with 2 locationTypes, 8 locationMappings, 56 preferredLocationDetails.
NRF Forwarding per NF Type/ Service Type	ENABLED for UDM	Forwarding is enabled for discovery services. 1 host is present in nrfHostConfig and 1 rule is present in forwardingRulesConfig.
EmptyList in Discovery Response	ENABLED	EmptyList feature is enabled and discoveryValidityPeriodCfg under nfDiscoveryOptions is configured with 5 rules. Note: No feature specific failures are induced in the performance test.
Access Token Request Authorization	ENABLED	AuthFeature is enabled with 8 rules in authRulesConfig under nfAccessTokenOptions.
Pod Protection Support for NRF Subscription Microservice	ENABLED	This feature must be configured with the recommended values as mentioned in the <i>Oracle Communications Cloud Native Core, Network Repository Function REST Specification Guide</i> . Note: No feature specific failures are induced in the performance test.
Subscription Limit	ENABLED	Set the subscription limit to 1000. For more information about the configuration, see <i>Oracle Communications Cloud Native Core, Network Repository Function REST Specification Guide</i> .
User-Agent Header for Outgoing Requests	ENABLED	For more information about the configuration, see <i>Oracle Communications Cloud Native Core, Network Repository Function REST Specification Guide</i> .
NRF to Pick Port from ipEndpoints and FQDN from the NF Service or NF level when Selecting SLF	ENABLED	For more information about the configuration, see <i>Oracle Communications Cloud Native Core, Network Repository Function REST Specification Guide</i> .
ASM Sidecar	ENABLED	For more information about the configuration, see <i>Oracle Communications Cloud Native Core, Network Repository Function REST Specification Guide</i> .
Maximum Number of Profiles in Discovery Response	ENABLED	Set the value as 12 for the profilesCountInDiscoveryResponse parameter. For more information about the configuration, see <i>Oracle Communications Cloud Native Core, Network Repository Function REST Specification Guide</i> . Note: If the value is greater than 12, there will be a performance impact.
NRF Georedundancy	ENABLED	For more information about the configuration, see <i>Oracle Communications Cloud Native Core, Network Repository Function REST Specification Guide</i> .
NRF Growth	ENABLED	For more information about the configuration, see <i>Oracle Communications Cloud Native Core, Network Repository Function REST Specification Guide</i> .
SLF Routing using SCP	ENABLED	For more information about the configuration, see <i>Oracle Communications Cloud Native Core, Network Repository Function REST Specification Guide</i> .

Table 4-1 (Cont.) NRF Features

NRF Features	Feature Status	Notes
Discovery Validity	ENABLED	Discovery validity feature is enabled with a value as 1 minute (for all NF Types). For more information about the configuration, see <i>Oracle Communications Cloud Native Core, Network Repository Function REST Specification Guide</i> .
Access Token Expiry Time	NA	The expiry duration for a NRF signed certificate. The limit is 24 hours. For more information about NRF Access Token Service Usage Details, see <i>Oracle Communications Cloud Native Core, Network Repository Function User Guide</i> .
Load Change Threshold for Notifications from NRF	NA	This feature is configured using <code>nfNotifyLoadThreshold</code> parameter. Set the value to 10%. For more information about the configuration, see <i>Oracle Communications Cloud Native Core, Network Repository Function REST Specification Guide</i> .
Monitoring the Availability of SCP Using SCP Health APIs	ENABLED	For more information about the configuration, see <i>Oracle Communications Cloud Native Core, Network Repository Function REST Specification Guide</i> . Note: No feature specific failures are induced in the performance test.
NRF to Retry Alternate Producer to Support Maximum Diversity	NA	The value of <code>maxSlfAttempts</code> parameter is set to 1 and the value of <code>useAlternateScp</code> parameter is set to true. For more information about the configuration, see <i>Oracle Communications Cloud Native Core, Network Repository Function REST Specification Guide</i> . Note: No feature specific failures are induced in the performance test.
Support for TLS	DISABLED	For more information about the configuration, see <i>Oracle Communications Cloud Native Core, Network Repository Function Installation, Upgrade, and Fault Recovery Guide</i> .
Error Response Enhancement	NA	For more information about the configuration, see <i>Oracle Communications Cloud Native Core, Network Repository Function REST Specification Guide</i> .
Error Log Enhancements	ENABLED	For more information about the configuration, see <i>Oracle Communications Cloud Native Core, Network Repository Function REST Specification Guide</i> .
NRF support for A_PFD DataSetId in registration and discovery (DataSetId Enhancements)	ENABLED	For more information about the configuration, see <i>Oracle Communications Cloud Native Core, Network Repository Function REST Specification Guide</i> .
NFService Priority Update	ENABLED	For more information about the configuration, see <i>Oracle Communications Cloud Native Core, Network Repository Function REST Specification Guide</i> .
Subscriber Identity in messages (3GPP-SBI-Correlation-Info)	ENABLED	For more information about the configuration, see <i>Oracle Communications Cloud Native Core, Network Repository Function REST Specification Guide</i> .
Ingress Gateway Pod Protection	DISABLED	For more information about the configuration, see <i>Oracle Communications Cloud Native Core, Network Repository Function REST Specification Guide</i> .

Table 4-1 (Cont.) NRF Features

NRF Features	Feature Status	Notes
Egress Gateway Pod Throttling	DISABLED	For more information about the configuration, see <i>Oracle Communications Cloud Native Core, Network Repository Function Installation, Upgrade and Fault Recovery Guide</i> .
NF Profile Size Limit	ENABLED	For more information about the configuration, see <i>Oracle Communications Cloud Native Core, Network Repository Function REST Specification Guide</i> .
Ingress Gateway Pod Protection Using Rate Limiting	ENABLED	This feature must be configured with the recommended values as mentioned in the <i>Oracle Communications Cloud Native Core, Network Repository Function REST Specification Guide</i> . Note: <ul style="list-style-type: none"> The Ingress Gateway Pod Protection using Rate Limiting feature is recommended over the Egress Gateway Pod Throttling feature. No feature specific failures are induced in the performance test.
Egress Gateway Pod Protection Using Rate Limiting	ENABLED	This feature must be configured with the recommended values as mentioned in the <i>Oracle Communications Cloud Native Core, Network Repository Function REST Specification Guide</i> . Note: <ul style="list-style-type: none"> The Egress Gateway Pod Protection using Rate Limiting feature is recommended over the Egress Gateway Pod Throttling feature. No feature specific failures are induced in the performance test.
Rerouting SLF Requests Based on Server Header Value	ENABLED	For more information about the configuration, see <i>Oracle Communications Cloud Native Core, Network Repository Function REST Specification Guide</i> . Note: No feature specific failures are induced in the performance test.

Note

- Apart from these features being enabled, rest of the feature configurations have been set to their default values. For more information about the feature configurations, see *Oracle Communications Cloud Native Core, Network Repository Function REST Specification Guide*. For more information about the features, see *Oracle Communications Cloud Native Core, Network Repository Function User Guide*.
- No feature specific failures are induced in the performance test.

4.3 Software Test Constraints

Table 4-2 Software Test Constraints

Test Constraint	Details
NRF Version	25.2.201
Average HTTP Request Packet Size	~630 to 15000 bytes (Based on NfRegister Request)
Average HTTP Response Packet Size	~5500-131000 bytes (Discovery Response with nfProfiles)
Average NF Heartbeats Timer per NF	30s
Average size of NF Profiles	~630 to 15000 bytes

4.4 Service Operations

This section list the various service operations with the required query parameters.

NfDiscover Service Operation

The NfDiscover service operations are sent with the following query parameters:

- Target-nf-type, guami, and requester-nssais
- Target-nf-type, requester-nf-type, routing-indicator, and group-id-list
- Target-nf-type, requester-nf-type, supi, and service-names
- Target-nf-type, requester-nf-type, and service-names
- Target-nf-type, requester-nf-type, and group-id-list
- Target-nf-type and pgw
- Target-nf-type, requester-nf-type, supi, and service-names
- Target-nf-type, requester-nf-type, and service-names
- Target-nf-type, dnn, snssai, TAI, and pgw-ind
- Target-nf-type, SUPI, group-id-list, and service-names
- Target-nf-type, dnn, snssai, preferred-locality, and service-names
- Target-nf-type and guami
- Target-nf-type, requester-nf-type, and preferred-locality
- Target-nf-type, requester-nf-type, preferred-locality, and SUPI
- Target-nf-type, requester-nf-type, SUPI, group-id-list, requester-nf-type, and data-set
- Target-nf-type, requester-nf-type, routing-indicator, group-id-list, and service-names
- Target-nf-type, servicenames, requester-nf-type, requester-nf-instance-fqdn, requester-plmn-list, and guami
- Target-nf-type, requester-nf-type, and target-nf-instance-id
- Target-nf-type, requester-nf-type, target-nf-instance-id, and service_name
- Target-nf-type, requester-nf-type, target-nf-instance-id, and notification-type
- Target-nf-type, requester-nf-type, dnn, snssais, preferred-locality, ue-ipv4-address, ip-domain, and ue-ipv6-prefix

- Target-nf-type, requester-nf-type, service-names, requester-nf-instance-fqdn, requester-plmn-list, preferred-locality, notification-type, n1-msg-class, and client-type
- Target-nf-type, requester-nf-type, gpsi, and preferred-locality
- Target-nf-type, requester-nf-type, service-names, requester-nf-instance-fqdn, requester-plmn-list, preferred-locality, and gpsi
- Target-nf-type, requester-nf-type, service-names, gpsi, and preferred-locality

4.5 53K TPS per NRF Set/106K TPS for Two NRF Sets

The per set capacity of NRF is 53K TPS. Higher TPS rate can be achieved by deploying additional NRF sets. This section provides the details for 106K TPS.

4.5.1 NRF Services Resource Requirements

This section provides the NRF services resource requirement per NRF deployment.

Note

Changes in the resource requirements are highlighted in bold.

Table 4-3 NRF Services Resource Requirements

Service Name	Min Pod Replica	Max Pod Replica	Min CPU/Pod	Max CPU/Pod	Min Memory/Pod (in Gi)	Max Memory/Pod (in Gi)	Min (Mi) Ephemeral Storage	Max (Gi) Ephemeral Storage
Helm test	1	1	1	2	1	2	78.1	1
<helm-release-name>-nfrgistration	2	2	4	4	3	3	78.1	1
<helm-release-name>-nfdiscovery	2	42	8	8	5	5	78.1	2
<helm-release-name>-nfsubscription	2	2	2	2	3	3	78.1	1
<helm-release-name>-nrfauditor	2	2	2	2	3	3	78.1	1
<helm-release-name>-nrfconfiguration	1	1	2	2	2	2	78.1	1

Table 4-3 (Cont.) NRF Services Resource Requirements

Service Name	Min Pod Replica	Max Pod Replica	Min CPU/Pod	Max CPU/Pod	Min Memory/Pod (in Gi)	Max Memory/Pod (in Gi)	Min (Mi) Ephemeral Storage	Max (Gi) Ephemeral Storage
<helm-release-name>-nfaccessstoken	2	2	2	2	2	2	78.1	1
<helm-release-name>-nrfartisan	1	1	2	2	2	2	78.1	1
<helm-release-name>-nrfcachedata	2	2	4	4	4	4	78.1	1
<helm-release-name>-ingressgateway	2	30	4	4	5	5	78.1	1
<helm-release-name>-egressgateway	2	30	4	4	4	4	78.1	1
<helm-release-name>-alternate-route	2	2	2	2	4	4	78.1	1
<helm-release-name>-appinfo	2	2	1	1	1	1	78.1	1
<helm-release-name>-perfinfo	2	2	1	1	1	1	78.1	1

Table 4-4 Service Mesh Sidecar Resource Requirements

Service Name	CPU/Pod		Memory/Pod (in G)		Concurrency
	Min	Max	Min	Max	
Helm test	0	0	0	0	NA
<helm-release-name>-nfreistration	2	2	3	3	2
<helm-release-name>-nfdiscovery	2	2	3	3	4
<helm-release-name>-nfsubscription	2	2	3	3	2
<helm-release-name>-nrfauditor	2	2	3	3	2

Table 4-4 (Cont.) Service Mesh Sidecar Resource Requirements

Service Name	CPU/Pod		Memory/Pod (in G)		Concurren- cy
	Min	Max	Min	Max	
<helm-release-name>-nrfconfiguration	2	2	3	3	2
<helm-release-name>-nrfaccesstoken	2	2	3	3	2
<helm-release-name>-nrfartisan	2	2	3	3	2
<helm-release-name>-nrfcachedata	2	2	3	3	2
<helm-release-name>-ingressgateway	4	4	3	3	8
<helm-release-name>-egressgateway	4	4	3	3	8
<helm-release-name>-alternate-route	2	2	3	3	2
<helm-release-name>-appinfo	2	2	3	3	2
<helm-release-name>-perfinfo	2	2	3	3	2

4.5.2 cnDBTier Resource Requirements per NRF Deployment

This section provides the details of cnDBTier services resource requirements.

Note

Changes in the resource requirements are highlighted in bold.

The following table describes the resources required by cnDBTier pods to perform NRF benchmark tests.

Table 4-5 cnDBTier Services Resource Requirements

Service Name	Pod Rep lica #	CPU/Pod		Memory/Pod (in Gi)		PVC Size (in Gi)		Ephemeral Storage	
		Min	Max	Min	Max	PVC1	PVC2	Min (Mi)	Max (Gi)
db-monitor-svc	1	4	4	4	4	0	NA	200	1
db-replication-svc (leader)	1	2	2	12	12	4	NA	200	1
db-replication-svc (other)	1	1.1	1.1	2	2	0	NA	200	1
db-backup-manager- svc	1	1.1	1.1	1	1	0	NA	200	1
ndbappmysqld (Access)	2	8	8	3	3	1	NA	200	1
ndbmgmd (Mgmt)	2	4	4	10	10	15	NA	200	1
ndbmt (Data)	4	4	4	5	5	4	5	200	1

Table 4-5 (Cont.) cnDBTier Services Resource Requirements

Service Name	Pod Rep lica #	CPU/Pod		Memory/Pod (in Gi)		PVC Size (in Gi)		Ephemeral Storage	
		Min	Max	Min	Max	PVC1	PVC2	Min (Mi)	Max (Gi)
ndbmysqld (SQL Replication)	4	4	4	20	20	90	NA	200	1

The following table describes the resources required by cnDBTier service mesh sidecar to perform NRF benchmark tests.

Table 4-6 Service Mesh Sidecar

Service Name	CPU		Memory (in Gi)		Concurre ncy
	Min	Max	Min	Max	
db-monitor-svc	2	2	1	1	2
db-replication-svc (leader)	2	2	1	1	2
db-replication-svc (other)	2	2	1	1	2
db-backup-manager-svc	2	2	1	1	2
ndbappmysqld (Access)	4	4	1	1	8
ndbmgmd (Mgmt)	2	2	1	1	8
ndbmtl (Data)	2	2	1	1	8
ndbmysqld (SQL Replication)	2	2	1	1	8

The following default values are added for the service mesh specific annotations in the `ocnrf_dbtier_CNDBTIER_VERSION_custom_values_NRF_VERSION.yaml` file.

The following table describes the default values for resources required by cnDBTier service mesh sidecar to perform NRF benchmark tests.

Table 4-7 Default Values for Service Mesh Specific Annotations

Parameter Name	Annotations
<code>mgm.annotations</code>	<pre> sidecar.istio.io/inject: "true" sidecar.istio.io/proxyCPU: "2" sidecar.istio.io/proxyCPULimit: "2" sidecar.istio.io/proxyMemory: "1Gi" sidecar.istio.io/proxyMemoryLimit: "1Gi" proxy.istio.io/config: '{concurrency: 8}' traffic.sidecar.istio.io/excludeInboundPorts: "8081" </pre>

Table 4-7 (Cont.) Default Values for Service Mesh Specific Annotations

Parameter Name	Annotations
ndb.annotations	<pre> sidecar.istio.io/inject: "true" sidecar.istio.io/proxyCPU: "2" sidecar.istio.io/proxyCPULimit: "2" sidecar.istio.io/proxyMemory: "1Gi" sidecar.istio.io/proxyMemoryLimit: "1Gi" proxy.istio.io/config: '{concurrency: 8}' traffic.sidecar.istio.io/excludeInboundPorts: "8081" </pre>
api.annotations	<pre> sidecar.istio.io/inject: "true" sidecar.istio.io/proxyCPU: "2" sidecar.istio.io/proxyCPULimit: "2" sidecar.istio.io/proxyMemory: "1Gi" sidecar.istio.io/proxyMemoryLimit: "1Gi" proxy.istio.io/config: '{concurrency: 8}' traffic.sidecar.istio.io/excludeInboundPorts: "8081" </pre>
ndbapp.annotations	<pre> sidecar.istio.io/inject: "true" sidecar.istio.io/proxyCPU: "4" sidecar.istio.io/proxyCPULimit: "4" sidecar.istio.io/proxyMemory: "1Gi" sidecar.istio.io/proxyMemoryLimit: "1Gi" proxy.istio.io/config: '{concurrency: 8}' traffic.sidecar.istio.io/excludeInboundPorts: "8081" </pre>
db-replication-svc.podAnnotations	<pre> sidecar.istio.io/inject: "true" sidecar.istio.io/proxyCPU: "2" sidecar.istio.io/proxyCPULimit: "2" sidecar.istio.io/proxyMemory: "1Gi" sidecar.istio.io/proxyMemoryLimit: "1Gi" proxy.istio.io/config: '{concurrency: 2}' </pre> <p>Note: The annotations for db-replication-svc.podAnnotations varies depending on the replication cardinality. podAnnotations should be added for each such replication deployment section in the ocnrf_dbtier_CNDBTIER_VERSION_custom_values_NRF_VERSION.yaml file.</p>
db-replication-svc.podAnnotations	<pre> traffic.sidecar.istio.io/excludeInboundPorts: "8081" </pre> <p>The inbound ports are added only to the db-replication-svc (leader pod).</p>

Table 4-7 (Cont.) Default Values for Service Mesh Specific Annotations

Parameter Name	Annotations
db-monitor-svc.podAnnotations	<pre> sidecar.istio.io/inject: "true" sidecar.istio.io/proxyCPU: "2" sidecar.istio.io/proxyCPULimit: "2" sidecar.istio.io/proxyMemory: "1Gi" sidecar.istio.io/proxyMemoryLimit: "1Gi" proxy.istio.io/config: '{concurrency: 2}' traffic.sidecar.istio.io/excludeInboundPorts: "8081" </pre>
db-backup-manager-svc.pod.annotations	<pre> sidecar.istio.io/inject: "true" sidecar.istio.io/proxyCPU: "2" sidecar.istio.io/proxyCPULimit: "2" sidecar.istio.io/proxyMemory: "1Gi" sidecar.istio.io/proxyMemoryLimit: "1Gi" proxy.istio.io/config: '{concurrency: 2}' </pre>

4.5.3 State Data Quantum at NRF

4.5.3.1 NF Count Per NF Types

The total number of NFs supported by per NRF set is 300. For two sets a total of 600 NFs are supported.

The following table lists the NF count per NF type along with the NF profile size.

Table 4-8 NF Count Per NF Types

Target-NF Types	NF Count	Profile Size (in bytes)
AMF	46	15000
AUSF	48	4640
BSF	3	1700
CBCF	3	11300
CHF	3	1800
GMLC	3	0630
LMF	3	2400
NEF	3	11300
PCF	79	1500-2500
SCP	27	1200
SEPP	3	300
SLF	3	2100
SMF	148	15000

Table 4-8 (Cont.) NF Count Per NF Types

Target-NF Types	NF Count	Profile Size (in bytes)
SMSF	3	11300
UDM	48	3200
UDR	177	1700-2000

4.5.3.2 NF Subscription Count

The total number of NF subscriptions supported by per NRF set is 750. For two sets a total of 1500 NFs are supported. The following table lists the NF subscription count.

Table 4-9 NF Subscription Count

Subscription Type	Subscription Count
Subscription to a set of NF Instances identified by their NF Type	412
Subscription to a set of NF Instances identified by their NF Instance Id	211
Subscription to a set of NF Instances identified by their NF Set Id	127

① Note

- Subscriptions are configured with `NotificationCondition` to monitor only the “/priority” or “/load” attributes.
- The `reqNotifEvents` include `NF_REGISTERED`, `NF_DEREGISTERED`, and `NF_PROFILE_CHANGE`, ensuring notifications for status change events and relevant profile updates.

4.5.4 Call-Mix

This section describes the various type of service requests with the required traffic distribution and targeted TPS per NRF set.

Table 4-10 Call-Mix

S.No	Type of Request	Traffic Distribution (in %)	Targeted TPS
A	NF Register	0.0094	5
B	NF Update	0.090	43
C	NF Heartbeat	-	155
D	NF Subscribe + Update Subscription	0.037	20
E	NF Discover	97.17	51600
F	NRF Forwarded Discovery (Incoming)	10% of Incoming discovery traffic	0

Table 4-10 (Cont.) Call-Mix

S.No	Type of Request	Traffic Distribution (in %)	Targeted TPS
G	% of traffic forwarded to remote PLMN (case Visited NRF)	0	Not included in this test.
H	% of traffic received from remote PLMN (case Home NRF)	0	Not included in this test.
I	SLF Query	88.4% of Discovery	45600
J	NF Access Token for SLF	Extra	1
K	NF Profile Retrieval	2.29	1210
L	NF List Retrieval	0.036	19

Table 4-10 (Cont.) Call-Mix

S.No	Type of Request	Traffic Distribution (in %)	Targeted TPS
M	NF Status Notify	4.55	2155
N	NF Status UnSubscribe	0.0057	3
O	NF DeRegister	0.0057	3
P	NF Access Token for Consumers	0.0188	10
Q	NRF Forwarding (Outgoing)	15% of Discovery	7740

Table 4-11 Total Traffic

Traffic Type	Message Considered	Targeted TPS
Total Incoming Traffic	A + B + C + D + E + F + G + H + J + K + L + N + O + P	53068
Total Outgoing Traffic	I + J + M + Q	55495

4.5.4.1 Discovery TPS per NfType and Number of Profiles in Response

The following table indicates the discovery TPS per NfType per NRF set and the number of profiles sent in the response.

Table 4-12 Discovery TPS per NfType and Number of Profiles in Response

Target-Nf Types	TPS (in %)	Absolute Discovery TPS	SLF Queries Posted	Number of Profiles in Response
AMF	4.24	2281	0	4
AUSF	0.39	209	0	6
BSF	0.00	0	0	3
CBCF	0.00	0	0	3
CHF	0.077	42.3	0	3
GMLC	0.00	0	0	4
LMF	0.025	14.5	0	4
NEF	0.026	14.1	0	3
PCF	0.39	211	0	10
SLF	0	0	0	3
SMF	8.09	4368	0	4
SMSF	0.025	13.5	0	3
UDM	85.5	46209	46209	6
UDR	1.21	652	652	12

4.5.5 Test Observations

This section following table provides the details of test duration used for the benchmarking test.

Table 4-13 Test Observations

Parameter	Values
Test Duration	48 hours
TPS Achieved	53608
Ingress Gateway incoming connections from the NF Peers (from the Ingress Gateway ASM sidecar to Ingress Gateway application container)	256
Ingress Gateway outgoing connections towards the backend services (from the Ingress Gateway application container to Ingress Gateway ASM sidecar)	2.4K
Egress Gateway incoming connections from the backend services (from the Egress Gateway ASM sidecar to Egress Gateway application container)	240
Egress Gateway outgoing connections towards NF Peers (from the Egress Gateway application container to Egress Gateway sidecar)	1.2K

4.5.6 NRF Resource Utilization

The following table describes NRF microservices and their utilization for 53K TPS per NRF set with 600 NFs in two sets.

Table 4-14 NRF Resource Utilization

NRF Microservices	Number of Pods	Min CPU	Max CPU	Average CPU	Min Memory	Max Memory	Average Memory
nfdiscovery	42	52.2%	53.6%	52.9%	51.5%	51.9%	51.7%
ingressgateway	30	56.3%	57.9%	57.4%	65.7%	67.4%	66.7%
egressgateway	30	51.5%	53.0%	52.3%	43.0%	44.1%	43.7%
nfregistration	2	26%	30.4%	30%	37.0%	37.8%	37.4%
nfsubscription	2	14.6%	17.5%	16.8%	32.7%	33.2%	33%
nrfauditor	2	0.434%	0.582%	0.54%	17.8%	18.1%	17.9%
nfaccessstoken	2	1.88%	2.05%	1.92%	37.4%	37.9%	37.6%
nrfcachedata	2	35.2%	36.6%	35.9%	49.8%	53.8%	53%
alternate-route	2	0.0758%	0.103%	0.08%	16.9%	16.9%	16.9%
perf-info	2	6.63%	7.71%	7.67%	14.1%	14.2%	14.1%
appinfo	2	4.42%	4.74%	4.48%	28.5%	28.5%	28.5%
nrfconfiguration	1	1.61%	1.89%	1.65%	37.4%	37.9%	37.7%
nrfartisan	1	0.106%	0.150%	0.10%	22.7%	22.8%	22.7%

Note

If you enable Message Feed feature at Ingress Gateway and Egress Gateway, approximately 33% pod capacity is impacted. The features enabled for this benchmark testing does not utilize alternate-route service.

4.5.7 cnDBTier Services Resource Utilization

The following table provides observed values of cnDBTier services.

Table 4-15 cnDBTier Services Resource Utilization

Service Name	Pod Count	Min CPU	Max CPU	Average CPU	Min Memory	Max Memory	Average Memory
ndbappmysqlid (Access)	2	6.45%	16.4%	11.5%	28.5%	28.7%	28.6%
ndbmtid (Data)	4	6.21%	9.87%	8.98%	91.1%	91.2%	91.2%
ndbmgmd (Mgmt)	2	0.116%	0.193%	0.172%	2.19%	2.20%	2.19%
db-monitor-svc	1	0.192%	1.03%	0.304%	21.4%	22.0%	21.6%
db-backup-manager-svc	1	0.152%	1.30%	0.217%	21.3%	22.2%	21.4%

Table 4-16 cnDBTier Statistics

DBTier Statistics	Value
Disk Write operations per second	1208
Disk Read operations per second	4897
Transaction rates on data nodes	2416

4.5.8 Latency Observations

This section provides information about the latency observations.

Table 4-17 Consolidated Latency Metrics Across Discovery Service

Latency Parameter	Details			
	99th Percentile (ms)	95th Percentile (ms)	90th Percentile (ms)	50th Percentile (ms)
Average Turnaround time at Ingress Gateway, Discovery Processing, and SLF Egress Transaction	206	195	191	151
Average Turnaround time for Discovery Processing and SLF Egress Transaction	200	195	190	150
Average Turnaround time for SLF Egress Transaction (see Note)	199	195	190	150

Note

Following simulators are used for performance measurement and their induced latency:

- NF simulator to NRF: 50 ms
- NRF to forwarding simulator and forwarding simulator to NRF: 150 ms
- NRF to SLF: 150 ms

The latency value of Min, Max, and Average for discovery service operation is maximum among all the other service operations.

Table 4-18 Consolidated Latency Metrics Across Non-discovery Service Operations

Service Operation	Latency Parameter	99th Percentile (ms)	95th Percentile (ms)	90th Percentile (ms)	50th Percentile (ms)
NfStatusSubscribe (POST)	Average Turnaround time at Ingress for all subscription service operations	41.9	1.90	1.80	1.00

Table 4-18 (Cont.) Consolidated Latency Metrics Across Non-discovery Service Operations

Service Operation	Latency Parameter	99th Percentile (ms)	95th Percentile (ms)	90th Percentile (ms)	50th Percentile (ms)
NfStatusSubscribe (PATCH)	Average Turnaround time at Ingress for all subscription service operations	79.5	53.9	42.8	26.2
NfStatusUnsubscribe	Average Turnaround time at Ingress for all subscription service operations	64.0	39.3	29.7	16.3
NfAccessToken	Latency observed at Ingress gateway for all access token service operations	29.7	19.3	17.8	8.91
NfRegistration (PUT)	Latency observed at Ingress for NfRegister/ NfUpdate	63.7	39.6	38.0	27.4
NfUpdate (PATCH)	Latency observed at Ingress for NfRegister/ NfUpdate	49.5	29.6	28.7	21.4
NfDeregister	Latency observed at Ingress for NfDeregister	50.1	20.2	19.4	15.2
NFProfileRetrieval	Latency observed at Ingress for NfProfileRetrieval	29.2	17.2	12.0	6.69
NfListRetrieval	Latency observed at Ingress for NfListRetrieval	29.2	17.2	12.0	6.69