

Oracle® Communications

Cloud Native Core, Unified Data Repository

Benchmarking Guide



Release 23.4.1
F90996-03
September 2025

ORACLE®

Copyright © 2022, 2025, 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
2.2	Deployment Resources	1
2.2.1	Resource Requirements for CNE Observability Services	2
3	UDR Benchmark Testing	
3.1	Test Scenario 1: SLF Call Deployment Model	1
3.1.1	SLF Call Model: 24.2K TPS for Performance-Medium Resource Profile for SLF Lookup	1
3.1.2	SLF Call Model: 50K lookup TPS + 1.26K Provisioning TPS	9
3.2	Test Scenario 2: EIR Deployment Model	14
3.3	Test Scenario 3: SOAP and Diameter Deployment Model	17
3.4	Test Scenario 4: Policy Data Traffic Deployment Model	25
3.4.1	Policy Data Large Profile 10K Mix Traffic with 3K Notifications	25
3.4.2	Policy Data: 10K TPS Signaling Traffic	27
3.4.3	Policy Data: Performance 17.2K N36 and 6.56K Notifications Policy	35

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 and Terminologies

Acronym	Description
AMF	Access and Mobility Management Function
AUSF	Authentication Server Function
CNE	Oracle Communications Cloud Native Core, Cloud Native Environment
CPU	Central Processing Unit
EIC	Equipment Identity Check
EIR	Equipment Identity Repository
GPSI	Generic Public Subscription Identifier
HTTP	Hypertext Transfer Protocol
MPS	Messages Per Second
NF	Network Function
NRF	Oracle Communications Cloud Native Core, Network Repository Function
PVC	Persistent Volume Claim
RAM	Random Access Memory
SLF	Subscriber Location Function
SOAP	Simple Object Access Protocol
SUPI	Subscription Permanent Identifier
TPS	Transactions Per Second
UDM	Unified Data Management
UDR	Oracle Communications Cloud Native Core, Unified Data Repository
vCNE	Virtual Cloud Native Environment

What's New in This Guide

This section introduces the documentation updates for Release 23.4.x.

Release 23.4.1 - F90996-03, September 2025

- Updated the Number of Pods for the APP SQL node in the [SLF Call Model: 24.2K TPS for Performance-Medium Resource Profile for SLF Lookup](#) section.
- Added nudr-dbc-r-auditor-service microservices values in the [SLF Call Model: 24.2K TPS for Performance-Medium Resource Profile for SLF Lookup](#) section.

Release 23.4.1 - F90996-02, March 2025

- Added the [SLF Call Model: 50K lookup TPS + 1.26K Provisioning TPS](#) section.

Release 23.4.0 - F90996-01, January 2024

- Updated cnDBTier and UDR resources data. For more information, see [Policy Data: 10K TPS Signaling Traffic](#).
- Updated cnDBTier, SLF, and Provision Gateway resources data. For more information, see [SLF Call Model: 24.2K TPS for Performance-Medium Resource Profile for SLF Lookup](#).

1

Introduction

Oracle Communications Cloud Native Core Unified Data Repository (UDR) is a key component of the 5G Service Based Architecture. It is implemented as a cloud native function and offers a unified database for storing application, subscription, authentication, service authorization, policy data, session binding, and application state information. For more information about UDR architecture, see *Oracle Communications Cloud Native Core, Unified Data Repository User Guide*.

1.1 Purpose and Scope

This document is designed to measure the performance and capacity of UDR, UDR microservices, SLF, EIR, and deployment environment setup software such as Cloud Native Environment (CNE) and cnDBTier.

It is recommended that UDR is run through a benchmark on the target cloud native infrastructure to determine the capacity and performance in the target infrastructure. This information can be used to adjust the initial deployment resources and to predict resource requirements when UDR is scaled up.

1.2 References

- *Oracle Communications Cloud Native Core, Unified Data Repository Installation, Upgrade, and Fault Recovery Guide*
- *Oracle Communications Cloud Native Core, Unified Data Repository User Guide*
- *Oracle Communications Cloud Native Core, Cloud Native Environment Installation, Upgrade, and Fault Recovery Guide*
- *Oracle Communications Cloud Native Core, cnDBTier Installation, Upgrade, and Fault Recovery Guide*

2

Deployment Environment

This section provides information about the cloud native infrastructure used for UDR, SLF, and EIR benchmarking.

2.1 Deployed Components

Deployment Platform

Oracle Communications Cloud Native Environment (OCCNE) 1.10 and BareMetal is used for performing benchmark tests.

Observability Services

The following table lists services that are used for UDR benchmark tests. These services fetch UDR metrics, alerts, logs, and traces. For more information about mentioned software, see *Oracle Communications Cloud Native Core, Unified Data Repository Installation, Upgrade, and Fault Recovery Guide*

Table 2-1 Observability Services

Service
Opensearch
Fluentd
Kibana
Prometheus
Grafana
Jaeger

Cloud Native Orchestrator

Kubernetes is used for managing application pods across the cluster.

cnDBTier

cnDBTier is used for performing benchmark tests.

For more information about above mentioned components, see *Oracle Communications Cloud Native Core, Unified Data Repository Installation, Upgrade, and Fault Recovery Guide*.

2.2 Deployment Resources

The performance and capacity of UDR can vary based on the chosen environment and how UDR is deployed. This section provides information about CNE resources used to perform benchmark tests.

2.2.1 Resource Requirements for CNE Observability Services

The following table provides information about number of pods required by each CNE service.

Table 2-2 Resource Requirements for CNE Observability Services

Service Name	Number of Pods
Prometheus Server	1
Prometheus-pushgateway	1
Alert Manager	2
Fluentd	1 per Worker node
Prom-node-exporter	1 per Worker node
MetalLB speaker	1 per Worker node
Opensearch Data/Master	3/3
Opensearch Client	1
Grafana	1
Kibana	1
kube-state-metrics	1
jaeger-agent	1 per Worker node
jaeger-collector	1
jaeger-query	1
rook-ceph-osd	1 for each raw disk available to OS on all Worker nodes
rook-ceph-mgr	1
rook-ceph-mon	3
rook-ceph-osd	1

3

UDR Benchmark Testing

This chapter describes UDR, SLF, and EIR test scenarios.

3.1 Test Scenario 1: SLF Call Deployment Model

This section provides information about SLF call deployment model test scenarios.

3.1.1 SLF Call Model: 24.2K TPS for Performance-Medium Resource Profile for SLF Lookup

This test scenario describes performance and capacity of SLF functionality offered by UDR and provides the benchmarking results for various deployment sizes.

The following features are enabled for this testcase:

- LCI and OCI handling
- User Agent in Egress Traffic
- Subscriber Activity Logging - Enable and configure 100 keys
- Overload and Rate Limiting

You can perform benchmark tests on UDR/SLF for compute and storage resources by considering the following conditions:

- Signaling (SLF Look Up): 24.2K TPS
- Provisioning: 1260 TPS
- Total Subscribers: 37M
- Profile Size: 450 bytes
- Average HTTP Provisioning Request Packet Size: 350
- Average HTTP Provisioning Response Packet Size: 250

The following table describes the benchmarking parameters and their values:

Table 3-1 Traffic Model Details

Request Type	Details	Provisioning %	TPS
Lookup 24.2k	SLF Lookup GET Requests	-	24.2K
Provisioning (1.26K using Provgw one site)	CREATE	10%	126
	DELETE	10%	126
	UPDATE	40%	504
	GET	40%	504

Note

- To run this model, one UDR site is brought down and 24.2K look up traffic and 1.26K provisioning traffic are run from one site.
- The values provided is for single site deployment

The following table describes the testcase parameters and their values:

Table 3-2 Testcase Parameters

Input Parameter Details	Configuration Values
UDR Version Tag	23.4.0
Target TPS	24.2K Lookup + 1.26K Provisioning
Traffic Profile	SLF 24.2K Profile
Notification Rate	OFF
UDR Response Timeout	5s
Client Timeout	30s
Signaling Requests Latency Recorded on Client	20ms
Provisioning Requests Latency Recorded on Client	45ms

The following table describes consolidated resource requirement and their utilization:

Table 3-3 Consolidated Resource Requirement

Resource	Site	CPU	Memory
cnDBTier	cnDBTier1	90	270
	cnDBTier2	90	270
SLF	Site1	186	118
	Site2	186	118
ProvGw	Site1	32	30
Buffer	-	50	50
Total Includes Site1 and Site2 resources	-	634	856

Note

All values are inclusive of ASM sidecar.

The following table describes cnDBTier1 resources and their utilization:

Note

- The same resources and usage are application for cnDBTier2
- For cnDBTier, you must use ocudr_slf_37msub_dbtier and ocudr_udr_10msub_dbtier custom value files for SLF and UDR respectively. For more information, see *Cloud Native Core, Unified Data Repository Installation, Upgrade, and Fault Recovery Guide*.

Table 3-4 cnDBTier1 Resources and Usage

Microservice Name	Container Name	Number of Pods	CPU Allocation Per Pod	Memory Allocation Per Pod	Total Resources
Management node (ndbmgmd)	mysqlIndbcluster	2	2 CPUs	9 GB	6 CPUs 26 GB
	istio-proxy		1 CPUs	4 GB	
Data node (ndbmtd)	mysqlIndbcluster	4	4 CPUs	33 GB	28 CPUs 156 GB
	istio-proxy		2 CPUs	4 GB	
	db-backup-executor-svc		1 CPU	2GB	
APP SQL node (ndbappmysqld)	mysqlIndbcluster	5	4 CPUs	2 GB	35 CPUs 30 GB
	istio-proxy		3 CPUs	4 GB	
SQL node (Used for Replication) (ndbmysqld)	mysqlIndbcluster	2	4 CPUs	16 GB	13 CPUs 41 GB
	istio-proxy		2 CPUs	4 GB	
	init-sidecar		100m CPU	256 MB	
DB Monitor Service (db-monitor-svc)	db-monitor-svc	1	200m CPUs	500 MB	3 CPUs 2 GB
	istio-proxy		1 CPUs	1 GB	
DB Backup Manager Service (backup-manager-svc)	backup-manager-svc	1	100m CPU	128 MB	2 CPUs 2 GB
	istio-proxy		1 CPUs	1 GB	
Replication Service (db-replication-svc)	db-replication-svc	1	2 CPU	12 GB	3 CPUs 13 GB
	istio-proxy		200m CPU	500MB	

Additional configurations in cnDBTier custom value file are shown below:

Table 3-5 cnDBTier Configuration Changes

cnDBTier Configurations
<pre> ndb: annotations: - sidecar.istio.io/inject: "true" - proxy.istio.io/config: "{concurrency: 8}" - sidecar.istio.io/proxyCPU: "2000m" - sidecar.istio.io/proxyCPULimit: "2000m" - sidecar.istio.io/proxyMemory: "4Gi" - sidecar.istio.io/proxyMemoryLimit: "4Gi" mgm: annotations: - sidecar.istio.io/inject: "true" - proxy.istio.io/config: "{concurrency: 8}" - sidecar.istio.io/proxyCPU: "1000m" - sidecar.istio.io/proxyCPULimit: "1000m" - sidecar.istio.io/proxyMemory: "4Gi" - sidecar.istio.io/proxyMemoryLimit: "4Gi" api: annotations: - sidecar.istio.io/inject: "true" - proxy.istio.io/config: "{concurrency: 8}" - sidecar.istio.io/proxyCPU: "2000m" - sidecar.istio.io/proxyCPULimit: "2000m" - sidecar.istio.io/proxyMemory: "4Gi" - sidecar.istio.io/proxyMemoryLimit: "4Gi" ndbapp: annotations: - sidecar.istio.io/inject: "true" - proxy.istio.io/config: "{concurrency: 8}" - sidecar.istio.io/proxyCPU: "3000m" - sidecar.istio.io/proxyCPULimit: "3000m" - sidecar.istio.io/proxyMemory: "4Gi" - sidecar.istio.io/proxyMemoryLimit: "4Gi" </pre>

The following table describes SLF resources and their utilization for Site1 (Lookup Latency: 40ms):

Table 3-6 SLF Resources and Usage for Site1

Microservice name	Container name	Number of Pods	CPU Allocation Per Pod	Memory Allocation Per Pod	Total Resources
Ingress-gateway-sig	ingressgateway-sig	8	6 CPUs	4 GB	80 CPUs 40 GB Memory
	istio-proxy		4 CPUs	1 GB	

Table 3-6 (Cont.) SLF Resources and Usage for Site1

Microservice name	Container name	Number of Pods	CPU Allocation Per Pod	Memory Allocation Per Pod	Total Resources
Ingress-gateway-prov	ingressgateway-prov	2	4 CPUs	4 GB	12 CPUs 10 GB Memory
	istio-proxy		2 CPUs	1 GB	
Nudr-dr-service	nudr-dr-service	6	6 CPUs	4 GB	54 CPUs 30 GB Memory
	istio-proxy		3 CPUs	1 GB	
Nudr-dr-provservice	nudr-dr-provservice	2	4 CPUs	4 GB	12 CPUs 10 GB Memory
	istio-proxy		2 CPUs	1 GB	
Nudr-nrf-client-nfmanagement	nrf-client-nfmanagement	2	1 CPU	1 GB	4 CPUs 4 GB Memory
	istio-proxy		1 CPUs	1 GB	
Nudr-egress-gateway	egressgateway	2	1 CPUs	1 GB	4 CPUs 4 GB Memory
	istio-proxy		1 CPUs	1 GB	
Nudr-config	nudr-config	1	2 CPUs	2 GB	3 CPUs 3 GB Memory
	istio-proxy		1 CPUs	1 GB	
Nudr-config-server	nudr-config-server	1	2 CPUs	2 GB	3 CPUs 3 GB Memory
	istio-proxy		1 CPUs	1 GB	
alternate-route	alternate-route	2	1 CPUs	1 GB	4 CPUs 4 GB Memory
	istio-proxy		1 CPUs	1 GB	
app-info	app-info	2	1 CPUs	1 GB	4 CPUs 4 GB Memory
	istio-proxy		1 CPUs	1 GB	
perf-info	perf-info	2	1 CPUs	1 GB	4 CPUs 4 GB Memory
	istio-proxy		1 CPUs	1 GB	
nudr-dbc-auditor-service	nudr-dbc-auditor-service	1	1 CPUs	1 GB	2 CPUs 2 GB Memory
	istio-proxy		1 CPUs	1 GB	

Note

The same resources and usage are used for Site2.

The following table describes provision gateway resources and their utilization (Provisioning Latency: 45ms):

Table 3-7 Provision Gateway Resources and their utilization

Microservice name	Container name	Number of Pods	CPU Allocation Per Pod	Memory Allocation Per Pod	Total Resources
provgw-ingress-gateway	ingressgateway	2	2 CPUs	2 GB	6 CPUs 6 GB Memory
	istio-proxy		1 CPUs	1 GB	

Table 3-7 (Cont.) Provision Gateway Resources and their utilization

Microservice name	Container name	Number of Pods	CPU Allocation Per Pod	Memory Allocation Per Pod	Total Resources
provgw-egress-gateway	egressgateway	2	3 CPUs	2 GB	6 CPUs
	istio-proxy		1 CPUs	1 GB	6 GB Memory
provgw-service	provgw-service	2	3 CPUs	2 GB	8 CPUs
	istio-proxy		1 CPUs	1 GB	6 GB Memory
provgw-config	provgw-config	2	2 CPUs	2 GB	6 CPUs
	istio-proxy		1 CPUs	1 GB	6 GB Memory
provgw-config-server	provgw-config-server	2	2 CPUs	2 GB	6 CPUs
	istio-proxy		1 CPUs	1 GB	6 GB Memory

Additional configurations in custom value file are shown below:

Table 3-8 SLF and Provgw Configuration Changes

SLF Configurations	Provgw Configurations
nudur-dr-service: deployment: customExtension: labels: {} annotations: sidecar.istio.io/proxyCPU: "2000m" sidecar.istio.io/ proxyCPULimit: "2000m" sidecar.istio.io/ proxyMemory: "1Gi" sidecar.istio.io/ proxyMemoryLimit: "1Gi" proxy.istio.io/config: "{concurrency: 8}" # MySQL connection pool configurations hikari: poolsize: "50" connectionTimeout: "60000" minimumIdleConnections: "30" idleTimeout: "420" nudur-dr-prov-service: deployment: customExtension: labels: {} annotations: sidecar.istio.io/proxyCPU: "2000m" sidecar.istio.io/ proxyCPULimit: "2000m" sidecar.istio.io/ proxyMemory: "1Gi" sidecar.istio.io/ proxyMemoryLimit: "1Gi" proxy.istio.io/config: "{concurrency: 8}" # MySQL connection pool configurations hikari: poolsize: "25" connectionTimeout: "60000" minimumIdleConnections: "25" idleTimeout: "420" ingressgateway-sig: deployment: # Microservice specific	prov-gw-service: deployment: customExtension: labels: {} annotations: sidecar.istio.io/proxyCPU: "1000m" sidecar.istio.io/ proxyCPULimit: "1000m" sidecar.istio.io/ proxyMemory: "1Gi" sidecar.istio.io/ proxyMemoryLimit: "1Gi" proxy.istio.io/config: "{concurrency: 8}" prov-ingressgateway: deployment: customExtension: labels: {} annotations: sidecar.istio.io/proxyCPU: "1000m" sidecar.istio.io/ proxyCPULimit: "1000m" sidecar.istio.io/ proxyMemory: "1Gi" sidecar.istio.io/ proxyMemoryLimit: "1Gi" proxy.istio.io/config: "{concurrency: 8}" # Jetty client settings maxConcurrentPushedStreams: 6000 maxRequestsQueuedPerDestination: 5000 maxConnectionsPerIp: 10 maxConnectionsPerDestination: 10 connectionTimeout: 250000 #(ms) requestTimeout: 21000 #(ms) prov-egressgateway: deployment: customExtension: labels: {} annotations: sidecar.istio.io/proxyCPU: "1000m" sidecar.istio.io/ proxyCPULimit: "1000m"

Table 3-8 (Cont.) SLF and Provgw Configuration Changes

SLF Configurations	Provgw Configurations
<pre> annotation for deployment customExtension: labels: {} annotations: sidecar.istio.io/proxyCPU: "3000m" sidecar.istio.io/ proxyCPULimit: "3000m" sidecar.istio.io/ proxyMemory: "1Gi" sidecar.istio.io/ proxyMemoryLimit: "1Gi" proxy.istio.io/config: "{concurrency: 8}" # Jetty client settings maxConcurrentPushedStreams: 6000 maxRequestsQueuedPerDestination: 5000 maxConnectionsPerIp: 10 maxConnectionsPerDestination: 10 connectionTimeout: 10000 #(ms) requestTimeout: 5000 #(ms) ingressgateway-prov: deployment: # Microservice specific annotation for deployment customExtension: labels: {} annotations: sidecar.istio.io/proxyCPU: "2000m" sidecar.istio.io/ proxyCPULimit: "2000m" sidecar.istio.io/ proxyMemory: "1Gi" sidecar.istio.io/ proxyMemoryLimit: "1Gi" proxy.istio.io/config: "{concurrency: 8}" # Jetty client settings maxConcurrentPushedStreams: 6000 maxRequestsQueuedPerDestination: 5000 maxConnectionsPerIp: 10 maxConnectionsPerDestination: 10 connectionTimeout: 10000 #(ms) requestTimeout: 5000 #(ms) </pre>	<pre> sidecar.istio.io/ proxyMemory: "1Gi" sidecar.istio.io/ proxyMemoryLimit: "1Gi" proxy.istio.io/config: "{concurrency: 8}" #jetty client configuration maxConcurrentPushedStreams: 6000 maxRequestsQueuedPerDestination: 5000 maxConnectionsPerDestination: 10 maxConnectionsPerIp: 10 connectionTimeout: 10000 #(ms) requestTimeout: 5000 #(ms) jettyIdleTimeout: 0 </pre>

The following table provides observation data for the performance test that can be used for the benchmark testing to scale up SLF performance:

Table 3-9 Result and Observation

Parameter	Values
Test Duration	48hr
TPS Achieved	24.2K SLF Lookup + 1.26K Provisioning
Success Rate	100%
Average SLF processing time (Request and Response)	40ms

3.1.2 SLF Call Model: 50K lookup TPS + 1.26K Provisioning TPS

This test scenario describes performance and capacity of SLF functionality offered by UDR and provides the benchmarking results for various deployment sizes.

The following features are enabled for this testcase:

- Overload Handling
- Controlled Shutdown of an Instance
- Network Function Scoring for a Site
- Support for User-Agent Header
- Support for Default Group ID in SLF

You can perform benchmark tests on SLF for compute and storage resources by considering the following conditions:

- Signaling (SLF Look Up): 50K TPS
- Provisioning: 1260 TPS
- Total Subscribers: 64M
- Profile Size: 450 bytes

The following table describes the benchmarking parameters and their values:

Table 3-10 Traffic Model Details

Request Type	Details	TPS
Lookup 50k	SLF Lookup GET Requests	50K
Provisioning (1.26K using Provgw one site)	CREATE	126
	DELETE	126
	UPDATE	504
	GET	504

The following table describes the testcase parameters and their values:

Table 3-11 Testcase Parameters

Input Parameter Details	Configuration Values
UDR Version Tag	23.4.1

Table 3-11 (Cont.) Testcase Parameters

Input Parameter Details	Configuration Values
Target TPS	50K Lookup + 1.26K Provisioning
Traffic Profile	SLF 50K Profile
Notification Rate	OFF
UDR Response Timeout	900ms
Client Timeout	30s
Signaling Requests Latency Recorded on Client	19ms
Provisioning Requests Latency Recorded on Client	42ms

The following table describes consolidated resource requirement and their utilization:

Table 3-12 Consolidated Resource Requirement

Resource	CPU	Memory	Ephemeral Storage	PVC
cnDBTier	130	436 GB	20 GB	946 GB
SLF	384	225 GB	53 GB	NA
ProvGw	45	39 GB	13 GB	NA
Buffer	50	50 GB	50 GB	50 GB
Total	609	750 GB	136 GB	996 GB

Note

All values are inclusive of ASM sidecar.

The following table describes cnDBTier1 resources and their utilization:

Note

- The same resources and usage are applicable for cnDBTier2.
- For cnDBTier, you must use the applicable custom values files shared as part of the custom templates. For more information, see *Cloud Native Core, Unified Data Repository Installation, Upgrade, and Fault Recovery Guide*.

Table 3-13 cnDBTier Resources and Usage

Microservice Name	Container Name	Number of Pods	CPU Allocation Per Pod (cnDBtier1)	Memory Allocation Per Pod (cnDBtier1)	Ephemeral Storage Per Pod	PVC Allocation Per Pod	Total Resources (cnDBtier)
Management node (ndbmgmd)	mysqlndbcluster	2	2 CPUs	12 GB	1 GB	16 GB	6 CPUs 26 GB

Table 3-13 (Cont.) cnDBTier Resources and Usage

Microservice Name	Container Name	Number of Pods	CPU Allocation Per Pod (cnDBTier1)	Memory Allocation Per Pod (cnDBTier1)	Ephemeral Storage Per Pod	PVC Allocation Per Pod	Total Resources (cnDBTier)
	istio-proxy		1 CPUs	1 GB			Ephemeral Storage: 2 GB PVC Allocation: 32 GB
Data node (ndbmt)	mysqlndbcluster	6	4 CPUs	50 GB	1 GB	65 GB (Backup: 63 GB)	42 CPUs 324 GB Ephemeral Storage: 6 GB PVC Allocation: 768 GB
	istio-proxy		2 CPUs	2 GB			
	db-backup-executor-svc		1 CPU	2 GB			
APP SQL node (ndbappmysqld)	mysqlndbcluster	7	6 CPUs	2 GB	1 GB	10 GB	63 CPUs 28 GB Ephemeral Storage: 7 GB PVC Allocation: 70 GB
	istio-proxy		3 CPUs	2 GB			
SQL node (Used for Replication) (ndbmysqld)	mysqlndbcluster	2	4 CPUs	16 GB	1 GB	16 GB	13 CPUs 41 GB Ephemeral Storage: 2 GB PVC Allocation: 32 GB
	istio-proxy		2 CPUs	4 GB			
	init-sidecar		100m CPU	256 MB			
DB Monitor Service (db-monitor-svc)	db-monitor-svc	1	200m CPUs	500 MB	1 GB	NA	3 CPUs 2 GB Ephemeral Storage: 1 GB
	istio-proxy		1 CPUs	1 GB			
DB Backup Manager Service (backup-manager-svc)	backup-manager-svc	1	100m CPU	128 MB	1 GB	NA	2 CPUs 2 GB Ephemeral Storage: 1 GB
	istio-proxy		1 CPUs	1 GB			
Replication Service (db-replication-svc)	db-replication-svc	1	2 CPU	12 GB	1 GB	42 GB	3 CPUs 13 GB Ephemeral Storage: 1 GB PVC Allocation: 42 GB
	istio-proxy		200m CPU	500 MB			

The following table describes SLF resources and their utilization for Site1 (Lookup Latency: 19ms):

Table 3-14 SLF Resources and Usage

Microservice name	Container name	Number of Pods	CPU Allocation Per Pod	Memory Allocation Per Pod	Ephemeral Storage Per Pod	Total Resources
Ingress-gateway-sig	ingressgateway-sig	19	6 CPUs	4 GB	1 GB	190 CPUs 95 GB Memory Ephemeral Storage: 19 GB
	istio-proxy		4 CPUs	1 GB		
Ingress-gateway-prov	ingressgateway-prov	2	4 CPUs	4 GB	1 GB	12 CPUs 10 GB Ephemeral Storage: 2 GB
	istio-proxy		2 CPUs	1 GB		
Nudr-dr-service	nudr-dr-service	15	6 CPUs	4 GB	1 GB	135 CPUs 75 GB Ephemeral Storage: 15 GB
	istio-proxy		3 CPUs	1 GB		
Nudr-dr-provservice	nudr-dr-provservice	2	4 CPUs	4 GB	1 GB	12 CPUs 10 GB Ephemeral Storage: 2 GB
	istio-proxy		2 CPUs	1 GB		
Nudr-nrf-client-nfmanagement	nrf-client-nfmanagement	2	1 CPU	1 GB	1 GB	4 CPUs 4 GB Ephemeral Storage: 2 GB
	istio-proxy		1 CPUs	1 GB		
Nudr-egress-gateway	egressgateway	2	1 CPUs	1 GB	1 GB	4 CPUs 4 GB Ephemeral Storage: 2 GB
	istio-proxy		1 CPUs	1 GB		
Nudr-config	nudr-config	2	2 CPUs	2 GB	1 GB	6 CPUs 6 GB Ephemeral Storage: 2 GB
	istio-proxy		1 CPUs	1 GB		
Nudr-config-server	nudr-config-server	2	2 CPUs	2 GB	1 GB	6 CPUs 6 GB Ephemeral Storage: 2 GB
	istio-proxy		1 CPUs	1 GB		
alternate-route	alternate-route	2	1 CPUs	1 GB	1 GB	4 CPUs 4 GB

Table 3-14 (Cont.) SLF Resources and Usage

Microservice name	Container name	Number of Pods	CPU Allocation Per Pod	Memory Allocation Per Pod	Ephemeral Storage Per Pod	Total Resources
	istio-proxy		1 CPUs	1 GB		Ephemeral Storage: 2 GB
app-info	app-info	2	1 CPUs	1 GB	1 GB	4 CPUs 4 GB Ephemeral Storage: 2 GB
	istio-proxy		1 CPUs	1 GB		
perf-info	perf-info	2	1 CPUs	1 GB	1 GB	4 CPUs 4 GB Ephemeral Storage: 2 GB
	istio-proxy		1 CPUs	1 GB		
Nudr-dbcrauditor	nudr-dbcrauditor-service	1	2 CPUs	2 GB	1 GB	3 CPUs 3 GB Ephemeral Storage: 1 GB
	istio-proxy		1 CPUs	1 GB		

Note

The same resources and usage are used for Site2.

The following table describes provision gateway resources and their utilization (Provisioning Latency: 42ms):

Table 3-15 Provision Gateway Resources and their utilization

Microservice name	Container name	Number of Pods	CPU Allocation Per Pod	Memory Allocation Per Pod	Ephemeral Storage Per Pod	Total Resources
provgw-ingress-gateway	ingressgateway	2	2 CPUs	2 GB	1 GB	6 CPUs 6 GB Memory Ephemeral Storage: 2 GB
	istio-proxy		1 CPUs	1 GB		
provgw-egress-gateway	egressgateway	3	3 CPUs	2 GB	1 GB	12 CPUs 9 GB Memory Ephemeral Storage: 3 GB
	istio-proxy		1 CPUs	1 GB		
provgw-service	provgw-service	3	3 CPUs	2 GB	1 GB	12 CPUs

Table 3-15 (Cont.) Provision Gateway Resources and their utilization

Microservice name	Container name	Number of Pods	CPU Allocation Per Pod	Memory Allocation Per Pod	Ephemeral Storage Per Pod	Total Resources
	istio-proxy		1 CPUs	1 GB		9 GB Memory Ephemeral Storage: 3 GB
provgw-config	provgw-config	2	2 CPUs	2 GB	1 GB	6 CPUs 6 GB Memory Ephemeral Storage: 2 GB
	istio-proxy		1 CPUs	1 GB		
provgw-config-server	provgw-config-server	2	2 CPUs	2 GB	1 GB	6 CPUs 6 GB Memory Ephemeral Storage: 2 GB
	istio-proxy		1 CPUs	1 GB		
provgw-auditor-service	auditor-service	1	2 CPUs	2 GB	1 GB	3 CPUs 3 GB Memory Ephemeral Storage: 1 GB
	istio-proxy		1 CPU	1 GB		

The following table provides observation data for the performance test that can be used for the benchmark testing to scale up SLF performance:

Table 3-16 Result and Observation

Parameter	Values
Test Duration	24hr
TPS Achieved	50K SLF Lookup + 1.26K Provisioning
Success Rate	100%
Average SLF processing time for signaling requests	19ms
Average SLF processing time for provisioning requests	42ms

3.2 Test Scenario 2: EIR Deployment Model

Performance Requirement - 300K subscriber DB size with 10K EIR lookup TPS

This test scenario describes performance and capacity improvements of EIR functionality offered by UDR and provides the benchmarking results for various deployment sizes.


The following features are disabled for this testcase:

- TLS
- OAuth2.0
- Default Response set to EQUIPMENT_UNKNOWN
- Header Validations like XFCC, server header, and user agent header

UDR/EIR is benchmarked for compute and storage resources under following conditions:

- Signaling (EIR Look Up): 10K TPS
- Provisioning: 2K TPS
- Total Subscribers: 300K
- Profile Size: 130 bytes
- Average HTTP Provisioning Request Packet Size: NA
- Average HTTP Provisioning Response Packet Size: NA

Figure 3-1 EIR Deployment Model



The following table describes the benchmarking parameters and their values:

Table 3-17 Traffic Model Details

Request Type	Details	TPS
Lookup 10k	EIR EIC	10k

The following table describes the testcase parameters and their values:

Table 3-18 Testcase Parameters

Input Parameter Details	Configuration Values
UDR Version Tag	22.3.0
Target TPS	10k Lookup
Traffic Profile	10k EIR EIC
Notification Rate	OFF
EIR Response Timeout	5s

Table 3-18 (Cont.) Testcase Parameters

Input Parameter Details	Configuration Values
Client Timeout	10s
Signaling Requests Latency Recorded on Client	NA
Provisioning Requests Latency Recorded on Client	NA

The following table describes the consolidated resource requirements and their utilization:

Table 3-19 Consolidated Resource Requirement

Resource	CPUs	Memory
EIR	32	30 GB
cnDBTier	177	616 GB
Total	600	903 GB

The following table describes cnDBTier resources and their utilization:

Table 3-20 cnDBTier Resources and their Utilization

Micro service name	Container name	Number of Pods	CPU Allocation Per Pod	Memory Allocation Per Pod	Total Resources
Management node	mysqlIndbcluster	2	4 CPUs	4 GB	4 CPU 8 GB Memory
Data node	mysqlIndbcluster	4	16 CPUs	32 GB	64 CPU 128 GB Memory
APP SQL node	mysqlIndbcluster	3	16 CPUs	32 GB	48 CPU 96 GB Memory
SQL node (Used for Replication)	mysqlIndbcluster	2	2 CPUs	4 GB	4 CPU 8 GB Memory
DB Monitor Service	db-monitor-svc	1	500m CPUs	500 MB	1 CPU 1 GB Memory
DB Backup Manager Service	replication-svc	1	250m CPUs	320 MB	1 CPU 1 GB Memory

The following table describes EIR resources and their utilization:

Table 3-21 EIR Resources and their Utilization (Lookup Latency: 16.9ms) without Aspen Service Mesh (ASM) Enabled

Micro service name	Container name	Number of Pods	CPU Allocation Per Pod	Memory Allocation Per Pod	Total Resources
Ingress-gateway-sig	Ingress-gateway-sig	5	6 CPUs	4 GB	30 CPUs 20 GB Memory

Table 3-21 (Cont.) EIR Resources and their Utilization (Lookup Latency: 16.9ms) without Aspen Service Mesh (ASM) Enabled

Micro service name	Container name	Number of Pods	CPU Allocation Per Pod	Memory Allocation Per Pod	Total Resources
Ingress-gateway-prov	Ingress-gateway-prov	2	6 CPUs	4 GB	12 CPUs 8 GB Memory
Nudr-dr-service	nudr-dr-service	6	4 CPUs	4 GB	24 CPUs 24 GB Memory
Nudr-dr-provservice	nudr-dr-provservice	2	4 CPUs	4 GB	8 CPUs 8 GB Memory
Nudr-egress-gateway	egressgateway	1	2 CPUs	2 GB	2 CPUs 2 GB Memory
Nudr-config	nudr-config	2	2 CPUs	2 GB	4 CPUs 4 GB Memory
Nudr-config-server	nudr-config-server	2	2 CPUs	2 GB	4 CPU 4 GB Memory

Note

The following table provides observation data for the performance test that can be used for the benchmark testing to scale up UDR/EIR performance:

Table 3-22 Result and Observation

Parameter	Values
Test Duration	8hr
TPS Achieved	10k
Success Rate	100%
Average EIR processing time (Request and Response)	16.9ms

3.3 Test Scenario 3: SOAP and Diameter Deployment Model

2K SOAP provisioning TPS for ProvGw for Medium profile + Diameter 25K with Large profile


The following features are disabled for this testcase:

- TLS
- OAuth2.0
- Header Validations like XFCC, server header, and user agent header

UDR is benchmarked for compute and storage resources under following conditions:

- Signaling : 10K TPS
- Provisioning: 2K TPS
- Total Subscribers: 1M - 10M range used for Diameter Sh and 1M range used for SOAP/XML
- Profile Size: 2.2KB
- Average HTTP Provisioning Request Packet Size: NA
- Average HTTP Provisioning Response Packet Size: NA

Figure 3-2 SOAP and Diameter Deployment Model



The following table describes the benchmarking parameters and their values:

Table 3-23 Traffic Model Details

Request Type	Details	TPS
Diameter SH Traffic	SH Traffic	25K
Provisioning (2K using Provgw)	SOAP Traffic	2K

Table 3-24 SOAP Traffic Model

Request Type	SOAP Traffic %
GET	33%
DELETE	11%
POST	11%
PUT	45%

Table 3-25 Diameter Traffic Model

Request Type	Diameter Traffic %
SNR	25%
PUR	50%
UDR	25%

The following table describes the benchmarking parameters and their values:

Table 3-26 Testcase Parameters

Input Parameter Details	Configuration Values
UDR Version Tag	22.2.0
Target TPS	25K + 2K
Traffic Profile	25K sh + 2K SOAP
Notification Rate	OFF
UDR Response Timeout	5s
Client timeout	10s
Signaling Requests Latency Recorded on Client	NA
Provisioning Requests Latency Recorded on Client	NA

Note

PNR scenarios are not tested because server stub is not used.

The following table describes cnDBTier resources and their utilization:

Table 3-27 cnDBTier Resources and their Utilization

Micro service name	Container name	Number of Pods	CPU Allocation Per Pod	Memory Allocation Per Pod	Total Resources
Management node	mysqlIndbcluster	3	4 CPUs	10 GB	12 CPUs 30 GB Memory
Data node	mysqlIndbcluster	4	15 CPUs	98 GB	64 CPU 408 GB Memory
	db-backup-executor-svc		100m CPU	128 MB	
APP SQL node	mysqlIndbcluster	4	16 CPUs	16 GB	64 CPUs 64 GB Memory
SQL node (Used for Replication)	mysqlIndbcluster	4	8 CPUs	16 GB	49 CPUs 81 GB Memory
DB Monitor Service	db-monitor-svc	1	200m CPUs	500 MB	3 CPUs 2 GB Memory

Table 3-27 (Cont.) cnDBTier Resources and their Utilization

Micro service name	Container name	Number of Pods	CPU Allocation Per Pod	Memory Allocation Per Pod	Total Resources
DB Backup Manager Service	replication-svc	1	200m CPU	500 MB	3 CPUs 2 GB Memory

cnDBTier Usage

Results for Kubectl top pods on cndbtier is shown below:

NAME	CPU (cores)	MEMORY (bytes)
mysql-cluster-db-backup-manager-svc-5f9956c869-44r9p	1m	41Mi
mysql-cluster-db-monitor-svc-74568b68c6-bgznr	2m	235Mi
ndbappmysqld-0	9518m	8827Mi
ndbappmysqld-1	9856m	8809Mi
ndbappmysqld-2	9874m	8811Mi
ndbappmysqld-3	8961m	8805Mi
ndbmcmd-0	8m	2064Mi
ndbmcmd-1	8m	2062Mi
ndbmcmd-2	8m	2065Mi
ndbmttd-0	5768m	92317Mi
ndbmttd-1	5526m	92312Mi
ndbmttd-2	6031m	92331Mi
ndbmttd-3	5788m	92317Mi

Results for Kubectl get hpa on cndbtier is shown below:

NAME	REFERENCE	TARGETS	MINPODS	MAXPODS	REPLICAS	AGE
ndbappmysqld	StatefulSet/ndbappmysqld	59%/80%	4	4	4	10d
ndbmcmd	StatefulSet/ndbmcmd	0%/80%	3	3	3	10d
ndbmttd	StatefulSet/ndbmttd	38%/80%	4	4	4	10d

cnDBTier Metrics:

- Data memory usage: 72GB (5.164GB used)
- DB Reads per second: 52k
- DB Writes per second: 24k

The following table describes UDR resources and their utilization:

Table 3-28 UDR Resources and their Utilization (Request Latency: 40ms)

Micro service name	Container name	Number of Pods	CPU Allocation Per Pod	Memory Allocation Per Pod	Total Resources
nudr-diameterproxy	nudr-diameterproxy	19	2.5 CPUs	4 GB	47.5 CPUs 76 GB Memory
nudr-diam-gateway	nudr-diam-gateway	3	6 CPUs	4 GB	18 CPUs 12 GB Memory
Ingress-gateway-sig	ingressgateway-sig	2	2 CPUs	2 GB	4 CPUs 4 GB Memory

Table 3-28 (Cont.) UDR Resources and their Utilization (Request Latency: 40ms)

Micro service name	Container name	Number of Pods	CPU Allocation Per Pod	Memory Allocation Per Pod	Total Resources
Ingress-gateway-prov	ingressgateway-prov	2	2 CPUs	2 GB	4 CPUs 4 GB Memory
Nudr-dr-service	nudr-drservice	2	2 CPUs	2 GB	4 CPUs 4 GB Memory
Nudr-dr-provservice	nudr-dr-provservice	2	2 CPUs	2 GB	4 CPUs 4 GB Memory
Nudr-nrf-client-nfmanagement	nrf-client-nfmanagement	2	1 CPUs	1 GB	2 CPUs 2 GB Memory
Nudr-egress-gateway	egressgateway	2	2 CPUs	2 GB	4 CPU 4 GB Memory
Nudr-config	nudr-config	2	1 CPUs	1 GB	2 CPU 2 GB Memory
Nudr-config-server	nudr-config-server	2	1 CPUs	1 GB	2 CPU 2 GB Memory
alternate-route	alternate-route	2	1 CPUs	1 GB	2 CPU 2 GB Memory
app-info	app-info	2	1 CPUs	1 GB	2 CPU 2 GB Memory
perf-info	perf-info	2	1 CPUs	1 GB	2 CPU 2 GB Memory

Resource Utilization

Diameter resource utilization is shown below:

ocudr-nudr-diam-gateway-0	2463m	1958Mi
ocudr-nudr-diam-gateway-1	2474m	1747Mi
ocudr-nudr-diam-gateway-2	2453m	1754Mi
ocudr-nudr-diameterproxy-b5c5f55b8-4vglt	1679m	984Mi
ocudr-nudr-diameterproxy-b5c5f55b8-5x94d	1688m	1095Mi
ocudr-nudr-diameterproxy-b5c5f55b8-8f7xg	1818m	1078Mi
ocudr-nudr-diameterproxy-b5c5f55b8-8kbgs	1789m	796Mi
ocudr-nudr-diameterproxy-b5c5f55b8-8wwhv	1845m	781Mi
ocudr-nudr-diameterproxy-b5c5f55b8-bwpzw	1719m	908Mi
ocudr-nudr-diameterproxy-b5c5f55b8-c8nk6	1690m	1054Mi
ocudr-nudr-diameterproxy-b5c5f55b8-cjlvq	1662m	983Mi
ocudr-nudr-diameterproxy-b5c5f55b8-dfkq4	1755m	983Mi
ocudr-nudr-diameterproxy-b5c5f55b8-fqjkw	1761m	979Mi
ocudr-nudr-diameterproxy-b5c5f55b8-gmpfn	1684m	865Mi
ocudr-nudr-diameterproxy-b5c5f55b8-hvmzk	1776m	907Mi
ocudr-nudr-diameterproxy-b5c5f55b8-q5qsp	1763m	863Mi
ocudr-nudr-diameterproxy-b5c5f55b8-qd5w2	1799m	996Mi
ocudr-nudr-diameterproxy-b5c5f55b8-qnkk5	1735m	945Mi
ocudr-nudr-diameterproxy-b5c5f55b8-rlm7k	1789m	811Mi
ocudr-nudr-diameterproxy-b5c5f55b8-tpqhf	1681m	964Mi
ocudr-nudr-diameterproxy-b5c5f55b8-vfq98	1839m	928Mi
ocudr-nudr-diameterproxy-b5c5f55b8-vtvcb	1795m	845Mi

UDR HPA resource utilization is shown below:

NAME	REFERENCE	TARGETS	MINPODS	MAXPODS	REPLICAS	AGE
ocudr-config-server-hpa	Deployment/ocudr-nudr-config-server	2%/80%	1	2	1	20h
ocudr-egressgateway-vl	Deployment/ocudr-egressgateway	0%/65%	1	4	1	20h
ocudr-ingressgateway-prov-vl	Deployment/ocudr-ingressgateway-prov	43%/65%	2	8	2	20h
ocudr-ingressgateway-sig-vl	Deployment/ocudr-ingressgateway-sig	0%/65%	1	4	1	20h
ocudr-nudr-config	Deployment/ocudr-nudr-config	0%/80%	1	1	1	20h
ocudr-nudr-diameterproxy	Deployment/ocudr-nudr-diameterproxy	70%/80%	19	19	19	20h
ocudr-nudr-dr-provservice	Deployment/ocudr-nudr-dr-provservice	70%/80%	2	2	2	20h
ocudr-nudr-drservice	Deployment/ocudr-nudr-drservice	0%/80%	1	4	1	20h
ocudr-nudr-notify-service	Deployment/ocudr-nudr-notify-service	9%/80%	1	1	1	20h
provgw-config-server-hpa	Deployment/provgw-provgw-config-server	0%/80%	1	2	1	5h24m
provgw-prov-egressgateway-vl	Deployment/provgw-prov-egressgateway	45%/80%	2	4	2	5h24m
provgw-prov-ingressgateway-vl	Deployment/provgw-prov-ingressgateway	64%/80%	2	4	3	5h24m
provgw-provgw-config	Deployment/provgw-provgw-config	0%/80%	1	1	1	5h24m
provgw-provgw-service	Deployment/provgw-provgw-service	72%/80%	4	4	4	5h24m

The following table describes provision gateway resources and their utilization:

Table 3-29 Provision Gateway Resources and their Utilization (Provisioning Request Latency: 40ms)

Micro service name	Container name	Number of Pods	CPU Allocation Per Pod	Memory Allocation Per Pod	Total Resources
provgw-ingress-gatewa	ingressgateway	3	2 CPUs	2 GB	6 CPUs 6 GB Memory
provgw-egress-gateway	egressgateway	2	2 CPUs	2 GB	4 CPUs 4 GB Memory

Table 3-29 (Cont.) Provision Gateway Resources and their Utilization (Provisioning Request Latency: 40ms)

Micro service name	Container name	Number of Pods	CPU Allocation Per Pod	Memory Allocation Per Pod	Total Resources
provgw-service	provgw-service	4	2.5 CPUs	3 GB	10 CPUs 12 GB Memory
provgw-config	provgw-config	2	1 CPUs	1 GB	2 CPUs 2 GB Memory
provgw-config-server	provgw-config-server	2	1 CPUs	1 GB	2 CPUs 2 GB Memory

Provisioning Gateway resource utilization is shown below:

ocudr-ingressgateway-prov-bc567cb6d-jsr7d	874m	666Mi
ocudr-ingressgateway-prov-bc567cb6d-v8lfp	862m	687Mi
ocudr-nudr-dr-provservice-7c855bb68-4bqns	1445m	815Mi
ocudr-nudr-dr-provservice-7c855bb68-gchlt	1348m	828Mi
provgw-prov-egressgateway-745c9d5d7d-bz7ct	921m	738Mi
provgw-prov-egressgateway-745c9d5d7d-xvj2k	904m	710Mi
provgw-prov-ingressgateway-7bd88db949-8qmps	1850m	1263Mi
provgw-prov-ingressgateway-7bd88db949-g7b25	4m	583Mi
provgw-prov-ingressgateway-7bd88db949-wdg4k	2003m	1281Mi
provgw-provgw-config-6fcc86cd78-jlfr	4m	536Mi
provgw-provgw-config-server-9dff7cf4f-6qwkb	5m	374Mi
provgw-provgw-service-c54d94bcb-q84gl	159m	582Mi
provgw-provgw-service-c54d94bcb-s8kpz	2501m	742Mi
provgw-provgw-service-c54d94bcb-sqwwr	2427m	1191Mi
provgw-provgw-service-c54d94bcb-zttr9	2103m	1082Mi

Resources calculation for UDR, Provisioning Gateway, and cnDbTier are shown below:

Table 3-30 cnUDR and ProvGw Resources Calculation

Resources	cnUDR			ProvGw		
	Core services used for traffic runs (Nudr-diamgw, Nudr-diamproxy, Nudr-ingressgateway-prov and Nudr-dr-prov) at 70% usage	Other Microservices	Total	Core services used for traffic runs (ProvGw-ingressgateway, ProvGw-provgw service and ProvGw-egressgateway) at 70% usage	Other Microservice	Total
CPU	73.5	24	97.5	20	4	24
Memory in GB	96	24	120	22	4	26
Disk Volume (Ephemeral storage) in GB	26	16	42	9	4	13

Table 3-31 cnDbTier Resources Calculation

Resources	cnDbTier					
	SQL nodes (at actual usage)	SQL Nodes (Overhead/ Buffer resources at 20%)	Data nodes (at actual usage)	Data nodes (Overhead/ Buffer resources at 10%)	MGM nodes and other resources (Default resources)	Total
CPU	76	16	23.2	5	18	138.5
Memory in GB	70.4	14	368	36	34	522
Disk Volume (Ephemeral storage) in GB	8	NA	960 (ndbdisksize = 240*4)	NA	20	988

Table 3-32 Total Resources Calculation

Resources	Total
CPU	260
Memory in GB	668 GB
Disk Volume (Ephemeral storage) in GB	104 GB

The following table provides observation data for the performance test that can be used for the benchmark testing to scale up UDR performance:

Table 3-33 Result and Observation

Parameter	Values
Test Duration	18hr

Table 3-33 (Cont.) Result and Observation

Parameter	Values
TPS Achieved	10K
Success Rate	100%
Average UDR processing time (Request and Response)	40ms

3.4 Test Scenario 4: Policy Data Traffic Deployment Model

This section provides information about policy data traffic deployment model test scenarios.


3.4.1 Policy Data Large Profile 10K Mix Traffic with 3K Notifications

The following features are disabled for this testcase:

- TLS
- OAuth2.0
- Header Validations like XFCC, server header, and user agent header

You can perform benchmark tests on UDR for compute and storage resources by considering the following conditions:

- Signaling : 10k (includes subscriptions)
- Provisioning: NA
- Total Subscribers: 4M
- Profile Size: NA
- Average HTTP Provisioning Request Packet Size: NA
- Average HTTP Provisioning Response Packet Size: NA

Figure 3-3 Policy Data Traffic Deployment Model

The following table describes the benchmarking parameters and their values:

Table 3-34 Traffic Model Details

Request Type	Details	TPS
Notifications 3k	Notifications Requests triggered from UDR	3k
Signaling (10k mix traffic)	GET	18%
	PUT	37%
	PATCH	15%
	DELETE	15%
	POST Subscription	7.5%
	DELETE Subscription	7.5%

The following table describes the testcase parameters and their values:

Table 3-35 Testcase Parameters

Input Parameter Details	Configuration Values
UDR Version Tag	1.15.0
Target TPS	10k SignalingTPS + 3k Notifications
Notification Rate	3k
UDR Response Timeout	45ms
Client Timeout	10s
Signaling Requests Latency Recorded on Client	NA
Provisioning Requests Latency Recorded on Client	NA

Table 3-36 Average Deployment Size to Achieve Higher TPS

Subscriber profile (number of subscribers on DB)	Microservice name	TPS rate	Number of pods	CPU allocation per pod	Memory allocation per pod	Average CPU used	Average memory used	Latency values on httpGo tool
4M	ocudr-ingress-gateway	10000	5	6	5Gi	2.8	1352Mi	45ms
	nudr-dr-service	10000	10	5	4Gi	4.15	1512Mi	NA
	nudr-notify-service	3.0 K	7	4	4Gi	4.5	1050Mi	NA
	ocudr-egress-gateway	3.0 K	4	4	3Gi	2.7	991Mi	NA

The following table describes cnDBTier pod details:

Table 3-37 cnDBTier Pod Details

VM name	vCPU	RAM	Storage
ndbmysqld-0	16	64 GB	90
ndbmysqld-1	16	64 GB	90
ndbmysqld-2	16	64 GB	90
ndbmysqld-3	16	64 GB	90
ndbmcmd-0	8	8 GB	50
ndbmtd-0	16	64 GB	190
ndbmtd-1	16	64 GB	190
ndbmtd-2	16	64 GB	190
ndbmtd-3	16	64 GB	190

Table 3-38 TPS Rate

TPS	TPS Rate
TPS of DB writes	3.91K x 4
TPS of DB reads	15.4K x 4
Total DB reads	29.42 Mil x 4
Total DB writes	12.31 Mil x 4

The following table provides observation data for the performance test that can be used for the benchmark testing to scale up UDR performance:

Table 3-39 Result and Observation

Parameter	Values
Test Duration	4h
TPS Achieved	10k Provisioning (Includes Subscription) + 3k Notifications
Success rate	100%
Average UDR processing time (Request and Response)	45ms

3.4.2 Policy Data: 10K TPS Signaling Traffic

The following feature is enabled for this testcase:

- Entity Tag (ETag) is enabled

You can perform benchmark tests on UDR for compute and storage resources by considering the following conditions:

- Signaling : 10K
- Provisioning: NA
- Total Subscribers: 1M
- Profile Size: 2.5KB
- Average HTTP Provisioning Request Packet Size: NA

- Average HTTP Provisioning Response Packet Size: NA

The following table describes the benchmarking parameters and their values:

Table 3-40 Traffic Model Details

Request Type	Details	TPS
N36 traffic (100%) 10K TPS for sm-data	GET	2500
	PUT	2500
	PATCH	2500
	DELETE	2500

Note

Provisioning and Egress traffic are not included in this model.

The following table describes the test case parameters and their values:

Table 3-41 Testcase Parameters

Input Parameter Details	Configuration Values
UDR Version Tag	23.4.0
Target TPS	10k Signaling
Notification Rate	NA
UDR Response Timeout	5s
Client Timeout	30s
Signaling Requests Latency Recorded on Client	NA
Provisioning Requests Latency Recorded on Client	NA

Table 3-42 Consolidated Resource Requirement

Resource	CPU	Memory
cnDBTier	50	434
UDR	73	55
Buffer	50	50
Total	173	539

The following table describes cnDBTier resources and their utilization:

Note

For cnDBTier, you must use `ocudr_slf_37msub_dbtier` and `ocudr_udr_10msub_dbtier` custom value files for SLF and UDR respectively. For more information, see *Cloud Native Core, Unified Data Repository Installation, Upgrade, and Fault Recovery Guide*.

Table 3-43 cnDBTier Resources and their Utilization

Microservice name	Container name	Number of Pods	CPU Allocation Per Pod	Memory Allocation Per Pod	Total Resources
Management node (ndbmcmd)	mysqlndbcluster	2	2 CPUs	9 GB	4 CPUs 18 GB
Data node (ndbmt)	mysqlndbcluster	4	4 CPUs	93 GB	16 CPUs 372 GB
APP SQL node (ndbappmysqld)	mysqlndbcluster	5	4 CPUs	2 GB	20 CPUs 10 GB
SQL node (ndbmysqld, used for replication)	mysqlndbcluster	2	4 CPUs	16 GB	8 CPUs 32 GB
DB Monitor Service	db-monitor-svc	1	200m CPUs	500 MB	1 CPUs 500 MB
DB Backup Manager Service	backup-manager-svc	1	100m CPU	128 MB	1 CPUs 128 MB

Additional cnDBTier configurations in custom value file are shown as follows:

```
# values for configuration files, cnf
ndbconfigurations:
  mgm:
    HeartbeatIntervalMgmdMgmd: 2000
    TotalSendBufferMemory: 16M
    startNodeId: 49
  ndb:
    MaxNoOfAttributes: 5000
    MaxNoOfOrderedIndexes: 1024
    NoOfFragmentLogParts: 4
    MaxNoOfExecutionThreads: 4
    StopOnError: 0
    MaxNoOfTables: 1024
    NoOfFragmentLogFiles: 64
  api:
    user: mysql
    max_connections: 4096
    all_row_changes_to_bin_log: 1
    binlog_expire_logs_seconds: '86400'
    auto_increment_increment: 2
    auto_increment_offset: 1
    wait_timeout: 600
    interactive_timeout: 600

additionalndbconfigurations:
  mgm: {}
  ndb:
    __TransactionErrorLogLevel: '0x0000'
```

```

TotalSendBufferMemory: '32M'
CompressedLCP: true
TransactionDeadlockDetectionTimeout: 1200
HeartbeatIntervalDbDb: 500
ConnectCheckIntervalDelay: 0
LockPagesInMainMemory: 0
MaxNoOfConcurrentOperations: 128K
MaxNoOfConcurrentTransactions: 65536
MaxNoOfUniqueHashIndexes: 16K
FragmentLogFileSize: 128M
ODirect: false
RedoBuffer: 1024M
SchedulerExecutionTimer: 50
SchedulerSpinTimer: 0
TimeBetweenEpochs: 100
TimeBetweenGlobalCheckpoints: 2000
TimeBetweenLocalCheckpoints: 6
TimeBetweenEpochsTimeout: 4000
TimeBetweenGlobalCheckpointsTimeout: 60000
# By default LcpScanProgressTimeout is configured to overwrite
configure LcpScanProgressTimeout
# with required value.
# LcpScanProgressTimeout: 180
RedoOverCommitLimit: 60
RedoOverCommitCounter: 3
StartPartitionedTimeout: '1800000'
CompressedBackup: 'true'
MaxBufferedEpochBytes: '26214400'
MaxBufferedEpochs: '100'
api:
  TotalSendBufferMemory: '32M'
  DefaultOperationRedoProblemAction: 'ABORT'
mysql:
  max_connect_errors: '4294967295'
  ndb_applier_allow_skip_epoch: 0
  ndb_batch_size: '2000000'
  ndb_blob_write_batch_bytes: '2000000'
  replica_allow_batching: 'ON'
  max_allowed_packet: '134217728'
  ndb_log_update_minimal: 1
  replica_parallel_workers: 0
  binlog_transaction_compression: 'ON'
  binlog_transaction_compression_level_zstd: '3'
  ndb_report_thresh_binlog_epoch_slip: 50
  ndb_eventbuffer_max_alloc: 0
  ndb_allow_copying_alter_table: 'ON'
  ndb_clear_apply_status: 'OFF'
tcp:
  SendBufferMemory: '2M'
  ReceiveBufferMemory: '2M'
  TCP_SND_BUF_SIZE: '0'
  TCP_RCV_BUF_SIZE: '0'

# specific mysql cluster node values needed in different charts
mgm:
  ndbdisksize: 15Gi

```

```

ndb:
  ndbdisksize: 132Gi
  ndbbackupdisksize: 164Gi
  datamemory: 69G
  KeepAliveSendIntervalMs: 60000
  use_separate_backup_disk: true
  restoreparallelism: 128
api:
  ndbdisksize: 12.6Gi
  startNodeId: 56
  startEmptyApiSlotNodeId: 222
  numOfEmptyApiSlots: 4
  ndb_extra_logging: 99
  general_log: 'OFF'
ndbapp:
  ndbdisksize: 2Gi
  ndb_cluster_connection_pool: 1
  ndb_cluster_connection_pool_base_nodeid: 100
  startNodeId: 70

```

The following table describes UDR resources and their utilization:

Table 3-44 UDR Resources and their Utilization (Average Latency: 10ms)

Micro service name	Container name	Number of Pods	CPU Allocation Per Pod	Memory Allocation Per Pod	Total Resources
Ingress-gateway-sig	ingressgateway-sig	4	6 CPUs	4 GB	24 CPUs 16 GB
Ingress-gateway-prov	ingressgateway-prov	2	4 CPUs	4 GB	8 CPUs 8 GB
Nudr-dr-service	nudr-dr-service	5	6 CPUs	4 GB	30 CPUs 20 GB
Nudr-dr-provservice	nudr-dr-provservice	2	4 CPUs	4 GB	8 CPUs 8 GB
Nudr-egress-gateway	egressgateway	1	1 CPUs	1 GB	1 CPUs 1 GB
Nudr-config	nudr-config	1	1 CPUs	1 GB	1 CPUs 1 GB
Nudr-config-server	nudr-config-server	1	1 CPUs	1 GB	1 CPU 1 GB

Subscriber payload:

```

{
  "sm-data": {
    "umData": {
      "mk1": {
        "scopes": {
          "11-abc123": {
            "dnn": [

```

```

        "dnn1"
      ],
      "snssai": {
        "sd": "abc123",
        "sst": 11
      }
    }
  },
  "limitId": "mk1",
  "umLevel": "SERVICE_LEVEL",
  "resetTime": "2018-01-02T08:17:14.090Z",
  "allowedUsage": {
    "duration": 9000,
    "totalVolume": 8888,
    "uplinkVolume": 6666,
    "downlinkVolume": 7777
  }
},
"umDataLimits": {
  "mk1": {
    "scopes": {
      "11-abc123": {
        "dnn": [
          "dnn1"
        ],
        "snssai": {
          "sd": "abc123",
          "sst": 11
        }
      }
    }
  },
  "endDate": "2018-11-05T08:17:14.090Z",
  "limitId": "mk1",
  "umLevel": "SESSION_LEVEL",
  "startDate": "2018-09-05T08:17:14.090Z",
  "usageLimit": {
    "duration": 6000,
    "totalVolume": 9000,
    "uplinkVolume": 5000,
    "downlinkVolume": 4000
  },
  "resetPeriod": {
    "period": "YEARLY"
  }
},
"smPolicySnssaiData": {
  "11-abc123": {
    "snssai": {
      "sd": "abc123",
      "sst": 11
    }
  },
  "smPolicyDnnData": {
    "dnn1": {
      "dnn": "dnn1",

```

```

"bdtRefIds": {
    "xyz": "bdtRefIds",
    "abc": "xyz"
},
"gbrDl": "7788 Kbps",
"gbrUl": "5566 Kbps",
"online": true,
"chfInfo": {
    "primaryChfAddress": "1.1.1.1",
    "secondaryChfAddress": "2.2.2.2"
},
"offline": true,
"praInfos": {
    "p1": {
        "praId": "p1",
        "trackingAreaList": [{
            "plmnId": {
                "mcc": "976",
                "mnc": "32"
            },
            "tac": "5CB6"
        },
        {
            "plmnId": {
                "mcc": "977",
                "mnc": "33"
            },
            "tac": "5CB7"
        }
    ],
    "ecgiList": [{
        "plmnId": {
            "mcc": "976",
            "mnc": "32"
        },
        "eutraCellId": "92FFdBE"
    },
    {
        "plmnId": {
            "mcc": "977",
            "mnc": "33"
        },
        "eutraCellId": "8F868C4"
    }
    ],
    "ncgiList": [{
        "plmnId": {
            "mcc": "976",
            "mnc": "32"
        },
        "nrCellId": "b2fB6fE9D"
    },
    {
        "plmnId": {
            "mcc": "977",

```

}

The following table provides observation data for the performance test that can be used for the benchmark testing to scale up UDR performance:

Table 3-45 Result and Observation

Parameter	Values
Test Duration	48h
TPS Achieved	10K Signaling
Success rate	100%
Average UDR processing time (Request and Response)	10ms

3.4.3 Policy Data: Performance 17.2K N36 and 6.56K Notifications Policy

You can perform benchmark tests on UDR for compute and storage resources by considering the following conditions:

- Signaling : 17.2K
- Provisioning: NA
- Total Subscribers: 10M

The following table describes the benchmarking parameters and their values:

Table 3-46 Traffic Model Details

Request Type	Details	TPS
N36 traffic (100%) 17.2K TPS for sm-data and subs-to-notify	subs-to-notify POST	4.57K (26%)
	sm-data GET	4.63K (27%)
	subs-to-notify DELETE	1.39K (8%)
	sm-data PATCH	6.56K (39%)
NOTIFICATIONS	POST Operation (Egress)	6.56K

Note

Provisioning traffic are not included in this model.

The following table describes the test case parameters and their values:

Table 3-47 Testcase Parameters

Input Parameter Details	Configuration Values
Target TPS	17.2K Signaling
Notification Rate	6.56K
UDR Response Timeout	2.7s
Client Timeout	3s
Signaling Requests Latency Recorded on Client	150ms
Provisioning Requests Latency Recorded on Client	150ms

Table 3-48 Consolidated Resource Requirement

Resource	CPU	Memory
cnDBTier	84 CPUs	451 GB
UDR	229 CPUs	177 GB
Buffer	50 CPUs	50 GB
Total	363 CPUs	678 GB

The following table describes cnDBTier resources and their utilization:

Note

For cnDBTier, you must use ocudr_udr_10msub17.2K_TPS_dbtier_24.1.0_custom_values_24.1.0 file. For more information, see *Oracle Communications Cloud Native Core, Unified Data Repository Installation, Upgrade, and Fault Recovery Guide*.

Table 3-49 cnDBTier Resources and their Utilization

Microservice name	Container name	Number of Pods	CPU Allocation Per Pod	Memory Allocation Per Pod	Total Resources
Management node (ndbmcmd)	mysqlndbcluster	2	2 CPUs	9 GB	4 CPUs 18 GB
Data node (ndbmt)	mysqlndbcluster	4	4 CPUs	93 GB	16 CPUs 372 GB
APP SQL node (ndbappmysqld)	mysqlndbcluster	9	6 CPUs	3 GB	54 CPUs 27 GB
SQL node (ndbmysqld, used for replication)	mysqlndbcluster	2	4 CPUs	16 GB	8 CPUs 32 GB
DB Monitor Service	db-monitor-svc	1	200 millicores CPUs	500 MB	1 CPU 500 MB
DB Backup Manager Service	backup-manager-svc	1	100 millicores CPUs	128 MB	1 CPU 128 MB

The following table describes UDR resources and their utilization:

Table 3-50 UDR Resources and their Utilization

Micro service name	Container name	Number of Pods	CPU Allocation Per Pod	Memory Allocation Per Pod	Total Resources
Ingress-gateway-sig	ingressgateway-sig	9	6 CPUs	4 GB	54 CPUs 36 GB

Table 3-50 (Cont.) UDR Resources and their Utilization

Micro service name	Container name	Number of Pods	CPU Allocation Per Pod	Memory Allocation Per Pod	Total Resources
Ingress-gateway-prov	ingressgateway-prov	2	4 CPUs	4 GB	8 CPUs 8 GB
Nudr-dr-service	nudr-dr-service	17	6 CPUs	4 GB	102 CPUs 68 GB
Nudr-dr-provservice	nudr-dr-provservice	2	4 CPUs	4 GB	8 CPUs 8 GB
Nudr-notify-service	nudr-notify-service	7	4 CPUs	4 GB	28 CPUs 28 GB
Nudr-egress-gateway	egressgateway	4	4 CPUs	4 GB	16 CPUs 16 GB
Nudr-config	nudr-config	2	1 CPU	1 GB	2 CPUs 2 GB
Nudr-config-server	nudr-config-server	2	1 CPU	1 GB	2 CPUs 2 GB
Alternate-route	alternate-route	2	1 CPU	1 GB	2 CPUs 2 GB
Nudr-nrf-client-nfmanagement-service	nrf-client-nfmanagement	2	1 CPU	1 GB	2 CPUs 2 GB
App-info	app-info	2	1 CPU	1 GB	2 CPUs 2 GB
Perf-info	perf-info	2	1 CPU	1 GB	2 CPUs 2 GB
Nudr-dbc-auditor-service	nudr-dbc-auditor-service	1	1 CPU	1 GB	1 CPU 1 GB

The following table provides observation data for the performance test that can be used for the benchmark testing to scale up UDR performance:

Table 3-51 Result and Observation

Parameter	Values
Test Duration	2h
TPS Achieved	17.2K Signaling
Success rate	100%
Average UDR processing time (Request and Response)	150ms