

# Oracle® Communications

## Cloud Native Core, Unified Data Repository

### Benchmarking Guide



Release 25.2.100  
G45153-02  
November 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 Model: 67K lookup + 1.44K Provisioning TPS (64M Subscribers)	1
3.2	Test Scenario 2: EIR 10K TPS and 10k Diameter S13 Interface TPS (600 K Subscribers)	7
3.3	Test Scenario 3: 25K N36 and 600 Provisioning (SOAP) Profile (35M Subscribers)	12
3.4	Test Scenario 4: 17.2K N36 + 10K SH and 600 Provisioning (SOAP) Profile (35M Subscribers)	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
<b>boldface</b>	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 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 25.2.1xx.

## Release 25.2.100 - G45153-02, November 2025

- Updated the UDR CPU value in [Table 3-17](#) in the [Test Scenario 3: 25K N36 and 600 Provisioning \(SOAP\) Profile \(35M Subscribers\)](#) section.
- Updated the ephemeral storage value for Nudr-dr-provservice in [Table 3-26](#) in the [Test Scenario 4: 17.2K N36 + 10K SH and 600 Provisioning \(SOAP\) Profile \(35M Subscribers\)](#) section.

## Release 25.2.100 - G45153-01, November 2025

- Added the [Test Scenario 1: SLF Call Model: 67K lookup + 1.44K Provisioning TPS \(64M Subscribers\)](#) section.
- Updated the TPS value in [Table 3-22](#) and Average Latency value in [Table 3-26](#) in the [Test Scenario 4: 17.2K N36 + 10K SH and 600 Provisioning \(SOAP\) Profile \(35M Subscribers\)](#) section.
- Updated the [Table 3-11](#) and [Table 3-12](#) in the [Test Scenario 2: EIR 10K TPS and 10k Diameter S13 Interface TPS \(600 K Subscribers\)](#) section.
- Removed the SOAP and Diameter Deployment Model (1M - 10M Subscribers) section from this document.

# 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, Unified Data Repository REST Specification Guide*
- *Oracle Communications Cloud Native Core, Unified Data Repository Troubleshooting Guide*
- *Oracle Communications Cloud Native Core, Provisioning Gateway Installation Guide*
- *Oracle Communications Cloud Native Core, Provisioning Gateway Interface Specification 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) 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 Model: 67K lookup + 1.44K Provisioning TPS (64M Subscribers)

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 SLF in this testcase:

- Support for Default Group ID in SLF
- Controlled Shutdown of an Instance
- TLS 1.3 Support for Kubernetes API Server Communication
- Error Logging Enhancement
- Error Response Enhancement
- OAuth2
- Alternate Routing Service
- Support for User-Agent Header
- Overload Handling
- Support for LCI and OCI Header
- Network Function Scoring for a Site
- Conflict Resolution
- Ingress Gateway Pod Protection Using Rate Limiting

The following features are enabled for Provisioning Gateway in this testcase:

- Auditor Service
- Provgw global configuration

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

- Signaling (SLF Look Up): 67K TPS
- Provisioning: 1.44 K
- Total Subscribers: 64 Million
- Profile Size: 450 bytes

The following table describes the benchmarking parameters and their values:

**Table 3-1 Traffic Model Details**

Request Type	Details	TPS
Lookup 67k	SLF Lookup GET Requests	67K
Provisioning (1.44K using Provgw)	CREATE	210
	DELETE	210
	UPDATE	510
	GET	510

The following table describes the testcase parameters and their values:

**Table 3-2 Testcase Parameters**

Input Parameter Details	Configuration Values
SLF Version Tag	25.2.100
Target TPS	67K Lookup + 1.44K Provisioning
Traffic Profile	SLF 67K Profile
Notification Rate	OFF
SLF Response Timeout	900ms
Signaling Requests Latency Recorded	31ms
Provisioning Requests Latency Recorded	50ms

The following table describes consolidated resource requirement:

**Table 3-3 Consolidated Resource Requirement per Site**

Resource	CPU	Memory	Ephemeral Storage	PVC
cnDBTier	177	531 GB	69 GB	1094 GB
SLF	517	295 GB	67 GB	NA
ProvGw	39	39 GB	9 GB	5 GB
Buffer	50	50 GB	50 GB	50 GB
Total	783	915 GB	195 GB	1149 GB

**Note**

All values are inclusive of ASM sidecar.

The following table provides OSO resources:

**Table 3-4 OSO Resources (Retention period: 14 days)**

Service	Replicas	CPU Limit	RAM Limit	PVC	Ephemeral storage
Prometheus (snapshot utility enabled)	1	4	8 GB	75 GB	12 GB
Prometheus AlertManager	2	4	4 GB	10 GB	1 GB
Total	3	12	16 GB	95 GB	14 GB

**Table 3-5 cnDBTier Resources**

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 (ndbmcmd)	mysqlndbcluster	2	2 CPUs	12 GB	1 GB	16 GB	6 CPUs 26 GB <b>Ephemeral Storage:</b> 2 GB <b>PVC Allocation:</b> 32 GB
	istio-proxy		1 CPUs	1 GB			
Data node (ndbmt)	mysqlndbcluster	6	4 CPUs	50 GB	1 GB	65 GB (Backup: 63 GB)	42 CPUs 324 GB <b>Ephemeral Storage:</b> 6 GB <b>PVC Allocation:</b> 768 GB
	istio-proxy		2 CPUs	2 GB			
	db-backup-executor-svc		1 CPU	2 GB			
APP SQL node (ndbappmysql)	mysqlndbcluster	10	6 CPUs	4 GB	1 GB	10 GB	90 CPUs 60 GB <b>Ephemeral Storage:</b> 10 GB <b>PVC Allocation:</b> 100 GB
	istio-proxy		3 CPUs	2 GB			
SQL node (Used for Replication) (ndbmysqld)	mysqlndbcluster	4	4 CPUs	16 GB	1 GB	16 GB	25 CPUs 41 GB <b>Ephemeral Storage:</b> 2 GB <b>PVC Allocation:</b> 32 GB
	istio-proxy		2 CPUs	4 GB			
	init-sidecar		100m CPU	256 MB			
DB Monitor Service (db-monitor-svc)	db-monitor-svc	1	4 CPUs	4 GB	1 GB	NA	5 CPUs 5 GB <b>Ephemeral Storage:</b> 1 GB
	istio-proxy		1 CPUs	1 GB			

Table 3-5 (Cont.) cnDBTier Resources

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)
DB Backup Manager Service (backup-manager-svc)	backup-manager-svc	1	1 CPU	1 GB	1 GB	NA	2 CPUs 2 GB
	istio-proxy		1 CPUs	1 GB			<b>Ephemeral Storage:</b> 1 GB
Replication Service (db-replication-svc)	db-replication-svc	2	2 CPU	12 GB	1 GB	160 GB	4 CPUs 13 GB
	istio-proxy		200m CPU	500 MB			<b>Ephemeral Storage:</b> 1 GB <b>PVC Allocation:</b> 160 GB

Additional cnDBTier configuration are as follows:

```

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"

The following table describes SLF resources (Lookup Latency: 31ms):

**Table 3-6 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	26	6 CPUs	4 GB	1 GB	260 CPUs 156 GB Memory <b>Ephemeral Storage: 26 GB</b>
	istio-proxy		4 CPUs	2 GB		
Ingress-gateway-prov	ingressgateway-prov	2	4 CPUs	4 GB	1 GB	12 CPUs 12 GB <b>Ephemeral Storage: 2 GB</b>
	istio-proxy		2 CPUs	2 GB		
Nudr-dr-service	nudr-dr-service	22	6 CPUs	4 GB	1 GB	198 CPUs 132 GB <b>Ephemeral Storage: 22 GB</b>
	istio-proxy		3 CPUs	2 GB		
Nudr-dr-provservice	nudr-dr-provservice	2	4 CPUs	4 GB	1 GB	12 CPUs 12 GB <b>Ephemeral Storage: 2 GB</b>
	istio-proxy		2 CPUs	2 GB		
Nudr-nrf-client-nfmanagement	nrf-client-nfmanagement	2	1 CPU	1 GB	1 GB	4 CPUs 4 GB <b>Ephemeral Storage: 2 GB</b>
	istio-proxy		1 CPUs	1 GB		
Nudr-egress-gateway	egressgateway	2	1 CPUs	1 GB	1 GB	4 CPUs 4 GB <b>Ephemeral Storage: 2 GB</b>
	istio-proxy		1 CPUs	1 GB		
Nudr-config	nudr-config	2	2 CPUs	2 GB	1 GB	6 CPUs 6 GB <b>Ephemeral Storage: 2 GB</b>
	istio-proxy		1 CPUs	1 GB		
Nudr-config-server	nudr-config-server	2	2 CPUs	2 GB	1 GB	6 CPUs 6 GB <b>Ephemeral Storage: 2 GB</b>
	istio-proxy		1 CPUs	1 GB		

Table 3-6 (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
alternate-route	alternate-route	2	1 CPUs	1 GB	1 GB	4 CPUs 4 GB <b>Ephemeral Storage: 2 GB</b>
	istio-proxy		1 CPUs	1 GB		
app-info	app-info	2	1 CPUs	1 GB	1 GB	4 CPUs 4 GB <b>Ephemeral Storage: 2 GB</b>
	istio-proxy		1 CPUs	1 GB		
perf-info	perf-info	2	1 CPUs	1 GB	1 GB	4 CPUs 4 GB <b>Ephemeral Storage: 2 GB</b>
	istio-proxy		1 CPUs	1 GB		
Nudr-dbcrauditor	nudr-dbcrauditor-service	1	2 CPUs	2 GB	1 GB	3 CPUs 3 GB <b>Ephemeral Storage: 1 GB</b>
	istio-proxy		1 CPUs	1 GB		

**Note**

The same resources and usage are used for all sites.

Table 3-7 Provision Gateway Resources (Provisioning Latency: 50 ms)

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	4 CPUs	4 GB	1 GB	10 CPUs 12 GB Memory <b>Ephemeral Storage: 2 GB</b>
	istio-proxy		1 CPUs	2 GB		
provgw-egress-gateway	egressgateway	2	4 CPUs	4 GB	1 GB	10 CPUs 12 GB Memory <b>Ephemeral Storage: 2 GB</b>
	istio-proxy		1 CPUs	2 GB		
provgw-service	provgw-service	2	4 CPUs	2 GB	1 GB	10 CPUs 8 GB Memory



**Table 3-7 (Cont.) Provision Gateway Resources (Provisioning Latency: 50 ms)**

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	2 GB		<b>Ephemeral Storage: 2 GB</b>
provgw-config	provgw-config	2	2 CPUs	2 GB	1 GB	6 CPUs 6 GB Memory <b>Ephemeral Storage: 2 GB</b>
	istio-proxy		1 CPUs	1 GB		
provgw-auditor-service	auditor-service	1	2 CPUs	2 GB	1 GB	3 CPUs 4 GB Memory <b>Ephemeral Storage: 1 GB</b>
	istio-proxy		1 CPU	2 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 3-8 Result and Observation**

Parameter	Values
TPS Achieved	67K SLF Lookup + 1.44K 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 10K TPS and 10k Diameter S13 Interface TPS (600 K Subscribers)

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 enabled for this testcase:

- Support for EIR International Mobile Equipment Identity Software Version (IMEISV) Fallback
- Overload Handling
- Ingress Gateway Pod Protection Using Rate Limiting
- Auto Create
- Diameter S13 Interface
- International Mobile Subscriber Identity (IMSI) Fallback Lookup

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

- EIR Look Up: 20K
- Total Subscribers: 600 K
- Profile Size: 130 bytes

The following table describes the benchmarking parameters and their values:

**Table 3-9 Traffic Model Details**

Request Type	Details	TPS
EIR GET	N17 GET Request	10K
ECR message	Diameter S13 Interface ECR	10K

The following table describes the testcase parameters and their values:

**Table 3-10 Testcase Parameters**

Input Parameter Details	Configuration Values
UDR Version Tag	25.2.100
Target TPS	20K Lookup
Traffic Profile	20K
EIR Response Timeout	2.7s
Client Timeout	10s
N17 Latency:	6 ms
S13 Latency	7 ms

The following table describes the consolidated resource requirements:

**Table 3-11 Consolidated Resource Requirement per Site**

Resource	CPUs	Memory	Ephemeral Storage	PVC
cnDBTier	55	141 GB	16 GB	516 GB
EIR	131	97 GB	34 GB	NA GB
Buffer	50	50 GB	20 GB	50 GB
Total	236	288 GB	70 GB	566 GB

The following table describes cnDBTier resources:

Table 3-12 cnDBTier Resources

Micro service name	Container name	Number of Pods (cnDBTier 1)	CPU Allocation Per Pod (cnDBTier 1)	Memory Allocation Per Pod (cnDBTier 1)	PVC Allocation Per Pod	Ephemeral Storage Per Pod	Total Resources (cnDBTier 1)
Management node	mysqlndbcluster	2	2 CPUs	12 GB	16 GB	1 GB	4 CPU 24 GB Memory <b>Ephemeral Storage:</b> 2 GB <b>PVC Allocation:</b> 32 GB
Data node	mysqlndbcluster	4	4 CPUs	20 GB	28 GB <b>Backup:</b> 56 GB	1 GB	16 CPU 80 GB Memory <b>Ephemeral Storage:</b> 4 GB <b>PVC Allocation:</b> 336 GB
APP SQL node	mysqlndbcluster	5	4 CPUs	4 GB	10 GB	1 GB	20 CPU 20 GB Memory <b>Ephemeral Storage:</b> 5 GB <b>PVC Allocation:</b> 50 GB
SQL node (Used for Replication)	mysqlndbcluster	2	4 CPUs	5 GB	16 GB	1 GB	8 CPU 10 GB Memory <b>Ephemeral Storage:</b> 2 GB <b>PVC Allocation:</b> 32 GB
DB Monitor Service (db-monitor-svc)	db-monitor-svc	1	4 CPUs	4 GB	NA	1 GB	4 CPUs 4 GB Memory <b>Ephemeral Storage:</b> 1 GB

Table 3-12 (Cont.) cnDBTier Resources

Micro service name	Container name	Number of Pods (cnDBTier 1)	CPU Allocation Per Pod (cnDBTier 1)	Memory Allocation Per Pod (cnDBTier 1)	PVC Allocation Per Pod	Ephemeral Storage Per Pod	Total Resources (cnDBTier 1)
DB Backup Manager Service (backup-manager-svc)	backup-manager-svc	1	100m CPUs	128 MB	NA	1 GB	100m CPUs 128 MB Memory <b>Ephemeral Storage: 1 GB</b>
Replication Service (db-replication-svc)	db-replication-svc	1	2 CPU	2 GB	NA	1 GB	2 CPUs 2 MB Memory <b>Ephemeral Storage: 1 GB</b> <b>PVC Allocation: 66 GB</b>

The following table describes EIR resources:

Table 3-13 EIR Resources

Micro service name	Container name	Number of Pods	CPU Allocation Per Pod	Memory Allocation Per Pod	Ephemeral Storage Per Pod	Total Resources
Ingress-gateway-sig	Ingress-gateway-sig	4	6 CPUs	4 GB	1 GB	24 CPUs 16 GB Memory <b>Ephemeral Storage: 4 GB</b>
Ingress-gateway-prov	Ingress-gateway-prov	2	4 CPUs	4 GB	1 GB	8 CPUs 8 GB Memory <b>Ephemeral Storage: 2 GB</b>
Nudr-dr-service	nudr-dr-service	3	6 CPUs	4 GB	1 GB	18 CPUs 12 GB Memory <b>Ephemeral Storage: 3 GB</b>

Table 3-13 (Cont.) EIR Resources

Micro service name	Container name	Number of Pods	CPU Allocation Per Pod	Memory Allocation Per Pod	Ephemeral Storage Per Pod	Total Resources
Nudr-dr-provservice	nudr-dr-provservice	2	4 CPUs	4 GB	1 GB	8 CPUs 8 GB Memory <b>Ephemeral Storage: 2 GB</b>
Nudr-diam-gateway	nudr-diam-gateway	2	6 CPUs	4 GB	1 GB	12 CPUs 8 GB Memory <b>Ephemeral Storage: 2 GB</b>
Nudr-diameterproxy	nudr-diameterproxy	8	6 CPUs	4 GB	1 GB	48 CPUs 32 GB Memory <b>Ephemeral Storage: 8 GB</b>
Nudr-config	nudr-config	2	1 CPUs	1 GB	1 GB	2 CPUs 2 GB Memory <b>Ephemeral Storage: 2 GB</b>
Nudr-config-server	nudr-config-server	2	1 CPUs	1 GB	1 GB	2 CPU 2 GB Memory <b>Ephemeral Storage: 2 GB</b>
Alternate-route	alternate-route	2	1 CPUs	1 GB	1 GB	2 CPU 2 GB Memory <b>Ephemeral Storage: 2 GB</b>
Nudr-nrf-client-nfmanagement-service	nrf-client-nfmanagement	2	1 CPUs	1 GB	1 GB	2 CPU 2 GB Memory <b>Ephemeral Storage: 2 GB</b>

**Table 3-13 (Cont.) EIR Resources**

Micro service name	Container name	Number of Pods	CPU Allocation Per Pod	Memory Allocation Per Pod	Ephemeral Storage Per Pod	Total Resources
App-info	app-info	2	1 CPUs	1 GB	1 GB	2 CPU 2 GB Memory <b>Ephemeral Storage: 2 GB</b>
Perf-info	perf-info	2	1 CPUs	1 GB	1 GB	2 CPU 2 GB Memory <b>Ephemeral Storage: 2 GB</b>
Nudr-dbcr-auditor-service	nudr-dbcr-auditor-service	1	1 CPUs	1 GB	1 GB	1 CPU 1 GB Memory <b>Ephemeral Storage: 1 GB</b>

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

**Table 3-14 Result and Observation**

Parameter	Values
TPS Achieved	20K
Success Rate	100%
Average EIR processing time (Request and Response)	9 ms

## 3.3 Test Scenario 3: 25K N36 and 600 Provisioning (SOAP) Profile (35M Subscribers)

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

- Signaling: 25K TPS
- Provisioning: 600 TPS
- Total Subscribers: 35 Million

The following features are enabled for this testcase:

- Auto Enrollment and Auto Create Features
- Overload Handling
- ETag (Entity Tag)

- Ingress Gateway Pod Protection Using Rate Limiting
- Support for User-Agent Header
- 3gpp-Sbi-Correlation-Info Header
- Suppress Notification
- Support for Post Operation for an Existing Subscription
- Subscriber Activity Logging

The following table describes the benchmarking parameters and their values:

**Table 3-15 Traffic Model Details**

Request Type	Details	TPS
N36 traffic 25K TPS	subs-to-notify POST	5K (20%)
	sm-data GET	5K (20%)
	subs-to-notify DELETE	5K (20%)
	sm-data PATCH	10K (40%)
SOAP PROVISIONING 600 TPS	GET	100
	UPDATE QUOTA	50
	UPDATE DYNAMIC QUOTA	50
	UPDATE STATE	100
	UPDATE SUBSCRIBER	100
	CREATE SUBSCRIBER	100
	DELETE SUBSCRIBER	100

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

**Table 3-16 Testcase Parameters**

Input Parameter Details	Configuration Values
UDR Version Tag	25.2.100
Target TPS	25K TPS Signaling + 600 SOAP
Notification Rate	300
UDR Response Timeout	2.7s
Signaling Requests Latency Recorded	25ms
Provisioning Requests Latency Recorded	40ms

**Table 3-17 Consolidated Resource Requirement per Site**

Resource	CPU	Memory	Ephemeral Storage	PVC
cnDBTier	163 CPUs	640 GB	29 GB	1955 GB
UDR	257 CPUs	184 GB	55 GB	NA
PROVGW	28	28	10	NA
Buffer	50 CPUs	50 GB	20 GB	200 GB
<b>Total</b>	498 CPUs	902 GB	114 GB	2155 GB

The following table describes cnDBTier resources:

**Table 3-18 cnDBTier Resources**

Microservice name	Container name	Number of Pods	CPU Allocation Per Pod	Memory Allocation Per Pod	Ephemeral Storage Per Pod	PVC Allocation Per Pod	Total Resources
Management node (ndbmcmd)	mysqlndbcluster	2	2 CPUs	9 GB	1 GB	15 GB	4 CPUs 18 GB <b>Ephemeral Storage:</b> 2 GB <b>PVC Allocation:</b> 30 GB
Data node (ndbmtnd)	mysqlndbcluster	4	9 CPUs	124 GB	1 GB	132 GB Backup: 220 GB	36 CPUs 496 GB <b>Ephemeral Storage:</b> 4 GB <b>PVC Allocation:</b> 1408 GB
APP SQL node (ndbappmysql)	mysqlndbcluster	18	6 CPUs	4 GB	1 GB	4 GB	108 CPUs 72 GB <b>Ephemeral Storage:</b> 18 GB <b>PVC Allocation:</b> 72 GB
SQL node (ndbmssql, used for replication)	mysqlndbcluster	2	4 CPUs	24 GB	1 GB	110 GB	8 CPUs 48 GB <b>Ephemeral Storage:</b> 2 GB <b>PVC Allocation:</b> 220 GB
DB Monitor Service (db-monitor-svc)	db-monitor-svc	1	4 CPUs	4 GB	1 GB	NA	4 CPU 4 MB <b>Ephemeral Storage:</b> 1 GB
DB Backup Manager Service (backup-manager-svc)	backup-manager-svc	1	100 millicores CPUs	128 MB	1 GB	NA	1 CPU 128 MB <b>Ephemeral Storage:</b> 1 GB



**Table 3-18 (Cont.) cnDBTier Resources**

Microservice name	Container name	Number of Pods	CPU Allocation Per Pod	Memory Allocation Per Pod	Ephemeral Storage Per Pod	PVC Allocation Per Pod	Total Resources
Replication service (Multi site cases)	replication-svc	1	2 CPUs	2 GB	1 GB	225 GB	2 CPUs 2 GB <b>Ephemeral Storage: 1 GB</b> <b>PVC Allocation: 225 GB</b>

The following table describes UDR resources:

**Table 3-19 UDR Resources (Average Latency: 33ms (N36/PROV) and N36: 25ms / PROV: 40ms)**

Micro service 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	13	6 CPUs	4 GB	1 GB	78 CPUs 52 GB <b>Ephemeral Storage: 13 GB</b>
Ingress-gateway-prov	ingressgateway-prov	2	4 CPUs	4 GB	1 GB	8 CPUs 8 GB <b>Ephemeral Storage: 2 GB</b>
Nudr-dr-service	nudr-dr-service	20	6 CPUs	4 GB	1 GB	120 CPUs 80 GB <b>Ephemeral Storage: 20 GB</b>
Nudr-dr-provservice	nudr-dr-provservice	2	4 CPUs	4 GB	1 GB	8 CPUs 8 GB <b>Ephemeral Storage: 2 GB</b>
Nudr-notify-service	nudr-notify-service	3	6 CPUs	5 GB	1 GB	18 CPUs 15 GB <b>Ephemeral Storage: 3 GB</b>

**Table 3-19 (Cont.) UDR Resources (Average Latency: 33ms (N36/PROV) and N36: 25ms / PROV: 40ms)**

Micro service name	Container name	Number of Pods	CPU Allocation Per Pod	Memory Allocation Per Pod	Ephemeral Storage Per Pod	Total Resources
Nudr-egress-gateway	egressgateway	2	6 CPUs	4 GB	1 GB	12 CPUs 8 GB <b>Ephemeral Storage: 2 GB</b>
Nudr-config	nudr-config	2	1 CPU	1 GB	1 GB	2 CPUs 2 GB <b>Ephemeral Storage: 2 GB</b>
Nudr-config-server	nudr-config-server	2	1 CPU	1 GB	1 GB	2 CPUs 2 GB <b>Ephemeral Storage: 2 GB</b>
Alternate-route	alternate-route	2	1 CPU	1 GB	1 GB	2 CPUs 2 GB <b>Ephemeral Storage: 2 GB</b>
Nudr-nrf-client-nfmanagement-service	nrf-client-nfmanagement	2	1 CPU	1 GB	1 GB	2 CPUs 2 GB <b>Ephemeral Storage: 2 GB</b>
App-info	app-info	2	1 CPU	1 GB	1 GB	2 CPUs 2 GB <b>Ephemeral Storage: 2 GB</b>
Perf-info	perf-info	2	1 CPU	1 GB	1 GB	2 CPUs 2 GB <b>Ephemeral Storage: 2 GB</b>
Nudr-dbcrauditor-service	nudr-dbcrauditor-service	1	1 CPU	1 GB	1 GB	1 CPU 1 GB <b>Ephemeral Storage: 1 GB</b>

The following table describes Provisioning Gateway resources:

**Table 3-20 Provisioning Gateway Resources**

Micro service name	Container name	Number of Pods	CPU Allocation Per Pod	Memory Allocation Per Pod	Total Resources
provgw-ingress-gateway	ingressgateway	2	4 CPUs	4 GB	8 CPUs 8 GB <b>Ephemeral Storage: 2 GB</b>
provgw-egress-gateway	egressgateway	2	4 CPUs	4 GB	8 CPUs 8 GB <b>Ephemeral Storage: 2 GB</b>
provgw-service	provgw-service	4	2 CPUs	2 GB	8 CPUs 8 GB <b>Ephemeral Storage: 4 GB</b>
provgw-config	provgw-config	2	2 CPUs	2 GB	4 CPUs 4 GB Memory <b>Ephemeral Storage: 2 GB</b>

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

**Table 3-21 Result and Observation**

Parameter	Values
TPS Achieved	25K Signaling + 600
UDR Request Average Latency	33ms
Success rate	100%

## 3.4 Test Scenario 4: 17.2K N36 + 10K SH and 600 Provisioning (SOAP) Profile (35M Subscribers)

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

- Signaling : 17.2K N36 + 10K SH
- Provisioning: 600 TPS
- Total Subscribers: 35 Million

The following features are enabled for this testcase:

- Auto Enrollment and Auto Create Features
- Overload Handling
- ETag (Entity Tag)
- Ingress Gateway Pod Protection Using Rate Limiting

- Support for User-Agent Header
- 3gpp-Sbi-Correlation-Info Header
- Suppress Notification
- Subscriber Activity Logging
- Diameter Gateway Pod Congestion Control
- Support for Post Operation for an Existing Subscription

The following table describes the benchmarking parameters and their values:

**Table 3-22 Traffic Model Details**

Request Type	Details	TPS
N36 17.2K TPS	subs-to-notify POST	5K (20%)
	sm-data GET	5K (20%)
	subs-to-notify DELETE	5K (20%)
	sm-data PATCH	5K (20%)
SH 10K TPS	UDR	4K
	PUR	1.2K
	SNR	4.8K
SH PNR 1.2K TPS	PNR	1.2K
SOAP PROVISIONING 600 TPS	GET	100
	UPDATE QUOTA	50
	UPDATE DYNAMIC QUOTA	50
	UPDATE STATE	100
	UPDATE SUBSCRIBER	100
	CREATE SUBSCRIBER	100
	DELETE SUBSCRIBER	100

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

**Table 3-23 Testcase Parameters**

Input Parameter Details	Configuration Values
UDR Version Tag	25.2.100
Target TPS	17.2K N36 + 10K SH + 600 SOAP
Notification Rate	1.2K
UDR Response Timeout	2.7s
Signaling Requests Latency Recorded	25ms
Provisioning Requests Latency Recorded	30ms
Diameter (SH) Requests Latency Recorded	18ms

**Table 3-24 Consolidated Resource Requirement per Site**

Resource	CPU	Memory	Ephemeral Storage	PVC
cnDBTier	139 CPUs	625 GB	25 GB	1939

**Table 3-24 (Cont.) Consolidated Resource Requirement per Site**

Resource	CPU	Memory	Ephemeral Storage	PVC
UDR	281 CPUs	201 GB	60 GB	NA
PROVGW	28	28	10	NA
Buffer	50 CPUs	50 GB	20 GB	200 GB
<b>Total</b>	<b>498 CPUs</b>	<b>901 GB</b>	<b>115 GB</b>	<b>2139 GB</b>

The following table describes cnDBTier resources:

**Table 3-25 cnDBTier Resource**

Microservice name	Container name	Number of Pods	CPU Allocation Per Pod	Memory Allocation Per Pod	Ephemeral Storage Per Pod	PVC Allocation Per Pod	Total Resources
Management node (ndbmcmd)	mysqlndbcluster	2	2 CPUs	9 GB	1 GB	15 GB	4 CPUs 18 GB <b>Ephemeral Storage: 2 GB</b> <b>PVC Allocation: 30 GB</b>
Data node (ndbmtd)	mysqlndbcluster	4	9 CPUs	124 GB	1 GB	132 GB Backup: 220 GB	36 CPUs 496 GB <b>Ephemeral Storage: 4 GB</b> <b>PVC Allocation: 1408 GB</b>
APP SQL node (ndbappmysqld)	mysqlndbcluster	14	6 CPUs	4 GB	1 GB	4 GB	84 CPUs 56 GB <b>Ephemeral Storage: 14 GB</b> <b>PVC Allocation: 56 GB</b>
SQL node (ndbmysqld, used for replication)	mysqlndbcluster	2	4 CPUs	24 GB	1 GB	110 GB	8 CPUs 48 GB <b>Ephemeral Storage: 2 GB</b> <b>PVC Allocation: 220 GB</b>

**Table 3-25 (Cont.) cnDBTier Resource**

Microservice name	Container name	Number of Pods	CPU Allocation Per Pod	Memory Allocation Per Pod	Ephemeral Storage Per Pod	PVC Allocation Per Pod	Total Resources
DB Monitor Service	db-monitor-svc	1	4 CPUs	4 GB	1 GB	NA	4 CPU 4 MB <b>Ephemeral Storage: 1 GB</b>
DB Backup Manager Service	backup-manager-svc	1	100 millicores CPUs	128 MB	1 GB	NA	1 CPU 128 MB <b>Ephemeral Storage: 1 GB</b>
Replication service (Multi site cases)	replication-svc	1	2 CPUs	2 GB	1 GB	225 GB	2 CPUs 2 GB <b>Ephemeral Storage: 1 GB</b> <b>PVC Allocation: 225 GB</b>

The following table describes UDR resources:

**Table 3-26 UDR Resources (Average Latency: 25ms [25ms (N36) / 18ms (SH) / 30ms Provisioning]).**

Micro service 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	9	6 CPUs	4 GB	1 GB	54 CPUs 36 GB <b>Ephemeral Storage: 9 GB</b>
Ingress-gateway-prov	ingressgateway-prov	2	4 CPUs	4 GB	1 GB	8 CPUs 8 GB <b>Ephemeral Storage: 2 GB</b>
Nudr-dr-service	nudr-dr-service	17	6 CPUs	4 GB	1 GB	102 CPUs 68 GB <b>Ephemeral Storage: 17 GB</b>

**Table 3-26 (Cont.) UDR Resources (Average Latency: 25ms [25ms (N36) / 18ms (SH) / 30ms Provisioning]).**

Micro service name	Container name	Number of Pods	CPU Allocation Per Pod	Memory Allocation Per Pod	Ephemeral Storage Per Pod	Total Resources
Nudr-dr-provservice	nudr-dr-provservice	2	4 CPUs	4 GB	1 GB	8 CPUs 8 GB <b>Ephemeral Storage: 2 GB</b>
Nudr-notify-service	nudr-notify-service	3	6 CPUs	5 GB	1 GB	18 CPUs 15 GB <b>Ephemeral Storage: 3 GB</b>
Nudr-egress-gateway	egressgateway	2	6 CPUs	4 GB	1 GB	12 CPUs 8 GB <b>Ephemeral Storage: 2 GB</b>
Nudr-diam-gateway	nudr-diam-gateway	2	6 CPUs	5 GB	1 GB	12 CPUs 10 GB <b>Ephemeral Storage: 2 GB</b>
Nudr-diameterproxy	nudr-diameterproxy	9	6 CPUs	4 GB	1 GB	54 CPUs 36 GB <b>Ephemeral Storage: 9 GB</b>
Nudr-config	nudr-config	2	1 CPU	1 GB	1 GB	2 CPUs 2 GB <b>Ephemeral Storage: 2 GB</b>
Nudr-config-server	nudr-config-server	2	1 CPU	1 GB	1 GB	2 CPUs 2 GB <b>Ephemeral Storage: 2 GB</b>
Alternate-route	alternate-route	2	1 CPU	1 GB	1 GB	2 CPUs 2 GB <b>Ephemeral Storage: 2 GB</b>
Nudr-nrf-client-nfmanagement-service	nrf-client-nfmanagement	2	1 CPU	1 GB	1 GB	2 CPUs 2 GB <b>Ephemeral Storage: 2 GB</b>

**Table 3-26 (Cont.) UDR Resources (Average Latency: 25ms [25ms (N36) / 18ms (SH) / 30ms Provisioning]).**

Micro service name	Container name	Number of Pods	CPU Allocation Per Pod	Memory Allocation Per Pod	Ephemeral Storage Per Pod	Total Resources
App-info	app-info	2	1 CPU	1 GB	1 GB	2 CPUs 2 GB <b>Ephemeral Storage: 2 GB</b>
Perf-info	perf-info	2	1 CPU	1 GB	1 GB	2 CPUs 2 GB <b>Ephemeral Storage: 2 GB</b>
Nudr-dbcrauditor-service	nudr-dbcrauditor-service	1	1 CPU	1 GB	1 GB	1 CPU 1 GB <b>Ephemeral Storage: 1 GB</b>

The following table describes Provisioning Gateway resources:

**Table 3-27 Provisioning Gateway Resources**

Micro service name	Container name	Number of Pods	CPU Allocation Per Pod	Memory Allocation Per Pod	Total Resources
provgw-ingress-gateway	ingressgateway	2	4 CPUs	4 GB	8 CPUs 8 GB <b>Ephemeral Storage: 2 GB</b>
provgw-egress-gateway	egressgateway	2	4 CPUs	4 GB	8 CPUs 8 GB <b>Ephemeral Storage: 2 GB</b>
provgw-service	provgw-service	4	2 CPUs	2 GB	8 CPUs 8 GB <b>Ephemeral Storage: 4 GB</b>
provgw-config	provgw-config	2	2 CPUs	2 GB	4 CPUs 4 GB Memory <b>Ephemeral Storage: 2 GB</b>

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



**Table 3-28 Result and Observation**

Parameter	Values
TPS Achieved	17.2K N36 + 10K SH + 600
UDR Request Average Latency	25ms
Success rate	100%