

Oracle® Communications

Cloud Native Core, Unified Data Repository Benchmarking Guide



Release 25.1.200

G37599-01

July 2025



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: 50K 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: SOAP and Diameter Deployment Model (1M - 10M Subscribers) | 12 |
| 3.4 | Test Scenario 4: 25K N36 and 600 Provisioning (SOAP) Profile (35M Subscribers) | 17 |
| 3.5 | Test Scenario 5: 17.2K N36 + 10K SH and 600 Provisioing (SOAP) Profile (35M Subscribers) | 22 |

Preface

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

Documentation Accessibility

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

Access to Oracle Support

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

Diversity and Inclusion

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

Conventions

The following text conventions are used in this document:

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

My Oracle Support

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

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

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

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

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

Acronyms

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

Table Acronyms 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 25.1.2xx.

Release 25.1.200 - G37599-01, July 2025

- Updated the values in the Consolidated Resource Requirement, CPU Allocation Per Pod, and Total Resources in the following sections:
 - [SLF Call Model: 50K lookup + 1.44K Provisioning TPS \(64M Subscribers\)](#)
 - [Test Scenario 4: 25K N36 and 600 Provisioning \(SOAP\) Profile \(35M Subscribers\)](#)
 - [Test Scenario 5: 17.2K N36 + 10K SH and 600 Provisioning \(SOAP\) Profile \(35M Subscribers\)](#)
- Removed the SLF Call Model: 34K TPS and 1.44K Provisioning TPS section from this document.
- Removed the Policy Data: 17.2K N36, 300 TPS Notifications and 500 TPS Provisioning 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 Deployment Model

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

3.1.1 SLF Call Model: 50K 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:

- Ingress Gateway Pod Protection Using Rate Limiting
- OAuth2
- Alternate Routing Service
- Support for User-Agent Header
- Overload Handling
- Support for LCI and OCI Header
- Auto Create
- Network Function Scoring for a Site
- Conflict Resolution
- Controlled Shutdown of an Instance
- Error Response and Logging Enhancement

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): 50K 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 50k | SLF Lookup GET Requests | 50K |
| 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 |
|--|---------------------------------|
| UDR Version Tag | 25.1.200 |
| Target TPS | 50K Lookup + 1.44K Provisioning |
| Traffic Profile | SLF 50K Profile |
| Notification Rate | OFF |
| UDR Response Timeout | 900ms |
| Client Timeout | 30s |
| Signaling Requests Latency Recorded on Client | 17ms |
| Provisioning Requests Latency Recorded on Client | 38ms |

The following table describes consolidated resource requirement:

Table 3-3 Consolidated Resource Requirement

| Resource | CPU | Memory | Ephemeral Storage | PVC |
|----------|-----|--------|-------------------|---------|
| cnDBTier | 134 | 453 GB | 20 GB | 1064 GB |
| SLF | 384 | 225 GB | 53 GB | NA |
| ProvGw | 39 | 39 GB | 9 GB | NA |
| Buffer | 50 | 50 GB | 50 GB | 50 GB |
| Total | 607 | 767 GB | 132 GB | 1114 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 | CPU Limit | RAM Limit | PVC |
|---------------------------------------|-----------|-----------|-------|
| Prometheus (snapshot utility enabled) | 4 | 8 GB | 75 GB |

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

| Service | CPU Limit | RAM Limit | PVC |
|-------------------------|-----------|-----------|-------|
| Prometheus AlertManager | 4 | 4 GB | NA |
| Total | 8 | 12 GB | 75 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 (ndbmgmd) | mysqlndbcluster | 2 | 2 CPUs | 12 GB | 1 GB | 16 GB | 6 CPUs 26 GB Ephemeral Storage: 2 GB PVC Allocation: 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 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 (ndbappmysql) | mysqlndbcluster | 7 | 6 CPUs | 4 GB | 1 GB | 10 GB | 63 CPUs 42 GB Ephemeral Storage: 7 GB PVC Allocation: 70 GB | istio-proxy | 3 CPUs | 2 GB | SQL node (Used for Replication) (ndbmssql) | 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 | 4 CPUs | 4 GB | 1 GB | NA | 5 CPUs 5 GB Ephemeral Storage: 1 GB | istio-proxy | 1 CPUs | 1 GB | DB Backup Manager Service | backup-manager-svc | 1 |
| 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 (ndbappmysql) | mysqlndbcluster | 7 | 6 CPUs | 4 GB | 1 GB | 10 GB | 63 CPUs 42 GB Ephemeral Storage: 7 GB PVC Allocation: 70 GB | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| | istio-proxy | | 3 CPUs | 2 GB | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| SQL node (Used for Replication) (ndbmssql) | 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 | 4 CPUs | 4 GB | 1 GB | NA | 5 CPUs 5 GB Ephemeral Storage: 1 GB | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| | istio-proxy | | 1 CPUs | 1 GB | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| DB Backup Manager Service | backup-manager-svc | 1 | 1 CPU | 1 GB | 1 GB | NA | 2 CPUs 2 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) |
|--|--------------------|----------------|------------------------------------|---------------------------------------|---------------------------|------------------------|---|
| (backup-manager-svc) | istio-proxy | | 1 CPUs | 1 GB | | | Ephemeral Storage: 1 GB |
| Replication Service (db-replication-svc) | db-replication-svc | 1 | 2 CPU | 12 GB | 1 GB | 160 GB | 3 CPUs 13 GB |
| | istio-proxy | | 200m CPU | 500 MB | | | Ephemeral Storage: 1 GB PVC Allocation: 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: 17ms):

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 | 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 | | |

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 Ephemeral Storage: 2 GB |
| | istio-proxy | | 1 CPUs | 1 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.

Table 3-7 Provision Gateway Resources (Provisioning Latency: 38 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 10 GB Memory Ephemeral Storage: 2 GB |
| | istio-proxy | | 1 CPUs | 1 GB | | |
| provgw-egress-gateway | egressgateway | 2 | 4 CPUs | 4 GB | 1 GB | 10 CPUs 10 GB Memory Ephemeral Storage: 2 GB |
| | istio-proxy | | 1 CPUs | 1 GB | | |
| provgw-service | provgw-service | 2 | 4 CPUs | 4 GB | 1 GB | 10 CPUs 10 GB Memory |

Table 3-7 (Cont.) Provision Gateway Resources (Provisioning Latency: 38 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 | 1 GB | | Ephemeral Storage: 2 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-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 3-8 Result and Observation

| Parameter | Values |
|---|-------------------------------------|
| TPS Achieved | 50K 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:

- Auto Create
- Diameter S13 Interface
- Subscriber Activity Logging
- 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 | 25.1.200 |
| 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

| Resource | CPUs | Memory | Ephemeral Storage | PVC |
|----------|------|--------|-------------------|---------|
| cnDBTier | 48 | 664 GB | 21 GB | 1893 GB |
| EIR | 215 | 155 GB | 48 GB | NA GB |
| Buffer | 50 | 50 GB | 20 GB | 200 GB |
| Total | 362 | 814 GB | 89 GB | 2015 GB |

The following table describes cnDBTier resources:

Table 3-12 cnDBTier Resources

| Micro service name | Container name | Number of Pods | CPU Allocation Per Pod | Memory Allocation Per Pod (cnDBTier1) | Total Resources (cnDBTier1) |
|--------------------|------------------|----------------|------------------------|---------------------------------------|---|
| Management node | mysqlInDbcluster | 2 | 2 CPUs | 12 GB | 4 CPU 23 GB Memory Ephemeral Storage: 2 GB PVC Allocation: 32 GB |

Table 3-12 (Cont.) cnDBTier Resources

| Micro service name | Container name | Number of Pods | CPU Allocation Per Pod | Memory Allocation Per Pod (cnDBTier1) | Total Resources (cnDBTier1) |
|--|--------------------|----------------|------------------------|---------------------------------------|---|
| Data node | mysqlndbcluster | 4 | 4 CPUs | 20 GB | 16 CPU 80 GB Memory Ephemeral Storage: 4 GB PVC Allocation: 336 GB |
| APP SQL node | mysqlndbcluster | 5 | 4 CPUs | 4 GB | 20 CPU 20 GB Memory Ephemeral Storage: 5 GB PVC Allocation: 50 GB |
| SQL node (Used for Replication) | mysqlndbcluster | 2 | 4 CPUs | 5 GB | 8 CPU 10 GB Memory Ephemeral Storage: 2 GB PVC Allocation: 32 GB |
| DB Monitor Service (db-monitor-svc) | db-monitor-svc | 1 | 4 CPUs | 4 GB | 4 CPUs 4 GB Memory Ephemeral Storage: 1 GB |
| DB Backup Manager Service (backup-manager-svc) | backup-manager-svc | 1 | 100m CPUs | 128 MB | 100m CPUs 128 MB Memory Ephemeral Storage: 1 GB |
| Replication Service (db-replication-svc) | db-replication-svc | 1 | 2 CPU | 2 GB | 2 CPUs 2 MB Memory Ephemeral Storage: 1 GB PVC Allocation: 66 GB |

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 Ephemeral Storage: 4 GB |
| Ingress-gateway-prov | Ingress-gateway-prov | 2 | 4 CPUs | 4 GB | 1 GB | 8 CPUs 8 GB Memory Ephemeral Storage: 2 GB |
| Nudr-dr-service | nudr-dr-service | 3 | 6 CPUs | 4 GB | 1 GB | 18 CPUs 12 GB Memory Ephemeral Storage: 3 GB |
| Nudr-dr-provservice | nudr-dr-provservice | 2 | 4 CPUs | 4 GB | 1 GB | 8 CPUs 8 GB Memory Ephemeral Storage: 2 GB |
| Nudr-diam-gateway | nudr-diam-gateway | 2 | 6 CPUs | 4 GB | 1 GB | 12 CPUs 8 GB Memory Ephemeral Storage: 2 GB |
| Nudr-diameterproxy | nudr-diameterproxy | 8 | 6 CPUs | 4 GB | 1 GB | 48 CPUs 32 GB Memory Ephemeral Storage: 8 GB |
| Nudr-config | nudr-config | 2 | 1 CPUs | 1 GB | 1 GB | 2 CPUs 2 GB Memory Ephemeral Storage: 2 GB |

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-config-server | nudr-config-server | 2 | 1 CPUs | 1 GB | 1 GB | 2 CPU 2 GB Memory Ephemeral Storage: 2 GB |
| Alternate-route | alternate-route | 2 | 1 CPUs | 1 GB | 1 GB | 2 CPU 2 GB Memory Ephemeral Storage: 2 GB |
| Nudr-nrf-client-nfmanagement-service | nrf-client-nfmanagement | 2 | 1 CPUs | 1 GB | 1 GB | 2 CPU 2 GB Memory Ephemeral Storage: 2 GB |
| App-info | app-info | 2 | 1 CPUs | 1 GB | 1 GB | 2 CPU 2 GB Memory Ephemeral Storage: 2 GB |
| Perf-info | perf-info | 2 | 1 CPUs | 1 GB | 1 GB | 2 CPU 2 GB Memory Ephemeral Storage: 2 GB |
| Nudr-dbcrauditor-service | nudr-dbcrauditor-service | 1 | 1 CPUs | 1 GB | 1 GB | 1 CPU 1 GB Memory Ephemeral Storage: 1 GB |

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% |

Table 3-14 (Cont.) Result and Observation

| Parameter | Values |
|--|--------|
| Average EIR processing time (Request and Response) | 9 ms |

3.3 Test Scenario 3: SOAP and Diameter Deployment Model (1M - 10M Subscribers)

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

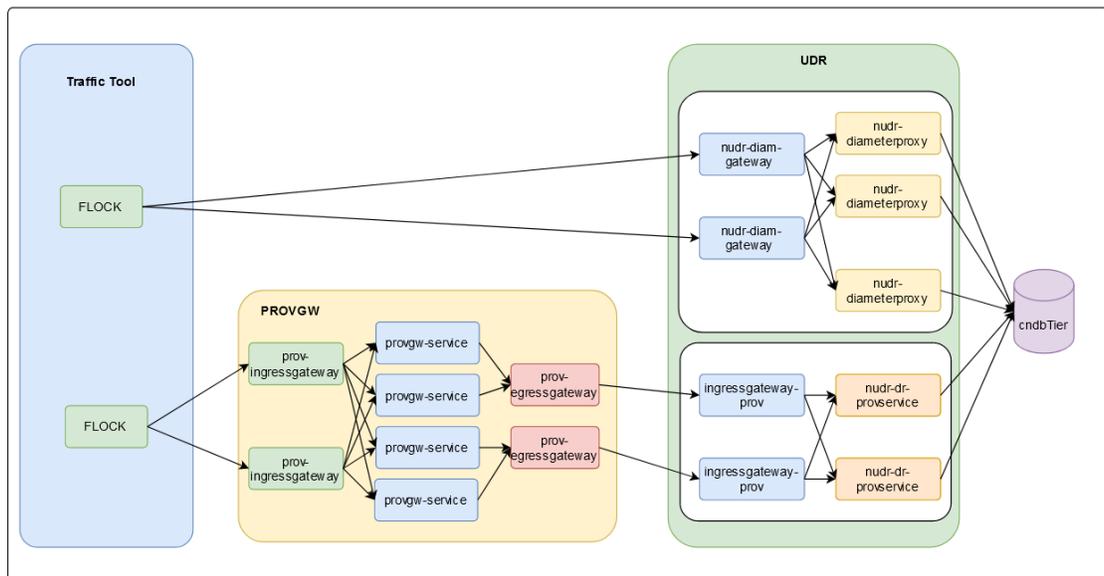
The following features are disabled for this testcase:

- 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 3-1 SOAP and Diameter Deployment Model



The following table describes the benchmarking parameters and their values:

Table 3-15 Traffic Model Details

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

Table 3-16 SOAP Traffic Model

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

Table 3-17 Diameter Traffic Model

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

The following table describes the benchmarking parameters and their values:

Table 3-18 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:

Table 3-19 cnDBTier Resources

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

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:

Table 3-20 UDR Resources (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 |
| Ingress-gateway-prov | ingressgateway-prov | 2 | 2 CPUs | 2 GB | 4 CPUs 4 GB Memory |
| Nudr-dr-service | nudr-dr-service | 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 |

Table 3-20 (Cont.) UDR Resources (Request Latency: 40ms)

| Micro service name | Container name | Number of Pods | CPU Allocation Per Pod | Memory Allocation Per Pod | Total Resources |
|------------------------------|-------------------------|----------------|------------------------|---------------------------|-----------------------|
| 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 |

The following table describes provisioning gateway resources:

Table 3-21 Provisioning Gateway Resources (Provisioning Request Latency: 40ms)

| Micro service name | Container name | Number of Pods | CPU Allocation Per Pod | Memory Allocation Per Pod | Total Resources |
|-----------------------|----------------------|----------------|------------------------|---------------------------|-------------------------|
| provgw-ingress-gatwa | 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 |
| 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 |

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

Table 3-22 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 3-23 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 3-24 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 3-25 Result and Observation

| Parameter | Values |
|--------------|--------|
| TPS Achieved | 10K |

Table 3-25 (Cont.) Result and Observation

| Parameter | Values |
|--|--------|
| Success Rate | 100% |
| Average UDR processing time (Request and Response) | 40ms |

3.4 Test Scenario 4: 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
- 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 3-26 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 3-27 Testcase Parameters

| Input Parameter Details | Configuration Values |
|--|----------------------|
| UDR Version Tag | 25.1.200 |
| Target TPS | 25K TPS Signaling |
| Notification Rate | 300 |
| UDR Response Timeout | 2.7s |
| Signaling Requests Latency Recorded on Client | 25ms |
| Provisioning Requests Latency Recorded on Client | 25ms |

Table 3-28 Consolidated Resource Requirement

| Resource | CPU | Memory | Ephemeral Storage | PVC |
|--------------|-----------------|---------------|-------------------|----------------|
| cnDBTier | 163 CPUs | 640 GB | 29 GB | 1875 GB |
| UDR | 247CPUs | 184 GB | 55 GB | NA |
| PROVGW | 28 | 28 | 10 | NA |
| Buffer | 50 CPUs | 50 GB | 20 GB | 200 GB |
| Total | 488 CPUs | 902 GB | 114 GB | 2075 GB |

The following table describes cnDBTier resources:

Table 3-29 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 (ndbmgmd) | mysqlndbcluster | 2 | 2 CPUs | 9 GB | 1 GB | 15 GB | 4 CPUs 18 GB Ephemeral Storage: 2 GB PVC Allocation: 30 GB |
| Data node (ndbmt) | mysqlndbcluster | 4 | 9 CPUs | 124 GB | 1 GB | 132 GB Backup: 220 GB | 36 CPUs 496 GB Ephemeral Storage: 4 GB PVC Allocation: 1408 GB |

Table 3-29 (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 |
|--|--------------------|----------------|------------------------|---------------------------|---------------------------|------------------------|--|
| APP SQL node (ndbappmysqld) | mysqlndbcluster | 18 | 6 CPUs | 4 GB | 1 GB | 4 GB | 108 CPUs 72 GB Ephemeral Storage: 18 GB PVC Allocation: 72 GB |
| SQL node (ndbmysqld, used for replication) | mysqlndbcluster | 2 | 4 CPUs | 24 GB | 1 GB | 110 GB | 8 CPUs 48 GB Ephemeral Storage: 2 GB PVC Allocation: 220 GB |
| DB Monitor Service (db-monitor-svc) | db-monitor-svc | 1 | 4 CPUs | 4 GB | 1 GB | NA | 4 CPU 4 MB Ephemeral Storage: 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 Ephemeral Storage: 1 GB |
| Replication service (Multi site cases) | replication-svc | 1 | 2 CPUs | 2 GB | 1 GB | 143 GB | 2 CPUs 2 GB Ephemeral Storage: 1 GB PVC Allocation: 143 GB |

The following table describes UDR resources:

Table 3-30 UDR Resources (Average Latency: 50ms (N36/PROV) and N36: 25ms / PROV: 75ms

| 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 Ephemeral Storage: 13 GB |
| Ingress-gateway-prov | ingressgateway-prov | 2 | 4 CPUs | 4 GB | 1 GB | 8 CPUs 8 GB Ephemeral Storage: 2 GB |
| Nudr-dr-service | nudr-dr-service | 20 | 6 CPUs | 4 GB | 1 GB | 120 CPUs 80 GB Ephemeral Storage: 20 GB |
| Nudr-dr-provservice | nudr-dr-provservice | 2 | 4 CPUs | 4 GB | 1 GB | 8 CPUs 8 GB Ephemeral Storage: 2 GB |
| Nudr-notify-service | nudr-notify-service | 3 | 6 CPUs | 5 GB | 1 GB | 18 CPUs 15 GB Ephemeral Storage: 3 GB |
| Nudr-egress-gateway | egressgateway | 2 | 6 CPUs | 4 GB | 1 GB | 12 CPUs 8 GB Ephemeral Storage: 2 GB |
| Nudr-config | nudr-config | 2 | 1 CPU | 1 GB | 1 GB | 2 CPUs 2 GB Ephemeral Storage: 2 GB |
| Nudr-config-server | nudr-config-server | 2 | 1 CPU | 1 GB | 1 GB | 2 CPUs 2 GB Ephemeral Storage: 2 GB |
| Alternate-route | alternate-route | 2 | 1 CPU | 1 GB | 1 GB | 2 CPUs 2 GB Ephemeral Storage: 2 GB |

Table 3-30 (Cont.) UDR Resources (Average Latency: 50ms (N36/PROV) and N36: 25ms / PROV: 75ms)

| Micro service name | Container name | Number of Pods | CPU Allocation Per Pod | Memory Allocation Per Pod | Ephemeral Storage Per Pod | Total Resources |
|--------------------------------------|--------------------------|----------------|------------------------|---------------------------|---------------------------|--|
| Nudr-nrf-client-nfmanagement-service | nrf-client-nfmanagement | 2 | 1 CPU | 1 GB | 1 GB | 2 CPUs 2 GB Ephemeral Storage: 2 GB |
| App-info | app-info | 2 | 1 CPU | 1 GB | 1 GB | 2 CPUs 2 GB Ephemeral Storage: 2 GB |
| Perf-info | perf-info | 2 | 1 CPU | 1 GB | 1 GB | 2 CPUs 2 GB Ephemeral Storage: 2 GB |
| Nudr-dbcrauditor-service | nudr-dbcrauditor-service | 1 | 1 CPU | 1 GB | 1 GB | 1 CPU 1 GB Ephemeral Storage: 1 GB |

The following table describes Provisioning Gateway resources:

Table 3-31 Provisioning Gateway Resources

| Micro service name | Container name | Number of Pods | CPU Allocation Per Pod | Memory Allocation Per Pod | Total Resources |
|-----------------------|----------------|----------------|------------------------|---------------------------|---|
| provgw-ingressgateway | ingressgateway | 2 | 4 CPUs | 4 GB | 8 CPUs 8 GB Ephemeral Storage: 2 GB |
| provgw-egressgateway | egressgateway | 2 | 4 CPUs | 4 GB | 8 CPUs 8 GB Ephemeral Storage: 2 GB |
| provgw-service | provgw-service | 4 | 2 CPUs | 2 GB | 8 CPUs 8 GB Ephemeral Storage: 2 GB |
| provgw-config | provgw-config | 2 | 2 CPUs | 2 GB | 4 CPUs 4 GB Memory Ephemeral Storage: 2 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 3-32 Result and Observation

| Parameter | Values |
|--------------------------------|-----------------------------------|
| TPS Achieved | 25K Signaling + 600K Provisioning |
| UDR Processing Average Latency | 50ms |
| Success rate | 100% |

3.5 Test Scenario 5: 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
- 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 3-33 Traffic Model Details

| Request Type | Details | TPS |
|-----------------|-----------------------|------------|
| N36 17.2K TPS | subs-to-notify POST | 3.6K (20%) |
| | sm-data GET | 3.6K (20%) |
| | subs-to-notify DELETE | 3.6K (20%) |
| | sm-data PATCH | 6.4K (40%) |
| SH 10K TPS | UDR | 4K |
| | PUR | 1.2K |
| | SNR | 4.8K |
| SH PNR 1.2K TPS | PNR | 1.2K |

Table 3-33 (Cont.) Traffic Model Details

| Request Type | Details | TPS |
|---------------------------|----------------------|-----|
| 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 3-34 Testcase Parameters

| Input Parameter Details | Configuration Values |
|---|----------------------|
| UDR Version Tag | 25.1.200 |
| Target TPS | 17.2K N36 + 10K SH |
| Notification Rate | 1.2K |
| UDR Response Timeout | 2.7s |
| Signaling Requests Latency Recorded on Client | 36ms |
| Provisioning Requests Latency Recorded on Client | 36ms |
| Diameter (SH) Requests Latency Recorded on Client | 40ms |

Table 3-35 Consolidated Resource Requirement

| Resource | CPU | Memory | Ephemeral Storage | PVC |
|--------------|-----------------|---------------|-------------------|----------------|
| cnDBTier | 139 CPUs | 625 GB | 25 GB | 1823 |
| UDR | 281 CPUs | 201 GB | 59 GB | NA |
| PROVGW | 28 | 28 | 10 | NA |
| Buffer | 50 CPUs | 50 GB | 20 GB | 200 GB |
| Total | 498 CPUs | 901 GB | 114 GB | 2023 GB |

The following table describes cnDBTier resources:

Table 3-36 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 (ndbmgmd) | mysqlndbcluster | 2 | 2 CPUs | 9 GB | 1 GB | 15 GB | 4 CPUs 18 GB Ephemeral Storage: 2 GB PVC Allocation: 30 GB |
| Data node (ndbmtdd) | mysqlndbcluster | 4 | 9 CPUs | 124 GB | 1 GB | 132 GB Backup: 220 GB | 36 CPUs 496 GB Ephemeral Storage: 4 GB PVC Allocation: 1408 GB |
| APP SQL node (ndbappmysqld) | mysqlndbcluster | 14 | 6 CPUs | 4 GB | 1 GB | 4 GB | 84 CPUs 56 GB Ephemeral Storage: 14 GB PVC Allocation: 20 GB |
| SQL node (ndbmysqld, used for replication) | mysqlndbcluster | 2 | 4 CPUs | 24 GB | 1 GB | 110 GB | 8 CPUs 48 GB Ephemeral Storage: 2 GB PVC Allocation: 220 GB |
| DB Monitor Service | db-monitor-svc | 1 | 4 CPUs | 4 GB | 1 GB | NA | 4 CPU 4 MB Ephemeral Storage: 1 GB |
| DB Backup Manager Service | backup-manager-svc | 1 | 100 millicores CPUs | 128 MB | 1 GB | NA | 1 CPU 128 MB Ephemeral Storage: 1 GB |

Table 3-36 (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 |
|--|-----------------|----------------|------------------------|---------------------------|---------------------------|------------------------|---|
| Replication service (Multi site cases) | replication-svc | 1 | 2 CPUs | 2 GB | 1 GB | 143 GB | 2 CPUs 2 GB Ephemeral Storage: 1 GB PVC Allocation: 143 GB |

The following table describes UDR resources:

Table 3-37 UDR Resources (Average Latency: 47ms [30ms (N36) / 30ms (SH) / 80ms 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 Ephemeral Storage: 9 GB |
| Ingress-gateway-prov | ingressgateway-prov | 2 | 4 CPUs | 4 GB | 1 GB | 8 CPUs 8 GB Ephemeral Storage: 2 GB |
| Nudr-dr-service | nudr-dr-service | 17 | 6 CPUs | 4 GB | 1 GB | 102 CPUs 68 GB Ephemeral Storage: 17 GB |
| Nudr-dr-provservice | nudr-dr-provservice | 2 | 4 CPUs | 4 GB | 1 GB | 8 CPUs 8 GB Ephemeral Storage: 2 GB |
| Nudr-notify-service | nudr-notify-service | 3 | 6 CPUs | 5 GB | 1 GB | 18 CPUs 15 GB Ephemeral Storage: 3 GB |

Table 3-37 (Cont.) UDR Resources (Average Latency: 47ms [30ms (N36) / 30ms (SH) / 80ms Provisioning]).

| 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 Ephemeral Storage: 2 GB |
| Nudr-diam-gateway | nudr-diam-gateway | 2 | 6 CPUs | 5 GB | 1 GB | 12 CPUs 10 GB Ephemeral Storage: 2 GB |
| Nudr-diameterproxy | nudr-diameterproxy | 9 | 6 CPUs | 4 GB | 1 GB | 54 CPUs 36 GB Ephemeral Storage: 9 GB |
| Nudr-config | nudr-config | 2 | 1 CPU | 1 GB | 1 GB | 2 CPUs 2 GB Ephemeral Storage: 2 GB |
| Nudr-config-server | nudr-config-server | 2 | 1 CPU | 1 GB | 1 GB | 2 CPUs 2 GB Ephemeral Storage: 2 GB |
| Alternate-route | alternate-route | 2 | 1 CPU | 1 GB | 1 GB | 2 CPUs 2 GB Ephemeral Storage: 2 GB |
| Nudr-nrf-client-nfmanagement-service | nrf-client-nfmanagement | 2 | 1 CPU | 1 GB | 1 GB | 2 CPUs 2 GB Ephemeral Storage: 2 GB |
| App-info | app-info | 2 | 1 CPU | 1 GB | 1 GB | 2 CPUs 2 GB Ephemeral Storage: 2 GB |
| Perf-info | perf-info | 2 | 1 CPU | 1 GB | 1 GB | 2 CPUs 2 GB Ephemeral Storage: 2 GB |

Table 3-37 (Cont.) UDR Resources (Average Latency: 47ms [30ms (N36) / 30ms (SH) / 80ms Provisioning]).

| Micro service name | Container name | Number of Pods | CPU Allocation Per Pod | Memory Allocation Per Pod | Ephemeral Storage Per Pod | Total Resources |
|--------------------------|--------------------------|----------------|------------------------|---------------------------|---------------------------|---|
| Nudr-dbcrauditor-service | nudr-dbcrauditor-service | 1 | 1 CPU | 1 GB | 1 GB | 1 CPU 1 GB Ephemeral Storage: 1 GB |

The following table describes Provisioning Gateway resources:

Table 3-38 Provisioning Gateway Resources

| Micro service name | Container name | Number of Pods | CPU Allocation Per Pod | Memory Allocation Per Pod | Total Resources |
|-----------------------|----------------|----------------|------------------------|---------------------------|---|
| provgw-ingressgateway | ingressgateway | 2 | 4 CPUs | 4 GB | 8 CPUs 8 GB Ephemeral Storage: 2 GB |
| provgw-egressgateway | egressgateway | 2 | 4 CPUs | 4 GB | 8 CPUs 8 GB Ephemeral Storage: 2 GB |
| provgw-service | provgw-service | 4 | 2 CPUs | 2 GB | 8 CPUs 8 GB Ephemeral Storage: 2 GB |
| provgw-config | provgw-config | 2 | 2 CPUs | 2 GB | 4 CPUs 4 GB Memory Ephemeral Storage: 2 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 3-39 Result and Observation

| Parameter | Values |
|--------------------------------|--|
| TPS Achieved | 17.2K N36 + 10K SH + 600K Provisioning |
| UDR Processing Average Latency | 47ms |
| Success rate | 100% |