Oracle® Communications

Network Analytics Data Director Benchmarking Guide

Release 25.2.200
G49043-01
December 2025

Oracle Communications Network Analytics Data Director Benchmarking Guide, Release 25.2.200

G49043-01

Contents

Preface

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.
italic	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 (MOS)

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 used in the document.

Table Acronyms

Acronym	Description
BSF	Oracle Communications Cloud Native Core, Binding Support Function
CNE	Oracle Communications Cloud Native Core, Cloud Native Environment
vCNE	Virtualized CNE
MPS	Messages Per Second
NDB	Network Data Broker
NRF	Oracle Communications Cloud Native Core, Network Repository Function
NVME	Non Volatile Memory Express
OCI	Oracle Cloud Infrastructure
OCNADD	Oracle Communications Network Analytics Data Director
OCPU	Oracle Compute Unit
PCF	Oracle Communications Cloud Native Core, Policy Control Function
SCP	Oracle Communications Cloud Native Core, Service Communication Proxy
SEPP	Oracle Communications Cloud Native Core, Security Edge Protection Proxy
xDR	Extended Detail Record
LB	Load Balancer
LBVM	Load Balancer Virtual Machine
CNLB	Cloud Native Load Balancer
PVC	Persistent Volume Claim

What's New in This Guide

This section lists the documentation updates for Release 25.2.2xx in *Oracle Communications Network Analytics Data Director Benchmarking Guide*.

Release 25.2.200 - G49043-01, December 2025

- This release introduces a fundamental transformation in the OCNADD architecture. The previously singular Worker Group has been decoupled and re-architected into two specialized functional groups: the Relay Agent Group and the Mediation Group. This architectural shift impacts the entire system workflow and is reflected throughout this document.
- Additionally OCNADD now supports 1500K MPS (1.5M MPS) traffic:
  - Added Resource Profile for 1500K MPS.
  - Added the following scenarios in OCNADD Benchmarking Testing chapter:
    * Performance Benchmarking for 1.5M MPS SCP traffic with Single Synthetic Feed with two endpoints

1 Introduction

Oracle Communications Network Analytics Data Director (OCNADD) is a specialized Network Data Broker (NDB) in 5G Network Architecture.

OCNADD receives network traffic data from various sources, 5G network functions (NFs), Non-5G NFs, and third-party producers, performs filtering, replication, and aggregation on the received data according to the rules implemented by the subscribed third-party consumers. OCNADD then sends the filtered, replicated, or aggregated data to the subscribed third-party consumers (third-party consumer applications or platforms) securely.

1.1 Purpose and Scope

This document provides detailed information about Oracle Communications Network Analytics Data Director (OCNADD) to support the documentation team in creating a customer-facing Planning Guide.

It includes the following:

- **OCNADD deployment overview**
- **OCNADD resource requirements**

1.2 References

For more information about OCNADD, see the following documents:

- *Oracle Communications Network Analytics Suite Release Notes*
- *Oracle Communications Network Analytics Suite Licensing Information User Manual*
- *Oracle Communications Network Analytics Automated Testing Suite Guide*
- *Oracle Communications Network Analytics Suite Security Guide*
- *Oracle Communications Network Analytics Data Director Installation, Upgrade, and Fault Recovery Guide*
- *Oracle Communications Network Analytics Data Director User Guide*
- *Oracle Communications Network Analytics Data Director Outbound Interface Specification Guide*
- *Oracle Communications Network Analytics Data Director Diameter User Guide*
- *Oracle Communications Network Analytics Data Director vCollector Installation Guide*
- *Oracle Communications Network Analytics Data Director Troubleshooting Guide*
- *Oracle Communications Cloud Native Core, cnDBTier Installation,Upgrade, and Fault Recovery Guide*
- *Oracle Communications Cloud Native Configuration Console Installation, Upgrade, and Fault Recovery Guide*
- *Oracle Communications Cloud Native Core, OCI Deployment Guide*

- *Oracle Communication Certificate Manager Installation, Upgrade and Fault Recovery Guide*
- *Oracle Communication Certificate Manager User Guide*

2 Deployment

OCNADD supports CNE deployment. For more information about OCNADD installation, see *Oracle Communications Network Analytics Data Director Installation, Upgrade, and Fault Recovery Guide*.

The following diagram depicts the OCNADD deployment in the 5G architecture:

Figure 2-1 OCNADD Deployment

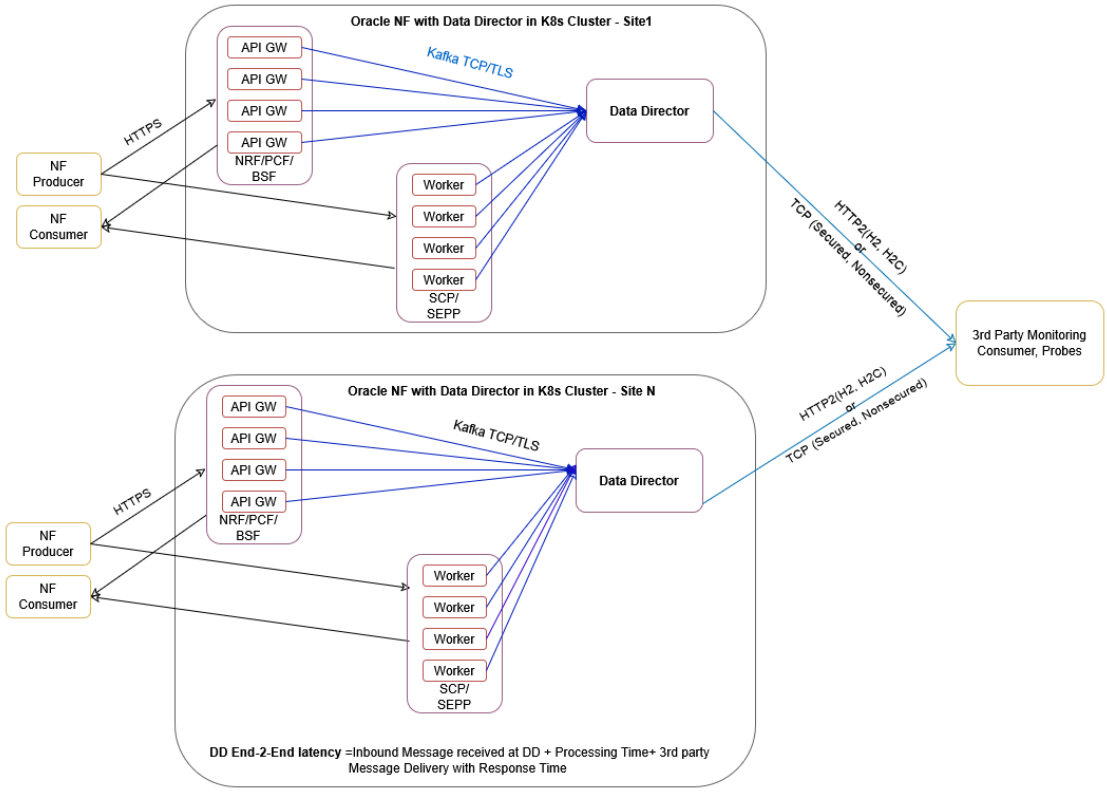

OCNADD uses the following common services of CNE:

- Kubernetes
- Prometheus
- CNLB / Metallb (Load balancer)
- cnDBTier

3
Resource Requirement

This chapter provides information about the resource requirements to install and run Oracle Communications Network Analytics Data Director (OCNADD) with the desired Message Per Second (MPS) profiles.

Cluster Details

The following tables provides information about the types of servers and the number of servers used in the test environment:

Table 3-1 Test Bed 1 - CNE on BareMetal

Type of Server	X9 Server and NVME
Master node	3
Worker node	46 (Note: During the test, nodes utilised for DD are 26. It can vary based on cluster performance)
Storage Class	Standard
LB Type	CNLB (with 3rd party in same cluster)
Network Interface Bandwidth	20Gbps Full duplex

Resource Requirements for OCI Environment

- OCI block volume is attached to the PVC with auto-tune based performance from balanced to high performance. To change block volume to auto-tune based performance (Balance to High Performance), see Changing the Performance of a Volume.
- All tests are performed with the default round-robin based ordering.
- Resource requirements may vary after enabling key or custom based ordering and running traffic with actual NFs.

Table 3-2 Test Bed 2 - OKE on OCI

Type of Server	OCI Hardware
Worker nodes	6
Instance Shape	VM.Standard.E4.Flex
OCPUs in worker node	50 (CPU: 100)
Memory in worker node	194 GB

3.1 Profile Resource Requirements

This section provides information about resource requirements to install and run the OCNADD with desired MPS profiles.

> **Note**
>
> It is recommended to have the following configurations for CNE Baremetal/vCNE setups to achieve the required throughput:
>
> - Jumbo frames should be enabled.
> - Ring buffer size should be increased to avoid packet drops at interfaces (not applicable for vCNE).
> - FluentD pods should not be in a `CrashLoopBackOff` state due to Out of Memory errors. For more information, see the *High Latency in Adapter Feeds Due to High Disk Latency* section in the *Oracle Communications Network Analytics Data Director Troubleshooting Guide*.
> - The benchmark tests were performed with a round-trip latency of up to 5 ms from third-party consumer applications. If the latency is greater than 5 ms, the resource profile footprint and the end-to-end latency will be higher.
> - Up to 95% of traffic is expected to complete within 100 ms latency; however, actual results depend on cluster performance. When the cluster is well-provisioned and operating efficiently, the percentage of traffic under 100 ms may increase. If the cluster is constrained or under heavy load, the percentage may decrease.

3.1.1 Resource Profile for Database

This section provides information about the database profile resource requirements to install and run Oracle Communications Network Analytics Data Director (OCNADD) with the desired Message Per Second (MPS) profiles.

Table 3-3 Resource Requirement

cnDBTier Pods	Min vCPU	Max vCPU	Min Memory	Max Memory	Total Replica
SQL (ndbmysqld) **Kubernetes Resource Type:** StatefulSet	1	1	1Gi	1Gi	2
SQL (ndbappmysqld) **Kubernetes Resource Type:** Statefulset	1	1	1Gi	1Gi	2
MGMT (ndbmgmd) **Kubernetes Resource Type:** StatefulSet	1	1	1Gi	1Gi	2
Database (ndbmtd) **Kubernetes Resource Type:** StatefulSet	1	1	4Gi	4Gi	2
Backup Manager Service (db-backup-manager-svc) **Kubernetes Resource Type:** Deployment	0.1	0.1	128Mi	128Mi	1
Monitor Service (db-monitor-svc) **Kubernetes Resource Type:** Deployment	0.2	0.2	500Mi	500Mi	1

Table 3-3 (Cont.) Resource Requirement

cnDBTier Pods	Min vCPU	Max vCPU	Min Memory	Max Memory	Total Replica
EXTENDED STORAGE is ENABLED in CORRELATION Feed(Per Correlation Feed) Rate Supported in current release: 1K MPS rate with 24 hours retention Update "global.ndb.datamemory=96G" in custom-value.yaml of cndbTier PVC of ndbmtd= 150GB					
Database (ndbmtd) **Kubernetes Resource Type:** StatefulSet	8	8	128Gi	128Gi	4

> **Note**
>
> Configure "**datamemory: 1G**" under "**ndbmtd"** section while deploying the CnDbTier for OCNADD. For more details on cnDBTier resource profile, see "cnDBTier Small Profile" section in *cnDBTier Resource Models Guide*.

3.1.2 Resource Profile for 1500K MPS

3.1.2.1 Resource Profile for OCNADD OAM Services

The following profile is used for management group services in all the performance scenarios.

Table 3-4 Resource Requirement

Service Name	vCPU	Memory Required (Gi)	Min Replica	Max Replica	Description
ocnaddconfiguration	1	1	1	1	-
ocnaddalarm	1	1	1	1	-
ocnaddhealthmonitoring	1	1	1	1	-
ocnaddgui	2	1	1	1	-
ocnadduirouter	2	1	1	2	-
ocnaddexport	0.5	1	1	1	Resource requirement will increase when export is configured. It can be disabled in case export is not required.
ocnaddredundancyagent	1	1	1	1	Required only when Geo Redundancy is enabled for OCNADD
ocnaddmanagementgateway	1	1	1	2	-

3.1.2.2 Resource Profile for OCNADD Worker Group Services

The following profile shall be used for worker group services. The resource profile for worker group services will vary based on the scenario to be executed.

> **Note**
>
> To support the increased throughput, first the number of Kafka instances must be increased followed by the number of topic partition changes based on the recommended MPS profile. For more details on this, see "Adding Partitions to an Existing Topic" in the *Oracle Communications Network Analytics Data Director User Guide*.

> **Note**
>
> After changing the topic partitions, perform the partition reassignment steps to redistribute the existing partitions across the new Kafka broker instances. For more details on partition reassignment, refer to the *Partitions Reassignment in Kafka Cluster* section of the *Oracle Communications Network Analytics Data Director User Guide*.

3.1.3 Egress 1500K MPS

> **Note**
>
> Test conducted using separate clusters: One is dedicated to OCNADD services, while the other is shared by SCP NF services and third-party consumer.

- Source Topic Replication Factor: 1
- MAIN Topic Replication Factor: 1
- Offset Topic Replication Factor: 2
- Message Size: 3500
- Feed Type: 1-TCP feeds
- FILTER: Egress filter OFF
- Message Sequencing/Metadata: OFF
- Test bed: CNLB-CNE Bare Metal Cluster Environment
- 1500K MPS SCP Profile(Ingress to DD)
- Storage Type: Volatile RAM Drive for Relay Agent and Mediation Kafka clusters

1500K MPS SCP Profile(Ingress to DD)

Table 3-5 1500K MPS SCP Profile (Ingress to DD)

Service	vCPU	Memory Required (Gi)	Min Replica	Max Replica	Topic Partitions	Topic Retention Time
Relay Agent Services						
kraft-controller	1	2	3	3	-	
ocnaddkafka (kafkaBroker)	6	160	20	20	-	
ocnaddscpaggregation	3	8	57	57	SCP = 342 (Each instance 6 partition)	5 minutes
ocnaddrelayagentgateway	1	1	1	2	-	
Mediation Services						
kraft-controller	1	2	3	3	-	
ocnaddkafka (kafkaBroker)	7	100	20	20	-	
ocnaddadapter-1(TCP) (consumeradapter)	6	6	59	59	MAIN=354 (Each instance 6 partition)	5 minutes
ocnaddadminservice	1	1	1	1	-	
ocnaddmediationgateway	1	1	1	2	-	

Note

- Additional memory and/or replicas are required for the aggregation service if the **Metadata Cache** feature is enabled and the values of the properties `METADATA_MAP_CACHE_EXPIRY_TIME_MS` and `METADATA_MAP_CACHE_SCHEDULER_TIME_MS` are increased to higher values.
- The end-to-end latency may increase based on:
  - Higher values of `METADATA_MAP_CACHE_EXPIRY_TIME_MS` and `METADATA_MAP_CACHE_SCHEDULER_TIME_MS`.
  - `Timer Expiry Value + Processing Time + RF2/RF1 Processing Time + Third-party Response Time (for HTTP2 feed).`
- Resource requirements may vary for the **consumeradapter** service based on the % of data allowed after filtering and the number of filter conditions with it's values used in the filter.
- It is recommended to configure a **replication factor (RF) of 2** for Kafka internal topics (**offsetTopicReplicationFactor** and **transactionStateReplicationFactor**) to improve cluster durability. This configuration enhances data availability and resilience.
- The profile mentioned here is with a retention time of 5 minutes. Increasing the retention time for the source NF topic or the MAIN topic will increase memory requirements.
- Setting a very low retention time (under 5 minutes) can increase Kafka's bookkeeping overhead, potentially degrading overall system performance. Hence it is not recommended to configure very low retention time less than 5 minutes.
- Depending on cluster performance, more instances of the KafkaBroker may be required when running with **RF=2**, and end-to-end latency may also increase if disk I/O is slow.
- For DISK I/O, see Disk Throughput Requirements.
- For Kafka PVC-Storage, see Kafka PVC Storage Requirements.

3.1.4 Default Deployment Profile

This profile can stream NFs (SCP, NRF, SEPP, PCF, BSF) data up to **15K MPS** and can be scaled to handle up to **100K MPS** when **weighted_lb** and **Filter (Ingress and Egress)** are disabled.

The replication factor should be set to **1**, and the incoming message size on OCNADD should be **less than or equal to 3500 bytes**.

- **Replication Factor:** 1
- **Message Size:** 3500
- **Feed Type:** HTTP2, Synthetic

Table 3-6 Default Deployment Profile

OCNADD Service	vCPU Req	vCPU Limit	Memory Req (Gi)	Memory Limit (Gi)	Min Replica	Max Replica	Partitions	Topic Name
Management Group Services								
ocnaddconfiguration	1	1	1	1	1	1		
ocnaddalarm	1	1	1	1	1	1		
ocnaddhealthmonitoring	1	1	1	1	1	1		
ocnaddgui	2	2	1	1	1	1		
ocnadduirouter	2	2	1	1	1	2		
ocnaddmanagementgateway	1	1	1	1	1	2		
Relay Agent Group Services								
kraft-controller	1	1	2	2	3	3		
ocnaddkafka (kafkaBroker)	5	5	150	150	4	4		
ocnaddscpaggregation (55K)	3	3	4	4	1	4	24	SCP
ocnaddnrfaggregation (15K)	2	3	2	2	1	1	6	NRF
ocnaddseppaggregation (30k)	2	3	4	4	1	2	12	SEPP
ocnaddbsfaggregation(9k)	2	3	4	4	1	1	6	BSF
ocnaddpcfaggregation(30k)	2	3	4	4	1	2	12	PCF
ocnaddrelayagentgateway	1	1	1	1	1	2		
Mediation Group Services								
kraft-controller	1	1	2	2	3	3		
ocnaddkafka (kafkaBroker)	5	5	48	48	4	4		
ocnaddadapter (consumeradapter)	3	3	HTTP2: 4 SYNTHETIC: 4	HTTP2: 24 SYNTHETIC: 6	HTTP2: 2 SYNTHETIC: 1	HTTP2: 13 SYNTHETIC: 9	117	MAIN
ocnaddadmin	1	1	1	1	1	1		
ocnaddmediationgateway	1	1	1	1	1	2		

Note:

- Four instances of Kafka broker might require while running RF=2 based on setup performance and end-2-end latency might get increased when DISK I/O is slow,
- For DISK I/O, see Disk Throughput Requirements.
- For Kafka PVC-Storage, see Kafka PVC Storage Requirements.

3.2 Pod Affinity (or Anti-affinity) Rules

In the Data Director, support for node affinity has been added. The rules are currently defined for the services mentioned in the table below. The rules are currently disabled; however, the user can enable them for the supported services. The rules are provided to control the deployment of certain traffic processing services on a particular set of identified nodes.

Relay Agent Group:

- Relay Agent Kafka Brokers
- ocnaddnrfaggregation
- ocnaddscpaggregation
- ocnaddseppaggregation
- ocnaddbsfaggregation
- ocnaddpcfaggregation

Mediation Group

- Mediation Kafka Brokers
- Consumer Adapter

Node Affinity Rules

Step 1: Update the **affinity** section in the `ocnadd-<ocnadd-group>-custom-values.yaml` file

```
affinity: {}
                # Node Affinity Configuration:
                #
                # To enable node affinity, remove the empty curly braces ({})
above and un-comment the nodeAffinity section below.
                # This allows you to specify rules for scheduling pods on
specific nodes.
                #
                # Example Configuration:
                ####################################
                # nodeAffinity:
                #   requiredDuringSchedulingIgnoredDuringExecution:
                #     nodeSelectorTerms:
                #     - matchExpressions:
                #       - key: kubernetes.io/hostname
                #         operator: NotIn
                #         values:
                #         - k8s-node-26
                #         - k8s-node-24
                #       - key: kubernetes.io/hostname
                #         operator: In
                #         values:
                #         - k8s-node-2
                #         - k8s-node-3
                ########################################
                # Explanation:
                #
                # - The 'NotIn' expression prevents pods from being scheduled
```

```
on nodes k8s-node-26 and k8s-node-24.
                # - The 'In' expression ensures pods are scheduled on nodes
k8s-node-2 and k8s-node-3.
                #
                # To customize, modify the 'key', 'operator', and 'values'
fields according to your needs.
                # You can add or remove 'matchExpressions' to create more
complex scheduling rules.
                #
                # Remember to remove the empty 'affinity: {}' and un-comment
the desired nodeAffinity configuration to enable it.
```

Step 2: Helm upgrade the corresponding OCNADD group (relayagent or mediation)

```
helm upgrade <release-name> -f ocnadd-common-custom-values.yaml -f ocnadd-
<ocnadd-group>-custom-values.yaml --namespace <release-namespace> <release-
helm-chart>
```

Where:

- `<release-name>` is the release name of the source release deployment
- `ocnadd-<ocnadd-group>-custom-values.yaml` is the custom values file created for relayagent or mediation group
- `<release-namespace>` is the OCNADD namespace of the source release
- `<release-helm-chart>` is the location of the Helm chart of the target release

Examples:

To upgrade relay agent group:

```
helm upgrade dd-rea -f ocnadd-common-custom-values.yaml -f ocnadd-relayagent-
custom-values.yaml --namespace ocnadd-rea ocnadd
```

To upgrade mediation group:

```
helm upgrade dd-med -f ocnadd-common-custom-values.yaml -f ocnadd-mediation-
custom-values.yaml --namespace ocnadd-med ocnadd
```

Step 3: Verify that the PODs of the modified services have been deployed as per the configured affinity rules

3.3 Ephemeral Storage Requirements

The following table describes the Ephemeral Storage requirements for OCNADD:

Table 3-7 Ephemeral Storage Requirements

Service Name	Ephemeral Storage (Request) in Mi	Ephemeral Storage (Limit) in Mi	Description
OAM Services			
ocnaddalarm	100	500	-
ocnaddhealthmonitoring	100	500	-
ocnaddconfiguration	100	500	-
ocnadduirouter	500	500	-
ocnaddexport	1000	2000	-
ocnaddredundancyagent	100	500	Required only when Geo Redundancy is enabled for OCNADD
ocnaddmanagementgateway	100	500	-
Relay Agent Services			
ocnaddscpaggregation	100	500	-
ocnaddseppaggregation	100	500	-
ocnaddnrfaggregation	100	500	-
ocnaddbsfaggregation	100	500	-
ocnaddpcfaggregation	100	500	-
ocnaddrelayagentgateway	100	500	-
Relay Agent Services			
<app-name>-adapter (consumeradapter)	1000	1000	-
ocnaddcorrelation	400	800	-
ocnaddstorageadapter	400	800	-
ocnaddingressadapter	400	800	-
ocnaddfilter	100	800	Required only when "*Filtered*" or "*Correlated Filtered*" feed is created
ocnaddadminservice	200	200	-
ocnaddmediationgateway	100	500	-

3.4 Disk Throughput Requirements

The following table describes the disk throughput requirements in OCNADD:

Table 3-8 Disk Throughput Requirements

Avg Size (in Bytes)	Rate	RF (Kafka Replication Factor)	Topic (NF+ MAIN)	Consumer Feed	Total Write Throughput (MB/s)	Total Read Throughput (MB/s)	No. of Broker	Per Broker Write Throughput (MB/s)	Per Broker Read Throughput (MB/s)	Total per Broker Throughput (MB/s) with 10% buffer	Total Disk Throughput (MB/s) for the Cluster with 10% Buffer
1941	39000	1	2	1	145	145	3	54	54	108	324
1941	39000	2	2	1	289	289	3	106	106	212	636
3769	39000	1	2	1	281	281	3	104	104	208	624
3769	39000	2	2	1	561	561	3	206	206	412	1236

Note

- The average size of OCNADD Ingress message captured in the table includes the size of `metadata list + header list` of original `5G HTTP2 header frame + 5G-SBI-Message`.
- Currently, it is recommended to set the Replication Factor (RF) value to **1** with the assumption that the underlying storage provides data redundancy. RF value of "2" will be supported in a future release.

The disk throughput calculations are as follows:

```
Writes: W * RF * T
Reads: ((RF*T)+C- 1) * W
Disk Throughput (Write + Read): (W * RF *T) + (L * W)
W -> MB/sec of data that will be written
RF -> Replication factor
T   -> No of topics to which data copied. As of now, each message will be
copied into two topics.
C   -> Number of consumer groups, that is the number of readers for each
write
L   -> (RF*T) + C -1
```

Average Message in Table:

```
Average Message Size= (a1b1+a2b2+..+a(n)b(n))/(a1+a2+..+a(n))
a1   -> SCP MPS
b1   -> SCP message size
a2   -> NRF MPS
b2   -> NRF message size
```

```
a(n) -> NF(n) MPS
b(n) -> NF(n) message size
```

Example:

Average message size for row 1 = ((1624*30000)+(3000*9000))/(30000+9000) = 1941 Bytes (approx)

Average message size for row 4 = ((4000*30000)+(3000*9000))/(30000+9000) = 3769 Bytes (approx)

The following table describes the disk throughput for SCP, NRF and SEPP:

Table 3-9 SCP, NRF, and SEPP Disk Throughput

SCP Message		NRF Message	SEPP Message		RF (Kafka Replication Factor)	Topic (NF+ MAIN)	Consumer Feed	Total Write Throughput (MB/s)	Total Read Throughput (MB/s)	No.of Broker	Per Broker Write Throughput (MB/s)	Per Broker Read Throughput (MB/s)	Total per Broker Throughput (MB/s) with 10% buffer	Total Disk Throughput (MB/s) for Cluster with 10% Buffer	
Avg Size (Bytes)	Rate	Avg Size (Bytes)	Rate	Avg Size (Bytes)											Rate
1624	30000	3000	9000	3000	15000	1	2	1	145	145	3	54	54	108	324
1624	30000	3000	9000	3000	15000	2	2	1	289	289	3	106	106	212	636
4000	30000	3000	9000	3000	15000	1	2	1	281	281	3	104	104	208	624
4000	30000	3000	9000	3000	15000	2	2	1	561	561	3	206	206	412	1236

> **Note**
>
> - The average size of OCNADD Ingress message captured in the table includes the size of `metadata list + header list` of original `5G HTTP2 header frame + 5G-SBI-Message`.
> - Currently, it is recommended to set the Replication Factor (RF) value to **1** with the assumption that the underlying storage provides data redundancy.

3.5 Kafka PVC Storage Requirements

> **Note**
>
> PVC is created only when Persistence (Disk) Kafka storage mode is configured. It is not created for Volatile (RAM Drive) Kafka storage mode.

The following table describes the retention period per topic for different NFs:

Table 3-10 Retention Period Per Topic

Topic Name	Retention Period
SCP	5 Minutes
NRF	5 Minutes
SEPP	5 Minutes
BSF	5 Minutes
PCF	5 Minutes
MAIN	6 Hours (Max) **Note**: Not applicable when RAM drive storage mode is enabled.

Storage Requirement for a topic:

Storage = MPS × Retention Period × RF × Average Message Size

Where:

- **MPS** = Message Per Second
- **RF** = Replication Factor

Example 1
Average Message Size = 1941 Bytes

Storage Requirement for SCP and NRF Topics

```
MPS * Retention * RF * Message Size
= 39000 * 5 minutes * 3 * 1941
= 39000 * 5 * 60 * 3 * 1941
≈ 63.45 GB
```

Storage Requirement for MAIN

```
= 39000 * 6 hours * 3 * 1941
= 39000 * 6 * 60 * 60 * 3 * 1941
≈ 4.46 TB
```

Total Storage Requirement for the Broker Cluster

```
= 63.45 GB + 4.46 TB
≈ 4.53 TB
```

Storage per Broker (3-broker cluster)

```
= 4.53 TB / 3
≈ 1.51 TB per broker
```

Example 2
Average Message Size = 3769 Bytes

Storage Requirement for SCP and NRF Topics

```
= 39000 * 5 minutes * 3 * 3769
= 39000 * 5 * 60 * 3 * 3769
≈ 123.20 GB
```

Storage Requirement for MAIN

```
= 39000 * 6 hours * 3 * 3769
= 39000 * 6 * 60 * 60 * 3 * 3769
≈ 8.66 TB
```

Total Storage Requirement for the Broker Cluster

```
= 123.20 GB + 8.66 TB
≈ 8.79 TB
```

Storage per Broker (3-broker cluster)

```
= 8.79 TB / 3
≈ 2.93 TB per broker
```

4 OCNADD Benchmarking Testing

This section describes the performance testing scenarios and results for the Message Feed functionality provided by Oracle Communications Network Analytics Data Director (OCNADD). The message feed feature is tested with SCP, NRF, SEPP, BSF, and PCF as the source of the message feed.

4.1 Bare Metal Cluster Environment

This section gives the details of the performance tests done in Bare Metal Cluster Environment.

4.1.1 Performance Benchmarking for 1.5M MPS SCP traffic with Single Synthetic Feed with two endpoints

This performance benchmarking evaluates the deployment of OCNADD handling a SCP traffic of 1.5M MPS. The benchmark utilizes single synthetic (TCP) feed, with tow endpoints.

The benchmarking results are provided below:

> **Note**
>
> One Ingress message from a NF is "1" MPS for OCNADD.

Benchmark Test Environment

- **Execution Time:** 24+ hours
- **Call Rate:** 1.5M MPS
- **Call Mix:** SCP Model-C + Model-D 20%

Setup Details

- **Environment:** OCCNE 25.1.200
- **LB Type**: Bare metal CNLB
- **Cluster Topology:**
  - **DD:** Deployed in Cluster-1
  - **NF:** Within the same cluster as DD
  - **Third-Party Consumers:** Within the same cluster as DD
- **Component Versions:**
  - **OCNADD:** 25.2.200
  - **cnDBTier:** 25.2.101
  - **SCP Version:** 25.2.100

  - **Third-Party Consumers:** Two endpoints (configured in round-robin)

Configuration Details

- **SCP:** `messageCopy` enabled
- **Security:** SASL/SSL enabled between SCP and OCNADD
- **OCNADD:**
  - Replication Factor: 1
  - Kafka RAM Storage(Relay): 160 GB
  - Kafka RAM Storage(Mediation): 100 GB
- **Data Feed Configuration:**
  - **TCP Feed (1):** Ingress 1.5M and Egress 1.5M
- **Message Size:** Approximately 3000–4000 bytes

Benchmark Testcase Specifications
The testcase parameters are as follows:

- **OCNADD**
  - A feed is configured using the OCNADD Console with SCP NF.
  - **Message Ingestion Rate:** 1.5M MPS

Resource Specifications:

Table 4-1 Resource Specifications

Services	Min CPU Per Pod	Max CPU Per Pod	Min Memory Per Pod (Gi)	Max Memory Per Pod (Gi)	Min Replicas	Max Replicas	Parition s/ Retenti on	Topics
ocnaddconfiguratio n	1	1	1	1	1	1	-	-
ocnaddalarm	1	1	1	1	1	1	-	-
ocnaddhealthmonit oring	1	1	1	1	1	1	-	-
ocnaddgui	1	1	1	1	1	1	-	-
ocnadduirouter	1	1	1	1	1	2	-	-
ocnaddmanageme ntgateway	1	1	1	1	1	2	-	-
Relay Agent Services								
kraft-controller	1	1	2	2	3	3	-	-
ocnaddkafka (kafkaBroker)	6	6	160	160	20	20	-	-
ocnaddscpaggrega tion	3	3	8	8	57	57	342	SCP
ocnaddrelayagentg ateway	1	1	1	1	1	2	-	-
Mediation Services								
kraft-controller	1	1	2	2	3	3	-	-
ocnaddkafka (kafkaBroker)	7	7	100	100	20	20	-	-

Table 4-1 (Cont.) Resource Specifications

Services	Min CPU Per Pod	Max CPU Per Pod	Min Memory Per Pod (Gi)	Max Memory Per Pod (Gi)	Min Replicas	Max Replicas	Parition s/ Retenti on	Topics
ocnaddadapter-1(TCP) (consumeradapter)	6	6	6	6	59	59	354	MAIN
ocnaddadminservice	1	1	1	1	1	1	-	-
ocnaddmediationgateway	1	1	1	1	1	2	-	-

- **SCP**: SCP Traffic is 750K TPS using two trigger points.

Benchmark Test Results

Traffic Feed Details

Table 4-2 Traffic Feed Details

NF	NF Traffic Copy Trigger Points	Traffic Rate	Duration in hours	E2E Traffic Feed Average Latency	E2E Traffic Feed Success Rate
OCNADD	NA	1.5M MPS	24h+	ocnaddadapter-1(TCP) 46.4 ms	99.97%
SCP	Ingress Gateway (Request) Egress Gateway (Request)	500K TPS	24h+	NA	NA

CPU and Memory Utilization

The following table describes the OCNADD CPU and memory utilization:

Table 4-3 OCNADD CPU and Memory Utilization

Micro-Service/ Container	CPU Utilization (%)	Memory Utilization (%)	Total Replica
ocnaddconfiguration	0.16%	40.63%	1
ocnaddalarm	0.08%	45.31%	1
ocnaddhealthmonitoring	0.28%	32.03%	1
ocnaddgui	0.00%	3.48%	1
ocnadduirouter	0.07%	17.87%	1
ocnaddmanagementgateway	0.15%	21.78%	1
Relay Agent Services			
kraft-controller	0.8%	59.8%	3

Table 4-3 (Cont.) OCNADD CPU and Memory Utilization

Micro-Service/ Container	CPU Utilization (%)	Memory Utilization (%)	Total Replica
ocnaddkafka (kafkaBroker)	48.9%	57.4%	20
ocnaddscpaggregation	52.7%	15.1%	57
ocnaddrelayagentgateway	0.2%	26.4%	1
Mediation Services			
kraft-controller	0.77%	57.17%	3
ocnaddkafka (kafkaBroker)	53.44%	52.66%	20
ocnaddadapter-1(TCP) (consumeradapter)	66.83%	21.47%	59
ocnaddadminservice	0.12%	25.98%	1
ocnaddmediationgateway	0.10%	22.17%	1