

Oracle® Communications

Network Analytics Data Director

Benchmarking Guide



Release 23.2.0

F83157-04

January 2024

ORACLE®

Copyright © 2023, 2024, 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	
3	Resource Requirement	
	3.1 Profile Resource Requirements	1
	3.2 Deployment Profiles	9
	3.3 Ephemeral Storage Requirements	12
	3.4 Disk Throughput Requirements	13
	3.5 Kafka PVC Storage Requirements	15
4	OCNADD Benchmarking Testing	
	4.1 Performance Benchmarking with HTTP2 Feed	1
	4.2 Performance Benchmarking with Synthetic Feed	4

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
OCNADD	Oracle Communications Network Analytics Data Director
SCP	Oracle Communications Cloud Native Core, Service Communication Proxy
SEPP	Oracle Communications Cloud Native Core, Security Edge Protection Proxy
NRF	Oracle Communications Cloud Native Core, Network Repository Function
CNE	Oracle Communications Cloud Native Core, Cloud Native Environment
MPS	Messages Per Second

What's New in This Guide

This section lists the documentation updates for Release 23.2.x in *Oracle Communications Network Analytics Data Director Benchmarking Guide*.

Release 23.2.0.0.3 - F83157-04, January 2024

There are no updates to this document in this release.

Release 23.2.0.0.2 - F83157-03, September 2023

There are no updates to this document in this release.

Release 23.2.0.0.1 - F83157-02, July 2023

There are no updates to this document in this release.

Release 23.2.0 - F83157-01, July 2023

Updated the following sections:

- [Resource Requirement](#)
- [Profile Resource Requirements](#)
- [Deployment Profiles](#)
- [OCNADD Benchmarking Testing](#)

Added the following sections:

- [Performance Benchmarking with HTTP2 Feed](#)
- [Performance Benchmarking with Synthetic Feed](#)

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 is designed to measure the performance and capacity of OCNADD deployment and resource requirements.

1.2 References

For more information about OCNADD, see the following documents:

- *Oracle Communications Network Analytics Data Director Installation, Upgrade, and Fault Recovery Guide*
- *Oracle Communications Network Analytics Data Director User Guide*


2

Deployment

OCNADD supports CNE and TANZU platforms. There are a few variations in the deployment process for both the platforms. 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



Note

From release 23.1.0 onwards OCNADD supports Synthetic Feeds along with HTTP2 Feeds.

OCNADD uses the following common services of CNE:

- Kubernetes
- Prometheus
- Metallb (Load balancer)

- cnDBTier

3

Resource Requirement

This section 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.

The resource requirement for the profiles is captured using Baremetal CNE with the following CNE and K8 cluster configuration:

Table 3-1 CNE Configuration for HTTP2 Traffic

Type of Server	Baremetal CNE
Master node	3
Worker node	21
Storage Class	Standard
Top of Rack Switch	1

Table 3-2 CNE Configuration for Synthetic Traffic

Type of Server	Baremetal CNE
Master node	3
Worker node	20
Storage Class	Standard
Top of Rack Switch	1

3.1 Profile Resource Requirements

Note

For information on increasing partitions in Kafka topics see, "*Adding Partitions to an Existing Topic*" section in *Oracle Communications Network Analytics Data Director User Guide*.

The following table displays the profile resource requirements when **HTTP2 FEED** is "65K" and **WEIGHTED_LB** is "OFF":

Table 3-3 Resource Requirements

OCNADD Service	9K NRF Profile				41K SCP Profile				15K SEPP Profile				65K Profile (SCP:41K NRF:9K SEPP:15K)				65K SCP Profile			
Replication Factor = 1 Message Size = 3500 FEED Type = HTTP2	vCPU	Total Replicas	Memory Required (Gi)	Topic Partitions	vCPU	Total Replicas	Memory Required (Gi)	Topic Partitions	vCPU	Total Replicas	Memory Required (Gi)	Topic Partitions	vCPU	Total Replicas	Memory Required (Gi)	Topic Partitions	vCPU	Total Replicas	Memory Required (Gi)	Topic Partitions
ocnadd configuration	1	1	1		1	1	1		1	1	1		1	1	1		1	1	1	
ocnadd alarm	1	1	1		1	1	1		1	1	1		1	1	1		1	1	1	
ocnadd admin	1	1	1		1	1	1		1	1	1		1	1	1		1	1	1	
ocnadd healthmonitoring	1	1	1		1	1	1		1	1	1		1	1	1		1	1	1	
ocnadd scpaggregation					2	2	2	SCP=12					2	2	2	SCP=12 (6 partitions in each instance)	2	2	2	SCP=12 (6 partitions in each instance)
ocnadd nrfaggregation	2	1	2	NRF=6									2	1	2	NRF=6 (6 partitions in each instance)				

Table 3-3 (Cont.) Resource Requirements

OCNADD Service	9K NRF Profile				41K SCP Profile				15K SEPP Profile				65K Profile (SCP:41K NRF:9K SEPP:15K)				65K SCP Profile			
ocnadd seppaggregation									2	1	2	SEPP=6	2	1	2	SEPP=6 (6 partitions in each instance)				
ocnadd adapter	3	2	3	MAN=18	3	4	3	MAN=36	3	2	3	MAN=18	3	8	4	MAN=72 (9 partitions in each instance)	3	8	4	MAN=72 (9 partitions in each instance)
ocnadd kafka	2	3	12		3	3	24		3	3	12		5	3	48		5	3	48	
zookeeper	1	3	1		1	3	1		1	3	1		1	3	2		1	3	2	
ocnadd gui	2	1	1		2	1	1		2	1	1		2	1	1		2	1	1	
ocnadd uirouter	2	1	1		2	1	1		2	1	1		2	1	1		2	1	1	
ocnadd cache	1	2	10		1	2	10		1	2	10		1	2	10		1	2	10	

Note

- Four instances of Kafka brokers might be required when running RF=2, based on the setup performance. The end-to-end latency might increase when DISK I/O is slow.
- For DISK I/O, see [Disk Throughput Requirements](#).
- For Kafka PVC-Storage, see [Kafka PVC Storage Requirements](#)

The following table displays the profile resource requirements when **HTTP2 FEED** is "65K" and **WEIGHTED_LB** is "ON", and/or **Filter** is "ON":

Note

For more information on weighted load balancing, see *Oracle Communications Network Analytics Data Director User Guide* and *Oracle Communications Network Analytics Data Director Installation, Upgrade, and Fault Recovery Guide*.

Note

If **ADAPTER_CORRID_LB_WITHOUT_LF_ENABLED** is "True", the resource requirements in [Table 3-3](#) table are applicable. If **ADAPTER_CORRID_LB_WITHOUT_LF_ENABLED** is "False", see the resource requirements in the following table. Resources are updated only for the *ocnaddadapter* and *ocnaddcache* services. See the [Table 3-3](#) table for all other service configurations.

Table 3-4 Resource Requirements

OCNADD Service	9K NRF Profile				41K SCP Profile				15K SEPP Profile				65K Profile or 65K SCP Profile (SCP:41K NRF:9K SEPP:15K)			
Replication Factor = 1 Message Size = 3500 FEED Type = HTTP2	vC PU	Total Replica	Memory Required (Gi)	Topic Partitions	vC PU	Total Replica	Memory Required (Gi)	Topic Partitions	vC PU	Total Replica	Memory Required (Gi)	Topic Partitions	vC PU	Total Replica	Memory Required (Gi)	Topic Partitions

Table 3-4 (Cont.) Resource Requirements

OCNADD Service	9K NRF Profile				41K SCP Profile				15K SEPP Profile				65K Profile or 65K SCP Profile (SCP:41K NRF:9K SEPP:15K)			
ocnaddadapter	3	3	3	MAIN=27	3	7	3	MAIN=63	3	3	3	MAIN=27	3	13	4	MAIN=117 (Each instance 9 partition)
ocnaddcache	1	3	24		1	4	24		1	3	24		1	6	30	

Note

- When advanced features such as Ingress filtering, Egress filtering, and Weighted load balancing are enabled simultaneously, the resource requirement for "ocnaddadapter" service may vary at higher throughput.
- Four instances of Kafka brokers might be required when running RF=2, based on the setup performance. The end-to-end latency might increase when DISK I/O is slow.
- For DISK I/O, see [Disk Throughput Requirements](#).
- For Kafka PVC-Storage, see [Kafka PVC Storage Requirements](#).

The following table displays the profile resource requirements when **SYNTHETIC FEED** is "30K" and **WEIGHTED_LB** is "OFF":

Table 3-5 Resource Requirements

OCNADD Service	6K NRF Profile				15K SCP Profile				9K SEPP Profile				30K Profile (SCP:15K NRF:6K SEPP:9K)				30K SCP Profile			
Replication Factor = 1 Message Size = 3500 FEED Type = SYNTHETIC	vCPU	Total Replicas	Memory Required (Gi)	Topic Partitions	vCPU	Total Replicas	Memory Required (Gi)	Topic Partitions	vCPU	Total Replicas	Memory Required (Gi)	Topic Partitions	vCPU	Total Replicas	Memory Required (Gi)	Topic Partitions	vCPU	Total Replicas	Memory Required (Gi)	Topic Partitions
ocnadd configuration	1	1	1		1	1	1		1	1	1		1	1	1		1	1	1	
ocnadd alarm	1	1	1		1	1	1		1	1	1		1	1	1		1	1	1	
ocnadd admin	1	1			1	1	1		1	1	1		1	1	1		1	1	1	
ocnadd healthmonitoring	1	1	1			1	1		1	1	1		1	1	1		1	1	1	
ocnadd scpaggregation					2	1	2	SCP=6					2	1	2	SCP=6 (Each instance 6 partitions)	2	1	2	SCP=6 (Each instance 6 partitions)
ocnadd nrfaggregation	2	1	2	NRF=6									2	1	2	NRF=6 (Each instance 6 partitions)				

Table 3-5 (Cont.) Resource Requirements

OCNADD Service	6K NRF Profile				15K SCP Profile				9K SEPP Profile				30K Profile (SCP:15K NRF:6K SEPP:9K)				30K SCP Profile			
ocnadd seppaggregation									2	1	2	SEPP=6	2	1	2	SEPP=6 (Each instance 6 partitions)				
ocnadd adapter	3	1	3	MAIN=6	3	1	3	MAIN=6	3	1	3	MAIN=6	3	2	4	MAIN=12 (Each instance 6 partitions)	3	2	4	MAIN=12 (Each instance 6 partitions)
ocnadd kafka	2	3	12		3	3	24		3	3	12		5	3	48		5	3	48	
zookeeper	1	3	1		1	3	1		1	3	1		1	3	2		1	3	2	
ocnadd gui	2	1	1		2	1	1		2	1	1		2	1	1		2	1	1	
ocnadd uirouter	2	1	1		2	1	1		2	1	1		2	1	1		2	1	1	
ocnadd cache	1	2	10		1	2	10		1	2	10		1	2	10		1	2	10	

Note

- Four instances of Kafka brokers might be required when running RF=2, based on the setup performance. The end-to-end latency might increase when DISK I/O is slow.
- For DISK I/O, see [Disk Throughput Requirements](#).
- For Kafka PVC-Storage, see [Kafka PVC Storage Requirements](#).

The following table displays the profile resource requirements when **SYNTHETIC FEED** is "30K" and **WEIGHTED_LB** is "ON", and/or **Filter** is "ON":

Note

For more information on weighted load balancing see, *Oracle Communications Network Analytics Data Director User Guide* and *Oracle Communications Network Analytics Data Director Installation, Upgrade, and Fault Recovery Guide*.

Note

If **ADAPTER_CORRID_LB_WITHOUT_LF_ENABLED** is "True", the resource requirements in the [Table 3-5](#) table are applicable. If **ADAPTER_CORRID_LB_WITHOUT_LF_ENABLED** is "False", see the resource requirements in the following table. Resources are updated only for the *ocnaddadapter* and *ocnaddcache* services, see the [Table 3-5](#) table for all other service configurations.

Table 3-6 Resource Requirements

OCNADD Service	6K NRF Profile				15K SCP Profile				9K SEPP Profile				30K Profile or 30K SCP Profile (SCP:15K NRF:9K SEPP:6K)			
Replication Factor = 1 Message Size = 3500 FEED Type = HTTP2, SYNTHTIC	vC PU	Total Replica	Memory Required (Gi)	Topic Partitions	vC PU	Total Replica	Memory Required (Gi)	Topic Partitions	vC PU	Total Replica	Memory Required (Gi)	Topic Partitions	vC PU	Total Replica	Memory Required (Gi)	Topic Partitions

Table 3-6 (Cont.) Resource Requirements

OCNADD Service	6K NRF Profile				15K SCP Profile				9K SEPP Profile				30K Profile or 30K SCP Profile (SCP:15K NRF:9K SEPP:6K)			
ocnaddadapter	3	1	3	MAIN=6	3	2	3	MAIN=12	3	1	3	MAIN=6	3	4	4	MAIN=24 (Each instance 6 partitions)
ocnaddcache	1	3	24		1	4	24		1	3	24		1	6	24	

Note

- When advanced features such as Ingress filtering, Egress filtering, and Weighted load balancing are enabled simultaneously, the resource requirement for "ocnaddadapter" service may vary at higher throughput.
- Four instances of Kafka brokers might be required when running RF=2, based on the setup performance. The end-to-end latency might increase 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 Deployment Profiles

Default Deployment Profile

This profile can stream NFs (SCP, NRF, SEPP) data up to 15K MPS and can be scaled to handle up to 65K MPS for HTTP2 feed when *weighted_lb* is "OFF".

Replication Factor should be "1" and the incoming message size on OCNADD should be less than or equal to 3500K.

Table 3-7 Default Deployment Profile

OCNADD Service Replication Factor = 1 Message Size = 3500K FEED Type = HTTP2, Synthetic	vCPU Req	vCPU Limit	Memory Req (Gi)	Memory Limit (Gi)	Min Replica	Max Replica	Partitio ns	Topic Name
ocnaddconfiguration	1	1	1	1	1	1		
ocnaddalarm	1	1	1	1	1	1		
ocnaddadmin	1	1	1	1	1	1		
ocnaddhealthmonit oring	1	1	1	1	1	1		
ocnaddscppaggrega tion (41K)	1	2	1	2	1	2	12	SCP
ocnaddnrfaggregati on (9K)	1	2	1	2	1	1	6	NRF
ocnaddseppaggreg ation (15k)	1	2	1	2	1	1	6	SEPP
ocnaddadapter	2	3	3	4	HTTP:2 SYNTH ETIC:1	8	72	MAIN
ocnaddkafka	2	5	4	24	3	3		
zookeeper	1	1	1	2	3	3		
ocnaddgui	1	2	1	1	1	2		
ocnadduirouter	1	2	1	1	1	2		
ocnadd-cache	1	1	20	24	2	3		

Note

- Four instances of Kafka brokers might be required when running RF=2, based on the setup performance. The end-to-end latency might increase when DISK I/O is slow.
- For DISK I/O, see [Disk Throughput Requirements](#).
- For Kafka PVC-Storage, see [Kafka PVC Storage Requirements](#)

Deployment Profile Kafka and Aggregation

Table 3-8 Deployment Profile Kafka and Aggregation

OCNADD Service Replication Factor = 1 Message Size = 3500K FEED Type = HTTP2	vCPU Req	vCPU Limit	Memory Req (Gi)	Memory Limit (Gi)	Min Replica	Max Replica	Partitio ns	Topic Name
ocnaddscppaggregation (41K)	1	2	1	2	1	2	12	SCP
ocnaddnrfaggregation (9K)	1	2	1	2	1	1	6	NRF
ocnaddseppaggregation (15k)	1	2	1	2	1	1	6	SEPP
ocnaddkafka	2	5	4	24	3	3		
zookeeper	1	1	1	2	3	3		

Note

- Four instances of Kafka brokers might be required when running RF=2, based on the setup performance. The end-to-end latency might increase when DISK I/O is slow.
- For DISK I/O, see [Disk Throughput Requirements](#).
- For Kafka PVC-Storage, see [Kafka PVC Storage Requirements](#)

Deployment Profile Kafka and Egress Adapter**Table 3-9 Deployment Profile Kafka and Egress Adapter**

OCNADD Service Replication Factor = 1 Message Size = 3500K FEED Type = HTTP2	vCPU Req	vCPU Limit	Memory Req (Gi)	Memory Limit (Gi)	Min Replica	Max Replica	Partitio ns	Topic Name
ocnaddconfiguration	1	1	1	1	1	1		
ocnaddalarm	1	1	1	1	1	1		
ocnaddadmin	1	1	1	1	1	1		
ocnaddhealthmonitoring	1	1	1	1	1	1		
ocnaddadapter	2	3	3	4	HTTP:2 SYNTH ETIC:1	8	72	MAIN
ocnaddkafka	2	5	4	24	3	3		
zookeeper	1	1	1	2	3	3		
ocnaddgui	1	2	1	1	1	2		

Table 3-9 (Cont.) Deployment Profile Kafka and Egress Adapter

OCNADD Service Replication Factor = 1 Message Size = 3500K FEED Type = HTTP2	vCPU Req	vCPU Limit	Memory Req (Gi)	Memory Limit (Gi)	Min Replica	Max Replica	Partitio ns	Topic Name
ocnadduirouter	2	2	1	1	1	2		
ocnaddcache	1	1	20	24	2	3		

Note

- Four instances of Kafka brokers might be required when running RF=2, based on the setup performance. The end-to-end latency might increase when DISK I/O is slow.
- For DISK I/O, see [Disk Throughput Requirements](#).
- For Kafka PVC-Storage, see [Kafka PVC Storage Requirements](#)

3.3 Ephemeral Storage Requirements

The following table describes the Ephemeral Storage requirements for OCNADD:

Table 3-10 Ephemeral Storage Requirements

Service Name	Ephemeral Storage (min) in Mi	Ephemeral Storage (max) in Mi
<app-name>-adapter	200	800
<app-name>-gw	400	800
ocnaddadminservice	100	200
ocnaddalarm	100	500
ocnaddhealthmonitoring	100	500
ocnaddscpaggregation	100	500

Note

Supported only in release 22.0.0

Table 3-10 (Cont.) Ephemeral Storage Requirements

Service Name	Ephemeral Storage (min) in Mi	Ephemeral Storage (max) in Mi
ocnaddseppaggregation	100	500
ocnaddnrfaggregation	100	500
ocnaddconfiguration	100	500

Note

Supported from release 23.1.0 onwards

3.4 Disk Throughput Requirements

The following table describes the disk throughput requirements in OCNADD:

Table 3-11 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.

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 = $(a_1b_1 + a_2b_2 + \dots + a(n)b(n)) / (a_1 + a_2 + \dots + a(n))$

a_1 -> SCP MPS
 b_1 -> SCP message size
 a_2 -> NRF MPS
 b_2 -> 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 and NRF:

Table 3-12 SCP and NRF Disk Throughput

SCP Message		NRF 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	Rate
Avg Size (in Bytes)	Rate	Avg Size (in Bytes)											
1624	30000	3000	9000	1	2	1	145	145	3	54	54	108	324
1624	30000	3000	9000	2	2	1	289	289	3	106	106	212	636
4000	30000	3000	9000	1	2	1	281	281	3	104	104	208	624

Table 3-12 (Cont.) SCP and NRF Disk Throughput

SCP Message		NRF 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	Rate
Avg Size (in Bytes)	Rate	Avg Size (in Bytes)											
4000	30000	3000	9000	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

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

Table 3-13 Retention Period Per Topic

Topic Name	Retention Period
SCP	5 Minutes
NRF	5 Minutes
SEPP	5 Minutes
MAIN	6 Hours (Max)

The following calculation is for storage requirements for a topic:

Important

For the 6 hrs storage in the MAIN topic, the storage requirement must be calculated using the following information:

Storage Requirement for a topic = MPS * Retention Period * RF * Average Message Size

Where,

MPS is "Message Per Second"

RF is "Replication Factor"

Examples:

1. Average Message Size = 1941 Bytes

The following example uses the values from the first row of the [Table 3-12](#) table. For more information about the table, see [Disk Throughput Requirements](#):

Storage Requirement for SCP and NRF Topics = MPS * Retention Period * RF * Message Size = 39000 * 5 Minutes * 3 * 1941

$$= 39000 * 5 * 60 * 3 * 1941$$

$$= \sim 63.45 \text{ GB}$$

Storage Requirement for MAIN = MPS * Retention Period * RF * Message Size = 39000 * 6 Hours * 3 * 1941

$$= 39000 * 6 * 60 * 60 * 3 * 1941$$

$$= \sim 4.46 \text{ TB}$$

Total Storage Requirement for the Broker Cluster = Storage for SCP + Storage for NRF + Storage for MAIN

$$= 63.45 \text{ GB} + 4.46 \text{ TB}$$

$$= \sim 4.53 \text{ TB}$$

Total Storage for each broker = (4.53/Number of Brokers) TB = (4.53/3) TB = $\sim 1.51 \text{ TB}$ [Assuming 3 Broker cluster]

2. Average Message Size = 3769 Bytes

The following example uses the values from the fourth row of the [Table 3-12](#) table. For more information about the table, see [Disk Throughput Requirements](#):

Storage Requirement for SCP and NRF Topics = MPS * Retention Period * RF * Message Size = 39000 * 5 Minutes * 3 * 3769

$$= 39000 * 5 * 60 * 3 * 3769$$

$$= \sim 123.20 \text{ GB}$$

Storage Requirement for MAIN = MPS * Retention Period * RF * Message Size = 39000 * 6 Hours * 3 * 3769

$$= 39000 * 6 * 60 * 60 * 3 * 3769$$

$$= \sim 8.66 \text{ TB}$$

Total Storage Requirement for the Broker Cluster = Storage for SCP + Storage for NRF + Storage for MAIN

$$= 123.20 \text{ GB} + 8.66 \text{ TB}$$

$$= \sim 8.79 \text{ TB}$$

Total Storage for each broker = (8.79/Number of Brokers) TB = (8.79/3) TB
= ~ 2.93 TB [Assuming 3 Broker cluster]

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. The message feed feature is tested with SCP, NRF, and SEPP as the source of the message feed.

- [Performance Benchmarking with HTTP2 Feed](#)
- [Performance Benchmarking with Synthetic Feed](#)

4.1 Performance Benchmarking with HTTP2 Feed

The performance benchmarking test is performed on OCNADD with 64K MPS HTTP2 egress feed which includes 40K MPS SCP traffic (20K TPS with 2 traffic copy trigger points on SCP), 9K MPS NRF traffic (4.5K TPS with 2 traffic copy trigger points on NRF) and 15K MPS SEPP traffic (7.5K TPS with 2 traffic copy trigger points on SEPP). The benchmarking results are provided below:

Note

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

Benchmark Test Environment

The test environment specifications are listed below:

- Ingress Traffic Rate: 64K MPS
- The latency, resource utilization, and MPS is observed for the period of 3 hrs or more.
- OCNADD single site deployment with ASM disabled
- OCNADD Release 23.2.0 deployed over CNE (Bare metal) 22.3.0 version
- Kafka Replication Factor (RF): 1
- OCNADD Data Feed: HTTP2/TLS
- OCNADD Kafka PVC: 300GB
For more information about PVC requirements, see [Kafka PVC Storage Requirements](#).
- 3rd Party Application: 1 Perfgo Server with 2 End-points
- cnDBTier Release 23.2.0
- NRF Release 23.1.0
- SCP Release 23.2.0
- SEPP Release 23.1.0
- SASL/SSL, enabled between NRF/SCP/SEPP and OCNADD
- CNCC Release 23.2.0
- Message Size: 3490 Bytes

- Environment: CNE (Bare metal)

Benchmark Testcase Specifications

The testcase parameters are as follows:

- **OCNADD**
 - A feed is configured using the OCNADD Console with aggregation rules set for SCP, NRF and SEPP.
 - Message Ingestion Rate: 64K MPS

Resource Specifications:

Table 4-1 OCNADD Resource Specifications

Services	CPU Request Per Pod	CPU Limit Per Pod	Memory Request Per Pod (Gi)	Memory Limit Per Pod (Gi)	min Replicas	max Replicas	Kafka Topic Partitions and Retention
ocnaddadpter	3	3	4	4	8	8	72 (MAIN)/ retention. ms=300000
ocnaddadmin	1	1	1	1	1	1	
ocnaddalarm	1	1	1	1	1	1	
ocnaddcache	1	1	10	10	2	2	
ocnaddconfiguration	1	1	1	1	1	1	
ocnaddgui	2	2	1	1	1	1	
ocnaddhealthmonitoring	1	1	1	1	1	1	
ocnaddkafka	5	5	48	48	3	3	
ocnaddnrfaggregation (9K)	2	2	2	2	1	1	6 (NRF)/ retention. ms=300000
ocnaddscpaggregation (40K)	2	2	2	2	2	2	12 (SCP)/ retention. ms=300000
ocnaddseppaggregation (15K)	2	2	2	2	1	1	6 (SEPP)/ retention. ms=300000
ocnaduiroter	2	2	1	1	1	1	
zookeeper	1	1	2	2	3	3	

- **SCP**
SCP Traffic: 20K TPS using two trigger points.
- **NRF**
NRF Traffic: 4.5K TPS using two trigger points.
- **SEPP**
SEPP Traffic: 7.5K 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	64K MPS	12	2.66 ms	99.94%
NRF	2 (IGW/EGW)	9K MPS	12		99.90%
SCP	2 (Request Ingress/ Request Egress)	40K MPS	12		99.98%
SEPP	2 (PLMN IGW)	15K MPS	12		99.98%

CPU and Memory Utilization

The following table describes the OCNADD CPU and memory utilization:

Table 4-3 OCNADD CPU and Memory Utilization

Microservice/ Container	Pod Count	CPU Utilization (%)	Memory Utilization (%)
ocnaddadapter	8	71.2	25.3
ocnaddkafka	3	87.6	18.1
ocnaddadmin	1	0.195	26.5
ocnaddalarm	1	0.546	18.3
ocnadduirouter	1	0.0410	24.6
ocnaddconfiguration	1	0.130	28.8
ocnaddhealthmonitoring	1	0.222	27.1
ocnaddgui	1	0.033	4.70
ocnaddnrfaggregation (9K)	1	73.8	44.3
ocnaddscpaggregation (40K)	2	79.9	59.2
ocnaddseppaggregation (15K)	1	80.7	35.4
zookeeper	3	0.270	14.2
ocnaddcache	2	0.202	6.89

4.2 Performance Benchmarking with Synthetic Feed

The performance benchmarking test is performed on OCNADD with 30K MPS synthetic feed which includes 16K MPS SCP traffic (8K TPS with 2 trigger points), 5K MPS NRF traffic (2.5K TPS with 2 trigger points) and 9K MPS SEPP traffic (4.5K TPS with 2 trigger points). The benchmarking results are provided below:

Note

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

Benchmark Test Environment

The test environment specifications are listed below:

- Ingress Traffic Rate: 30K MPS
- The latency, resource utilization, and MPS is observed for the period of 4 hrs or more.
- OCNADD single site deployment with ASM disabled
- OCNADD Release 23.2.0 deployed over CNE (Bare metal) 22.3.0 version
- Kafka Replication Factor (RF): 1
- OCNADD Data Feed: Secure TCP feed (TCP with TLS enabled)
- OCNADD Kafka PVC: 300GB
For more information about PVC requirements, see [Kafka PVC Storage Requirements](#).
- 3rd Party Application: 1 Perfgo Server with 2 End-points
- cnDBTier Release 23.2.0
- NRF Release 23.1.0
- SCP Release 23.2.0
- SEPP Release 23.1.0
- SASL/SSL, enabled between NRF/SCP/SEPP and OCNADD
- CNCC Release 23.2.0
- Message Size: 3000 Bytes
- Environment: CNE (Bare metal)

Benchmark Testcase Specifications

The testcase parameters are as follows:

- **OCNADD**
 - A feed is configured using the OCNADD Console with aggregation rules set for SCP, NRF and SEPP.
 - Message Ingestion Rate: 30K MPS

Resource Specifications:

Table 4-4 OCNADD Resource Specifications

Services	CPU Request Per Pod	CPU Limit Per Pod	Memory Request Per Pod (Gi)	Memory Limit Per Pod (Gi)	min Replicas	max Replicas	Kafka Topic Partitions and Retention
ocnaddconfiguration	1	1	1	1	1	1	
ocnaddalarm	1	1	1	1	1	1	
ocnaddadmin	1	1	1	1	1	1	
ocnaddhealthmonitoring	1	1	1	1	1	1	
ocnaduirouter	1	1	1	1	1	1	
ocnaddscppaggregation (16K)	2	2	2	2	2	2	12 (SCP)/ retention. ms=300000
ocnaddnrffaggregation (5K)	2	2	2	2	1	1	6 (NRF)/ retention. ms=300000
ocnaddseppaggregation (9K)	2	2	2	2	1	1	6 (SEPP)/ retention. ms=300000
ocnaddadapter	3	3	4	4	2	2	12 (MAIN)/ retention. ms=300000
ocnaddkafka	5	5	48	48	3	3	
zookeeper	1	1	2	2	3	3	
ocnaddgui	2	2	1	1	1	1	
ocnaddcache	1	1	10	10	2	2	

- **SCP**
SCP Traffic: 8K TPS using two trigger points.
- **NRF**
NRF Traffic: 2.5K TPS using two trigger points.
- **SEPP**
SEPP Traffic: 4.5K TPS using two trigger points.

Benchmark Test Results

Traffic Feed Details

Table 4-5 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	30K MPS	4	3.15 ms	99.99%
NRF	2 (IGW/ EGW)	5K MPS	4		99.96%
SCP	2 (Request Ingress/ Request Egress)	16K MPS	4		99.8%
SEPP	2 (PLMN IGW)	9K MPS	4		99.99%

CPU and Memory Utilization

The following table describes the OCNADD CPU and memory utilization:

Table 4-6 OCNADD CPU and Memory Utilization

Microservice/ Container	Pod Count	CPU Utilization (%)	Memory Utilization (%)
ocnaddadapter	2	80.6	27.5
ocnaddkafka	3	53.7	10.9
ocnaddadmin	1	0.247	27.2
ocnaddalarm	1	0.306	18.9
ocnadduirouter	1	0.0190	24.6
ocnaddconfiguration	1	0.154	28.0
ocnaddhealthmonitoring	1	0.242	27.1
ocnaddgui	1	0.0301	5.15
ocnaddnrfaggregation (5K)	1	49.3	22.4
ocnaddscpaggregation (16K)	2	35.8	30.9
ocnaddseppaggregation (9K)	1	64.7	28.0
zookeeper	3	0.248	15.0
ocnaddcache	2	0.0656	8.09