Oracle® Communications

Network Analytics Data Director Benchmarking Guide

Release 23.4.0.0.1
F91516-02
July 2024

Oracle Communications Network Analytics Data Director Benchmarking Guide, Release 23.4.0.0.1

F91516-02

Contents

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
XDR	Extended Detection and Response

What's New in This Guide

This section lists the documentation updates for Release 23.4.x in *Oracle Communications Network Analytics Data Director Benchmarking Guide*.

Release 23.4.0.0.1 - F91516-02, July 2024

There are no updates to this document in this release.

Release 23.4.0 - F91516-01, January 2024

Updated the following sections:

- Updated the following sections in the Profile Resource Requirements chapter:
  - Profile Resource Requirements for HTTP2 Feed
  - Profile Resource Requirements for Synthetic Feed
  - Profile Resource Requirements for Aggregated Kafka Feed
  - Profile Resource Requirements for Correlated Kafka Feed
- Updated the Deployment Profiles section.
- Added the Pod Affinity (or Anti-affinity) Rules section.
- Added the following sections in the OCNADD Benchmarking Testing chapter:
  - Performance Benchmarking with Correlated Feed for XDR Generation
  - Performance Benchmarking with Synthetic Feed with Replication
  - Performance Benchmarking with 60K MPS SEPP Traffic with HTTP2 Feed Replication

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 deployment. 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

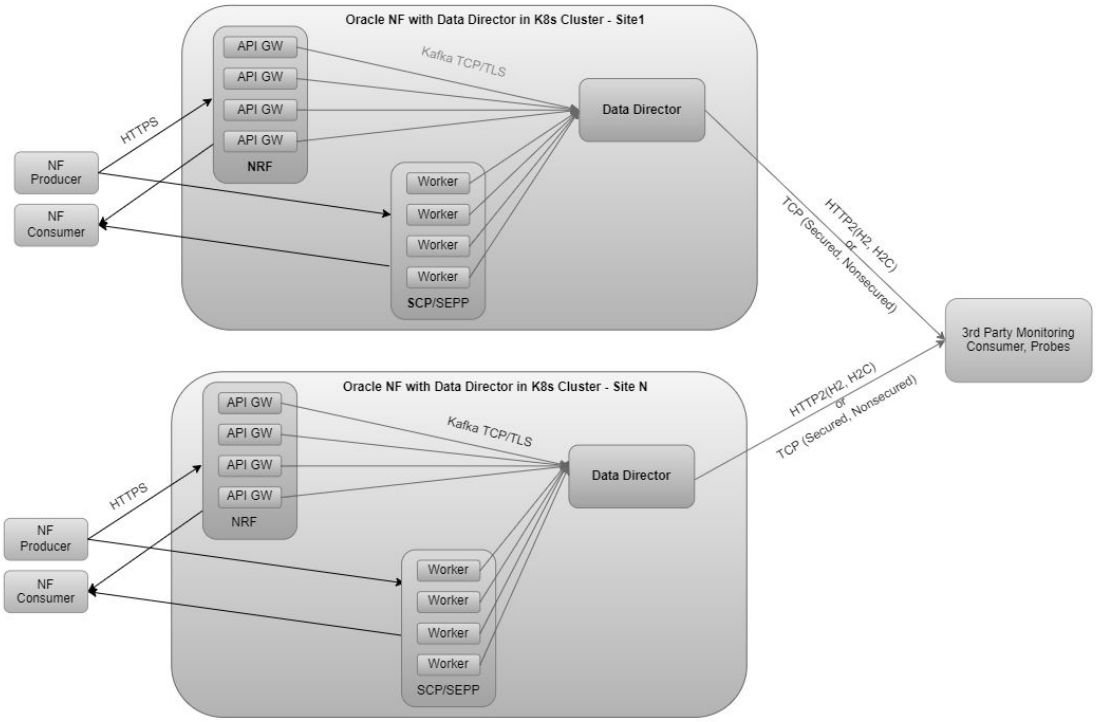

> **Note**
>
> From release 23.1.0 onwards OCNADD supports Synthetic Feeds along with HTTP2 Feeds.

OCNADD uses the following common services of CNE:

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

3
Resource Requirement

This 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.

The resource requirement for the profiles (up to 100K MPS for HTTP2 feed and up to 64K MPS for synthetic feed) 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	7
Storage Class	Standard
Top of Rack Switch	1

Table 3-2 CNE Configuration for Kafka Traffic

Type of Server	Baremetal CNE
Master node	3
Worker node	7
Storage Class	Standard
Top of Rack Switch	1

Table 3-3 CNE Configuration for Synthetic Traffic

Type of Server	Baremetal CNE
Master node	3
Worker node	7
Storage Class	Standard
Top of Rack Switch	1

3.1 Deployment Profiles

Default Deployment Profile

This profile can stream NFs (SCP, NRF, SEPP) data up to 100K MPS and can be scaled to handle up to 100K MPS for HTTP2 feed (and 64K for synthetic feed) when *weighted_lb* and Filter (Ingress and Egress) are "OFF".

Replication Factor should be "1" or "2" (if SSD is supported) and the incoming message size on OCNADD should be less than or equal to 3500K.

Table 3-4 Default Deployment Profile

OCNADD Service Replication Factor = 1 or 2 (when SSD is enabled) Message Size = 3500K FEED Type = HTTP2, Synthetic	vCPU Req	vCPU Limit	Memory Req (Gi)	Memory Limit (Gi)	Min Replica	Max Replica	Partitions	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		
ocnaddscpaggregation (55K)	2	2	2	2	1	2	12	SCP
ocnaddnrfaggregation (15K)	2	2	2	2	1	1	6	NRF
ocnaddseppaggregation (30K)	2	2	2	2	1	2	12	SEPP
ocnaddadapter	3	3	4	4	HTTP:2 SYNTHETIC:1	HTTP:13 SYNTHETIC:5	117	MAIN
ocnaddkafka	6	6	48	48	4	4		
zookeeper	1	1	1	2	3	3		
ocnaddgui	1	2	1	1	1	2		
ocnadduirouter	1	2	1	1	1	2		

> **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.

3.2 Pod Affinity (or Anti-affinity) Rules

The ocnaddkafka and Zookeeper services use the POD anti-affinity rules. The rules are intended to support even distribution of Kafka and Zookeeper pods across the available nodes.

Zookeeper Service Anti-affinity Rules

Zookeeper service anti-affinity rules are listed below:

```
spec:
      affinity:
```

```
      podAntiAffinity:
        preferredDuringSchedulingIgnoredDuringExecution:
        - weight: 100
          podAffinityTerm:
            labelSelector:
              matchExpressions:
              - key: app
                operator: In
                values:
                - zookeeper
            topologyKey: topology.kubernetes.io/zone
```

***ocnaddkafka* Service Anti-affinity Rules**

ocnaddkafka service anti-affinity rules are listed below:

```
spec:
      topologySpreadConstraints:
        - maxSkew: 1
          topologyKey: kubernetes.io/hostname
          whenUnsatisfiable: ScheduleAnyway
          labelSelector:
            matchLabels:
              app1 : nodeselection
```

The configuration ensures the scheduler keeps equal number of pods matching the constraint on every node. However, the parameter `whenUnsatisfiable` set to `ScheduleAnyway`, helps to schedule the pod, but the scheduler prioritizes honoring the skew to not make the cluster more imbalanced.

3.3 Profile Resource Requirements

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

> **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*.

- Profile Resource Requirements for HTTP2 Feed
- Profile Resource Requirements for Synthetic Feed
- Profile Resource Requirements for Aggregated Kafka Feed
- Profile Resource Requirements for Correlated Kafka Feed

3.3.1 Profile Resource Requirements for HTTP2 Feed

The following table displays the profile resource requirements when HTTP2 FEED is "100K" MPS.

Table 3-5 Resource Requirements

OCNADD Service	15K NRF Profile				55K SCP Profile				30K SEPP Profile				100K Profile (SCP:55K NRF:15K SEPP:30K)				100K SCP Profile			
Replication Factor = 1 or 2 (SSD is enabled) **Message Size = 3500** **FEED Type = HTTP2**	**vCPU**	**Total Replica**	**Memory Required** (Gi)	**Topic Partitions**	**vCPU**	**Total Replica**	**Memory Required** (Gi)	**Topic Partitions**	**vCPU**	**Total Replica**	**Memory Required** (Gi)	**Topic Partitions**	**vCPU**	**Total Replica**	**Memory Required** (Gi)	**Topic Partitions**	**vCPU**	**Total Replica**	**Memory Required** (Gi)	**Topic Partitions**
OCNADD microservices resource requirements:																				
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 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	-
zookeeper	1	3	1	-	1	3	1	-	1	3	1	-	1	3	2	-	1	3	2	-
ocnadd kafka	2	3	24	-	4	3	48	-	3	3	48	-	6	4	48	-	6	4	48	-
Replicated Feed																				
ocnadd kafka	2	3	24	-	4	3	48	-	4	3	48	-	10	4	96	-	10	4	96	-
WLB is ON/OFF and Filter is OFF																				

Table 3-5 (Cont.) Resource Requirements

OCNADD Service	15K NRF Profile				55K SCP Profile				30K SEPP Profile				100K Profile (SCP:55K NRF:15K SEPP:30K)				100K SCP Profile			
ocnaddscpaggregation	-	-	-	-	2	2	2	SCP=12	-	-	-	-	2	3	2	SCP=18 (Each instance 6 partitions)	2	4	2	SCP=24 (Each instance 6 partitions)
ocnaddnrfaggregation	2	1	2	NRF=6	-	-	-	-	-	-	-	-	2	1	2	NRF=6 (Each instance 6 partitions)	-	-	-	-
ocnaddseppaggregation	-	-	-	-	-	-	-	-	2	2	2	SEPP=12	2	2	2	SEPP=12 (Each instance 6 partitions)	-	-	-	-

Table 3-5 (Cont.) Resource Requirements

OCNADD Service	15K NRF Profile				55K SCP Profile				30K SEPP Profile				100K Profile (SCP:55K NRF:15K SEPP:30K)				100K SCP Profile			
ocnaddadapter	3	3	3	MAIN=27	3	8	3	MAIN=72	3	5	3	MAIN=45	3	14	4	MAIN=126 (Each instance 9 partitions)	3	14	4	MAIN=126 (Each instance 9 partitions)
WLB is ON/OFF and Filter is ON																				
Ingress Filter is ON:																				
ocnaddscpaggregation	-	-	-	-	2	3	2	SCP=18	-	-	-	-	2	4	2	SCP=24 (Each instance 6 partitions)	2	5	2	SCP=30 (Each instance 6 partitions)
ocnaddnrfaggregation	2	2	2	NRF=12	-	-	-	-	-	-	-	-	2	2	2	NRF=12 (Each instance 6 partitions)	-	-	-	-

Table 3-5 (Cont.) Resource Requirements

OCNADD Service	15K NRF Profile				55K SCP Profile				30K SEPP Profile				100K Profile (SCP:55K NRF:15K SEPP:30K)				100K SCP Profile			
ocnaddseppaggregation	-	-	-	-	-	-	-	-	2	3	2	SEPP=18	2	3	2	SEPP=18 (Each instance 6 partitions)	-	-	-	-
Egress Filter is ON:																				
ocnaddadapter	3	4	3	MAIN=36	3	8	3	MAIN=72	3	5	3	MAIN=45	3	15	4	MAIN=135 (Each instance 9 partitions)	3	15	4	MAIN=135 (Each instance 9 partitions)

> **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.
> - Filter is not tested with replicated feed enabled due to disk I/O limitation and coherence performance issue.
> - Resource requirement may vary when Filter is ON based on % data allowed after filtering and number of filter conditions with configured values.
> - 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.

3.3.2 Profile Resource Requirements for Synthetic Feed

The following table displays the profile resource requirements when **Synthetic FEED** is "100K" MPS.

> **Note**
>
> Set the value of the parameter *TCP_CONN_INSPECTOR_FIXED_DELAY_MS* to "60000" in the *custom-value.yaml* file.

Table 3-6 Resource Requirements

OCNADD Service	**15K NRF Profile**				**55K SCP Profile**				**30K SEPP Profile**				**100K Profile (SCP:55K NRF:15K SEPP:30K)**				**100K SCP Profile**			
Replication Factor = 1 or 2 (when SSD is enabled) **Message Size = 3500** **FEED Type = SYNTHETIC**	**vCPU**	**Total Replica**	**Memory Required** (Gi)	**Topic Partitions**	**vCPU**	**Total Replica**	**Memory Required** (Gi)	**Topic Partitions**	**vCPU**	**Total Replica**	**Memory Required** (Gi)	**Topic Partitions**	**vCPU**	**Total Replica**	**Memory Required** (Gi)	**Topic Partitions**	**vCPU**	**Total Replica**	**Memory Required** (Gi)	**Topic Partitions**
OCNADD microservices resource requirements:																				
ocnadd configuration	1	1	1	-	1	1	1	-	1	1	1	-	1	1	1	-	1	1	1	-

Table 3-6 (Cont.) Resource Requirements

OCNADD Service	15K NRF Profile				55K SCP Profile				30K SEPP Profile				100K Profile (SCP:55K NRF:15K SEPP:30K)				100K SCP Profile			
ocnaddalarm	1	1	1	-	1	1	1	-	1	1	1	-	1	1	1	-	1	1	1	-
ocnaddadmin	1	1	-	-	1	1	1	-	1	1	1	-	1	1	1	-	1	1	1	-
ocnaddhealthmonitoring	1	1	1	-		1	1	-	1	1	1	-	1	1	1	-	1	1	1	-
ocnaddgui	2	1	1	-	2	1	1	-	2	1	1	-	2	1	1	-	2	1	1	-
ocnadduirouter	2	1	1	-	2	1	1	-	2	1	1	-	2	1	1	-	2	1	1	-
zookeeper	1	3	1	-	1	3	1	-	1	3	1	-	1	3	2	-	1	3	1	-
ocnaddkafka	2	3	24	-	3	3	24	-	3	3	24	-	6	4	48	-	6	4	48	-
Replicated Feed																				
ocnaddkafka	2	3	24	-	4	3	48	-	4	3	24	-	10	4	96	-	10	4	96	-
WEIGHTED-LB is ON/OFF, L3L4 Mapping is ON/OFF and FILTER is OFF																				
ocnaddscpaggregation	-	-	-	-	2	2	2	SCP=12	-	-	-	-	2	3	2	SCP=18 (Each instance 6 partitions)	2	4	2	SCP=24 (Each instance 6 partitions)
ocnaddnrfaggregation	2	1	2	NRF=6	-	-	-	-	-	-	-	-	2	1	2	NRF=6 (Each instance 6 partitions)	-	-	-	-

Table 3-6 (Cont.) Resource Requirements

OCNADD Service	15K NRF Profile				55K SCP Profile				30K SEPP Profile				100K Profile (SCP:55K NRF:15K SEPP:30K)				100K SCP Profile			
ocnaddseppaggregation	-	-	-	-	-	-	-	-	2	2	2	SEPP=12	2	2	2	SEPP=12 (Each instance 6 partitions)	-	-	-	-
ocnaddadapter	3	1	3	MAIN=6	3	3	3	MAIN=18	3	2	3	MAIN=12	3	9	4	MAIN=54 (Each instance 6 partitions)	3	9	4	MAIN=54 (Each instance 6 partitions)
WEIGHTED-LB is ON/OFF, L3L4 Mapping is ON/OFF and FILTER is ON																				
Ingress Filter is ON																				
ocnaddscpaggregation	-	-	-	-	2	4	2	SCP=18	-	-	-	-	2	4	2	SCP=24 (Each instance 6 partitions)	2	5	2	SCP=30 (Each instance 6 partitions)

Table 3-6 (Cont.) Resource Requirements

OCNADD Service	15K NRF Profile				55K SCP Profile				30K SEPP Profile				100K Profile (SCP:55K NRF:15K SEPP:30K)				100K SCP Profile			
ocnaddnrfaggregation	2	2	2	NRF=6	-	-	-	-	-	-	-	-	2	2	2	NRF=12 (Each instance 6 partitions)	-	-	-	-
ocnaddseppaggregation	-	-	-	-	-	-	-	-	2	3	2	SEPP=12	2	3	2	SEPP=18 (Each instance 6 partitions)	-	-	-	-
Egress Filter is ON:																				
ocnaddadapter	3	2	3	MAIN=6	3	4	3	MAIN=24	3	3	3	MAIN=18	3	9	4	MAIN=54 (Each instance 6 partitions)	3	9	4	MAIN=54 (Each instance 6 partitions)

> **Note**
>
> - When advanced features such as Ingress filtering, Egress filtering, L3-L4 and Weighted load balancing are enabled simultaneously, the resource requirement for "*ocnaddadapter*" service may vary at higher throughput.
> - Resource requirement may vary when Filter is ON based on % data allowed after filtering and number of filter condition with values configured.
> - Resource requirement may vary when L3L4 is ON and size of global l3l4 configuration is huge.
> - Five or six 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.

3.3.3 Profile Resource Requirements for Aggregated Kafka Feed

The following table displays the profile resource requirements when Aggregated Kafka Feed is "100K" MPS:

Table 3-7 Resource Requirements

OCNADD Service	15K NRF Profile				55K SCP Profile				30K SEPP Profile				100K Profile (SCP:55K NRF:15K SEPP:30K)				100K SCP Profile			
Replication Factor = 1 or 2 (when SSD is enabled) Message Size = 3500 FEED Type = HTTP2	**vCPU**	**Total Replica**	**Memory Required** (Gi)	**Topic Partitions**	**vCPU**	**Total Replica**	**Memory Required** (Gi)	**Topic Partitions**	**vCPU**	**Total Replica**	**Memory Required** (Gi)	**Topic Partitions**	**vCPU**	**Total Replica**	**Memory Required** (Gi)	**Topic Partitions**	**vCPU**	**Total Replica**	**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	-

Table 3-7 (Cont.) Resource Requirements

OCNADD Service	15K NRF Profile				55K SCP Profile				30K SEPP Profile				100K Profile (SCP:55K NRF:15K SEPP:30K)				100K SCP Profile			
ocnaddgui	2	1	1	-	2	1	1	-	2	1	1	-	2	1	1	-	2	1	1	-
ocnadduirouter	2	1	1	-	2	1	1	-	2	1	1	-	2	1	1	-	2	1	1	-
zookeeper	1	3	1	-	1	3	1	-	1	3	1	-	1	3	2	-	1	3	2	-
ocnaddkafka	2	3	24	-	4	3	48	-	4	3	48	-	10	4	96	-	10	4	96	-
Replicated Feed																				
ocnaddkafka	2	3	24	-	4	3	48	-	4	3	48	-	10	4	96	-	10	4	96	-
Ingress filter is "OFF"																				
ocnaddscpaggregation	-	-	-	-	2	2	2	SCP=12	-	-	-	-	2	2	2	SCP=12 (Each instance 6 partitions)	2	4	2	SCP=24 (Each instance 6 partitions)
ocnaddnrfaggregation	2	1	2	NRF=6	-	-	-	-	-	-	-	-	2	1	2	NRF=6 (Each instance 6 partitions)	-	-	-	-

Table 3-7 (Cont.) Resource Requirements

OCNADD Service	15K NRF Profile				55K SCP Profile				30K SEPP Profile				100K Profile (SCP:55K NRF:15K SEPP:30K)				100K SCP Profile			
ocnaddseppaggregation	-	-	-	-	-	-	-	-	2	2	2	SEPP=12	2	2	2	SEPP=12 (Each instance 6 partitions)	-	-	-	-
Ingress filter is "ON"																				
ocnaddscpaggregation	-	-	-	-	2	4	2	SCP=24	-	-	-	-	2	4	2	SCP=24 (Each instance 6 partitions)	2	5	2	SCP=30 (Each instance 6 partitions)
ocnaddnrfaggregation	2	2	2	NRF=12	-	-	-	-	-	-	-	-	2	2	2	NRF=12 (Each instance 6 partitions)	-	-	-	-

Table 3-7 (Cont.) Resource Requirements

OCNADD Service	15K NRF Profile				55K SCP Profile				30K SEPP Profile				100K Profile (SCP:55K NRF:15K SEPP:30K)				100K SCP Profile			
ocnaddseppaggregation	-	-	-	-	-	-	-	-	2	3	2	SEPP=18	2	3	2	SEPP=18 (Each instance 6 partitions)	-	-	-	-

3.3.4 Profile Resource Requirements for Correlated Kafka Feed

The following table displays the profile resource requirements when Correlated Kafka Feed is "30K" MPS:

> **Note**
>
> - Creating a CORRELATED or CORRELATED_FILTERED ACL feed is mandatory before making a correlation configuration.
> - Calculate the PVC size of a Kafka broker in advance, as each new CORRELATED ACL feed type correlation configuration creates one new topic, and each new CORRELATED_FILTERED ACL feed type correlation configuration creates two new topics.
> - The CPU and memory requirement in Kafka increase based on the number of CORRELATED or CORRELATED_FILTERED configurations; the resources mentioned in the below table are for a maximum of two configurations.
> - Resource requirements for correlation service vary when all the transaction messages are not received and the value of *maxTransactionWaitTime* is set to higher value (it is recommended to have a lower value for this scenario).

Table 3-8 Resource Requirements

OCNADD Service	15K NRF Profile				30K SCP Profile				30K SEPP Profile				30K Profile (SCP:15K NRF:5K SEPP:10K)			
Replication Factor = 1 or 2 (when SSD is enabled) Message Size = 3500 ACL Feed Type: Correlated	**vCPU**	**Total Replica**	**Memory Required** (Gi)	**Topic Partitions**	**vCPU**	**Total Replica**	**Memory Required** (Gi)	**Topic Partitions**	**vCPU**	**Total Replica**	**Memory Required** (Gi)	**Topic Partitions**	**vCPU**	**Total Replica**	**Memory Required** (Gi)	**Topic Partitions**
ocnaddconfiguration	1	1	1	-	1	1	1	-	1	1	1	-	1	1	1	-
ocnaddalarm	1	1	1	-	1	1	1	-	1	1	1	-	1	1	1	-
ocnaddadmin	1	1	1	-	1	1	1	-	1	1	1	-	1	1	1	-
ocnaddhealthmonitoring	1	1	1	-	1	1	1	-	1	1	1	-	1	1	1	-
ocnaddgui	2	1	1	-	2	1	1	-	2	1	1	-	2	1	1	-
ocnadduirouter	2	1	1	-	2	1	1	-	2	1	1	-	2	1	1	-
zookeeper	1	3	1	-	1	3	1	-	1	3	1	-	1	3	2	-
ocnaddkafka	4	4	48	-	6	4	48	-	6	4	48	-	6	4	48	-
Replicated Feed																
ocnaddkafka	4	4	48	-	8	4	72	-	8	4	72	-	8	4	72	-
ocnaddscpaggregation	-	-	-	-	2	2	2	SCP=12	-	-	-	-	2	1	2	SCP=6 Each Instance 6 partition)

Table 3-8 (Cont.) Resource Requirements

OCNADD Service	15K NRF Profile				30K SCP Profile				30K SEPP Profile				30K Profile (SCP:15K NRF:5K SEPP:10K)			
ocnaddnrfaggregation	2	1	2	NRF=6	-	-	-	-	-	-	-	-	2	1	2	NRF=6 Each Instance 6 partition)
ocnaddseppaggregation	-	-	-	-	-	-	-	-	2	2	2	SEPP=12	2	1	2	SEPP=6 (Each Instance 6 partition)
Feed Type = CORRELATED																
ocnaddcorrelation	3	3	64	MAIN=18 <feed-name>-CORRELATED=18 or as per consumer	3	4	64	MAIN=24 <feed-name>-CORRELATED=24 or as per consumer	3	4	64	MAIN=24 <feed-name>-CORRELATED=24 or as per consumer	3	4	64	MAIN=24 <feed-name>-CORRELATED=24 or as per consumer
FEED TYPE = CORRELATED_FILTERED (Minimum 20% data filter is configured)																

Table 3-8 (Cont.) Resource Requirements

OCNADD Service	15K NRF Profile				30K SCP Profile				30K SEPP Profile				30K Profile (SCP:15K NRF:5K SEPP:10K)			
ocnaddfilter	2	3	3	MAIN=18 <feed-name>-FILTERED=18	2	4	3	MAIN=24 <feed-name>-FILTERED=24	2	4	3	MAIN=24 <feed-name>-FILTERED=24	2	4	3	MAIN=24 <feed-name>-FILTERED=24
ocnaddcorrelation	3	3	48	<feed-name>-FILTERED-CORRELATED=18 or as per consumer	3	4	64	<feed-name>-FILTERED-CORRELATED=24 or as per consumer	3	4	64	<feed-name>-FILTERED-CORRELATED=24 or as per consumer	3	4	64	<feed-name>-FILTERED-CORRELATED=24 or as per consumer

3.4 Ephemeral Storage Requirements

The following table describes the Ephemeral Storage requirements for OCNADD:

Table 3-9 Ephemeral Storage Requirements

Service Name	Ephemeral Storage (min) in Mi	Ephemeral Storage (max) in Mi
<app-name>-adapter	200	800

Table 3-9 (Cont.) Ephemeral Storage Requirements

Service Name	Ephemeral Storage (min) in Mi	Ephemeral Storage (max) in Mi
<app-name>-gw	400	800 **Note** Supported only in release 22.4.0
ocnaddadminservice	100	200
ocnaddalarm	100	500
ocnaddhealthmonitoring	100	500
ocnaddscpaggregation	100	500
ocnaddseppaggregation	100	500 **Note** Supported from release 23.1.0 onwards
ocnaddnrfaggregation	100	500
ocnaddconfiguration	100	500

3.5 Disk Throughput Requirements

The following table describes the disk throughput requirements in OCNADD:

Table 3-10 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

Table 3-10 (Cont.) 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	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 and NRF:

Table 3-11 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 (in Bytes)	Rate	Avg Size (in Bytes)	Rate	Avg Size (in 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.

3.6 Kafka PVC Storage Requirements

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

Table 3-12 Retention Period Per Topic

Topic Name	Retention Period
SCP	5 Minutes

Table 3-12 (Cont.) Retention Period Per Topic

Topic Name	Retention Period
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 3-11 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

       = ~ 63.45 GB
   Storage Requirement for MAIN = MPS * Retention Period * RF * Message Size
   =  39000 * 6 Hours * 3 * 1941

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

       =  ~ 4.46 TB
   Total Storage Requirement for the Broker Cluster = Storage for SCP +
   Storage for NRF + Storage for MAIN

       = 63.45 GB + 4.46 TB

       = ~ 4.53 TB
   Total Storage for each broker = (4.53/Number of Brokers) TB = (4.53/3) TB
   = ~ 1.51 TB [Assuming 3 Broker cluster]
   ```

2. **Average Message Size = 3769 Bytes**

The following example uses the values from the fourth row of the Table 3-11 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

             = ~ 123.20 GB

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

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

             = ~ 8.66 TB
Total Storage Requirement for the Broker Cluster = Storage for SCP +
Storage for NRF + Storage for MAIN

             = 123.20 GB + 8.66 TB

             = ~ 8.79 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 Synthetic Feed with Replication
- Performance Benchmarking with 60K MPS SEPP Traffic with HTTP2 Feed Replication
- Performance Benchmarking with Correlated Feed for XDR Generation

4.1 Performance Benchmarking with Synthetic Feed with Replication

The performance benchmarking test is performed on OCNADD with 100K MPS traffic (feed replication enabled) which includes 55K MPS SCP traffic (27.5K TPS with 2 traffic copy trigger points on SCP), 15K MPS NRF traffic (7.5K TPS with 2 traffic copy trigger points on NRF) and 30K MPS SEPP traffic (15K 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: 100K 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 23.4.0 deployed over CNE (Bare metal) 22.3.0 version
- Kafka Replication Factor (RF): 1
- OCNADD TCP Feed: HTTP2/TLS
- OCNADD Kafka PVC: 400GB
  For more information about PVC requirements, see Kafka PVC Storage Requirements.
- 3rd Party Application: With 2 End-points
- cnDBTier Release 23.4.0
- NRF Release 23.4.0
- SCP Release 23.3.0
- SEPP Release 23.4.0
- SASL/SSL, enabled between NRF/SCP/SEPP and OCNADD

- CNCC Release 23.4.0
- Message Size: 3580 Bytes
- Environment: CNE (Bare metal)
- Execution time: 12 Hours
- NRF: IGW and EGW messageCopy feature is enabled
- SCP: Default configuration
- cSEPP: PLMN IGW messageCopy feature is enabled
- pSEPP: The messageCopy feature is disabled

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: 100K 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
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	-
ocnadduirouter	2	2	1	1	1	1	-
ocnaddscpaggregation	2	2	2	2	2	2	18 (SCP)/ retention.ms = 300000
ocnaddnrfaggregation	2	2	2	2	1	1	6 (NRF)/ retention.ms = 300000
ocnaddseppaggregation	2	2	2	2	2	2	12 (SEPP)/ retention.ms = 300000

Table 4-1 (Cont.) 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
ocnaddadapter	3	3	4	4	9	9	54 (MAIN)/ retention.ms = 300000
ocnaddkafka	10	6	96	48	4	4	-
zookeeper	1	1	2	2	3	3	-
ocnaddgui	2	2	1	1	1	1	-

- **SCP**
  SCP Traffic: 55K TPS using two trigger points.
- **NRF**
  NRF Traffic: 15K TPS using two trigger points.
- **SEPP**
  SEPP Traffic: 30K 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	100K MPS	12	Feed-1: 4.17 msec Feed-2: 4.07 sec	Feed-1: 99.95% Feed-2: 99.95%
NRF	2 - IGW/EGW (Request and Response)	15K MPS	12	NA	99.93%
SCP	2 (Request Ingress/ Request Egress)	55K MPS	12	NA	99.99%
SEPP	2 (PLMN IGW) (Request and Response)	30K MPS	12	NA	99.93%

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 (%)
ocnaddadminservice	1	0.16%	28.60%
ocnaddalarm	1	0.26%	18.60%
ocnaddconfiguration	1	0.15%	35.00%
ocnaddgui	1	0.00%	3.47%
ocnaddhealthmonitoring	1	0.40%	35.30%
ocnadduirouter	1	0.01%	27.80%
feed-dd1-adapter	9	61.00%	29.70%
feed-dd2-adapter	9	62.00%	28.92%
kafka-broker	4	87.70%	72.50%
zookeeper	4	0.16%	17.00%
ocnaddfilter	4	49.30%	13.90%
ocnaddnrfaggregation	1	62.40%	64.70%
ocnaddscpaggregation	3	71.60%	66.20%
ocnaddseppaggregation	2	54.20%	52.50%

4.2 Performance Benchmarking with Correlated Feed for XDR Generation

The performance benchmarking test is performed on OCNADD with correlated feed for XDR generation in a centralized deployment. The traffic includes 20K MPS SCP traffic (10K TPS with 2 traffic copy trigger points on SCP), 4K MPS NRF traffic (2K TPS with 2 traffic copy trigger points on NRF), and 6K MPS SEPP traffic (3K 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: 100K 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 23.4.0 deployed over CNE (Bare metal) 22.3.0 version
- Kafka Replication Factor (RF): 1
- OCNADD Correlated Feed: HTTP2/TLS
- OCNADD Kafka PVC: 400GB
  For more information about PVC requirements, see Kafka PVC Storage Requirements.
- 3rd Party Application: With 2 End-points

- cnDBTier Release 23.4.0
- NRF Release 23.4.0
- SCP Release 23.3.0
- SEPP Release 23.4.0
- SASL/SSL, enabled between NRF/SCP/SEPP and OCNADD
- CNCC Release 23.4.0
- Message Size: 3300 bytes
- Environment: CNE (Bare metal)
- Execution time: 12 Hours
- NRF: IGW and EGW messageCopy feature is enabled
- SCP: Default configuration
- cSEPP: PLMN IGW messageCopy feature is enabled
- pSEPP: The messageCopy feature is disabled

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: 100K 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	-
ocnadduiouter	2	2	1	1	1	1	-
ocnaddscpaggregation	2	2	2	2	2	2	18 (SCP)/ retention.ms = 300000

Table 4-4 (Cont.) 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
ocnaddnrfaggregation	2	2	2	2	1	1	6 (NRF)/ retention.ms = 300000
ocnaddseppaggregation	2	2	2	2	2	2	12 (SEPP)/ retention.ms = 300000
ocnaddadapter	3	3	4	4	9	9	54 (MAIN)/ retention.ms = 300000
ocnaddkafka	8	8	64	64	4	4	-
zookeeper	1	1	2	2	3	3	-
ocnaddgui	2	2	1	1	1	1	-

- **SCP**
  SCP Traffic: 20K TPS using two trigger points.
- **NRF**
  NRF Traffic: 4K TPS using two trigger points.
- **SEPP**
  SEPP Traffic: 6K 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	12	NA	99.99%
NRF	2 - IGW/EGW (Request and Response)	4K MPS	12	NA	99.99%
SCP	2 (Request Ingress/ Request Egress)	20K MPS	12	NA	99.99%
SEPP	2 (PLMN IGW) (Request and Response)	6K MPS	12	NA	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 (%)
ocnaddadminservice	1	0.13%	29.40%
ocnaddalarm	1	0.34%	21.80%
ocnaddconfiguration	1	0.14%	34.80%
ocnaddgui	1	0.00%	5.09%
ocnaddhealthmonitoring	1	0.29%	35.30%
ocnadduirouter	1	0.06%	29.20%
kafka-broker	4	52.10%	35.90%
zookeeper	4	0.14%	18.10%
ocnaddfilter	4	49.50%	15.40%
ocnaddnrfaggregation	1	63.60%	55.50%
ocnaddscpaggregation	3	49.20%	46.60%
ocnaddseppaggregation	2	39.00%	39.10%
feed-dd1-correlation	16	23.40%	26.50%

4.3 Performance Benchmarking with 60K MPS SEPP Traffic with HTTP2 Feed Replication

The performance benchmarking test is performed on OCNADD with SEPP HTTP2 feed (replication enabled) in a centralized deployment with one worker group. The traffic includes 60K MPS SEPP traffic (15K TPS with 4 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: 60K 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 23.4.0 deployed over CNE (Bare metal) 22.3.0 version
- Kafka Replication Factor (RF): 1
- OCNADD Data Feed: HTTP2/TLS
- OCNADD Kafka PVC: 400GB
  For more information about PVC requirements, see Kafka PVC Storage Requirements.
- 3rd Party Application: With 2 End-points
- cnDBTier Release 23.4.0

- SEPP Release 23.4.0
- SASL/SSL, enabled between SEPP and OCNADD
- CNCC Release 23.4.0
- Message Size: 3300 bytes
- Environment: CNE (Bare metal)
- Execution time: 12 Hours
- cSEPP: PLMN IGW and N32 EGW messageCopy feature enabled
- pSEPP: The messageCopy feature is disabled

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: 60K MPS

  Resource Specifications:

Table 4-7 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	-
ocnadduirouter	2	2	1	1	1	1	-
ocnaddscpaggregation	2	2	2	2	4	4	6 (SCP)/ retention.ms = 300000
ocnaddnrfaggregation	2	2	2	2	1	1	6 (NRF)/ retention.ms = 300000
ocnaddseppaggregation	2	2	2	2	1	1	24 (SEPP)/ retention.ms = 300000

Table 4-7 (Cont.) 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
ocnaddadapter	3	3	4	4	14	14	126 (MAIN)/ retention.ms = 300000
ocnaddkafka	8	8	64	64	4	4	-
zookeeper	1	1	2	2	3	3	-
ocnaddgui	2	2	1	1	1	1	-

- **SEPP**
  SEPP Traffic: 60K TPS using four trigger points.

Benchmark Test Results

Traffic Feed Details

Table 4-8 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	60K MPS	12	Feed-1: 2.39 msec Feed-2: 2.27 sec	Feed-1: 99.99% Feed-2: 99.99%
SEPP	2 (PLMN IGW) (Request and Response) 2 - N32 EGW (Request and Response)	60K MPS	12	NA	99.99%

CPU and Memory Utilization

The following table describes the OCNADD CPU and memory utilization:

Table 4-9 OCNADD CPU and Memory Utilization

Microservice/ Container	Pod Count	CPU Utilization (%)	Memory Utilization (%)
ocnaddadminservice	1	0.15%	29.00%
ocnaddalarm	1	1.42%	22.00%
ocnaddconfiguration	1	0.16%	35.20%
ocnaddgui	1	0.00%	3.02%
ocnaddhealthmonitoring	1	0.35%	37.70%
ocnadduirouter	1	0.23%	30.50%

Table 4-9 (Cont.) OCNADD CPU and Memory Utilization

Microservice/ Container	Pod Count	CPU Utilization (%)	Memory Utilization (%)
feed-dd1-adapter	14	42.50%	17.00%
kafka-broker	4	84.10%	39.70%
zookeeper	4	0.16%	16.70%
feed-dd2-adapter	14	41.00%	17.60%
ocnaddfilter	4	49.40%	15.20%
ocnaddnrfaggregation	1	0.06%	14.10%
ocnaddscpaggregation	1	0.06%	14.80%
ocnaddseppaggregation	4	75.20%	59.50%