Oracle® Communications

Cloud Native Core, Network Exposure Function Troubleshooting Guide

Release 24.2.2
G10005-03
September 2025

ORACLE

Oracle Communications Cloud Native Core, Network Exposure Function Troubleshooting Guide, Release 24.2.2

G10005-03

Contents

5 NEF Alerts

6 CAPIF Alerts

My Oracle Support

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

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

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

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

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

Acronyms

The following table lists the acronyms and the terminologies used in the document:

Table Acronyms and Terminologies

Acronym	Description
3GPP	3rd Generation Partnership Project
5GC	5G Core Network
5GS	5G System
AF	Application Function
API	Application Programming Interface
CAPIF	Oracle Communications Common API Framework
CNC	Oracle Communications Cloud Native Core
CNE	Oracle Communications Cloud Native Core, Cloud Native Environment
ECUR	Event Charging with Unit Reservation
FQDN	Fully Qualified Domain Name
GPSI	Generic Public Subscription Identifier
GMLC	Gateway Mobile Location Center
HA	High Availability
IEC	Immediate Event Charging
IGW	Internet Gateway
IMSI	International Mobile Subscriber Identity
K8s	Kubernetes
ME	Monitoring Events
MSISDN	Mobile Station International Subscriber Directory Number
NEF	Oracle Communications Cloud Native Core, Network Exposure Function
NF	Network Function
NRF	Oracle Communications Cloud Native Core, Network Repository Function
NSSF	Oracle Communications Cloud Native Core, Network Slice Selection Function
OAM	Operations, Administration, and Maintenance
OCI	Oracle Cloud Infrastructure
OKE	Oracle Kubernetes Engine
PDB	Pod Disruption Budget
QoS	Quality of Service
SBA	Service Based Architecture
SBI	Service Based Interface
SCEF	Service Capability Exposure Function
S-NSSAI	Single Network Slice Selection Assistance Information
SUPI	Subscription Permanent Identifier
TPS	Traffic Per Second
UDM	Unified Data Management

Table (Cont.) Acronyms and Terminologies

Acronym	Description
URI	Uniform Resource Identifier

What's New in This Guide

This section introduces the documentation updates for Release 24.2.x.

Release 24.2.2 - G10005-03, September 2025

There are no changes in this release.

Release 24.2.1 - G10005-02, October 2024

There are no changes in this release.

Release 24.2.0 - G10005-01, July 2024

The following changes are made in this release:

- Added the following alerts to configure OCI as part of the Deployment in OCI feature:
  - OccapifTotalExternalIngressTrafficRateAboveMinorThreshold
  - OccapifTotalNetworkIngressTrafficRateAboveMinorThreshold
  - OccapifTotalExternalIngressTrafficRateAboveMajorThreshold
  - OccapifTotalNetworkIngressTrafficRateAboveMajorThreshold
  - OccapifTotalExternalIngressTrafficRateAboveCriticalThreshold
  - OccapifTotalNetworkIngressTrafficRateAboveCriticalThreshold
  - OccapifExternalIngressTransactionErrorRateAboveZeroPointOnePercent
  - OccapifNetworkIngressTransactionErrorRateAboveZeroPointOnePercent
  - OccapifExternalIngressTransactionErrorRateAbove1Percent
  - OccapifNetworkIngressTransactionErrorRateAbove1Percent
  - OccapifExternalIngressTransactionErrorRateAbove10Percent
  - OccapifNetworkIngressTransactionErrorRateAbove10Percent
  - OccapifExternalIngressTransactionErrorRateAbove25Percent
  - OccapifNetworkIngressTransactionErrorRateAbove25Percent
  - OccapifExternalIngressTransactionErrorRateAbove50Percent
  - OccapifNetworkIngressTransactionErrorRateAbove50Percent
  - OcnefTotalExternalIngressTrafficRateAboveMinorThreshold
  - OcnefTotalFivegcIngressTrafficRateAboveMinorThreshold
  - OcnefTotalExternalIngressTrafficRateAboveMajorThreshold
  - OcnefTotalFivegcIngressTrafficRateAboveMajorThreshold
  - OcnefTotalExternalIngressTrafficRateAboveCriticalThreshold
  - OcnefTotalFivegcIngressTrafficRateAboveCriticalThreshold
  - OcnefExternalIngressTransactionErrorRateAboveZeroPointOnePercent
  - OcnefFivegcIngressTransactionErrorRateAboveZeroPointOnePercent
  - OcnefExternalIngressTransactionErrorRateAbove1Percent

  - OcnefFivegcIngressTransactionErrorRateAbove1Percent
  - OcnefExternalIngressTransactionErrorRateAbove10Percent
  - OcnefFivegcIngressTransactionErrorRateAbove10Percent
  - OcnefExternalIngressTransactionErrorRateAbove25Percent
  - OcnefFivegcIngressTransactionErrorRateAbove25Percent
  - OcnefExternalIngressTransactionErrorRateAbove50Percent
  - OcnefFivegcIngressTransactionErrorRateAbove50Percent
- Removed the following alerts as part of the Deployment in OCI feature:
  - OccapifTotalIngressTrafficRateAboveMinorThreshold
  - OccapifTotalIngressTrafficRateAboveMajorThreshold
  - OccapifTotalIngressTrafficRateAboveCriticalThreshold
  - OccapifTransactionErrorRateAbove0.1Percent
  - OccapifTransactionErrorRateAbove1Percent
  - OccapifTransactionErrorRateAbove10Percent
  - OccapifTransactionErrorRateAbove25Percent
  - OccapifTransactionErrorRateAbove50Percent
  - OcnefTotalIngressTrafficRateAboveMinorThreshold
  - OcnefTotalIngressTrafficRateAboveMajorThreshold
  - OcnefTotalIngressTrafficRateAboveCriticalThreshold
  - OcnefTransactionErrorRateAbove0.1Percent
  - OcnefTransactionErrorRateAbove1Percent
  - OcnefTransactionErrorRateAbove10Percent
  - OcnefTransactionErrorRateAbove25Percent
  - OcnefTransactionErrorRateAbove50Percent

1 Introduction

This document provides information about troubleshooting Oracle Communications Network Exposure Function (NEF).

1.1 Overview

NEF is an important component of the 5G Service Based Architecture. NEF provides a platform to securely expose the network services and capabilities offered by the 5G Network Functions (NFs) to either external third-party applications or the internal Application Functions (AFs). As an interface between 5G core network and different application functions, NEF enables the operators to customize their network and provide innovative services to the end users.

> **Note**
>
> The performance and capacity of the NEF system may vary based on the call model, feature or interface configuration, and underlying CNE and hardware environment.

1.2 Audience

The intended audiences for this document are the network administrators and the professionals responsible for NEF deployment and maintenance.

1.3 Reference

For more information about NEF, refer to the following documents:

- Oracle Communications Cloud Native Core, Network Exposure Function Installation, Upgrade, and Fault Recovery Guide
- Oracle Communications Cloud Native Core, Network Exposure Function User Guide
- Oracle Communications Cloud Native Core, Cloud Native Environment Installation, Upgrade, and Fault Recovery Guide

2
Logs

This chapter explains the process to retrieve the logs and status that can be used for effective troubleshooting.

2.1 Log Levels

Logs register system events along with their date and time of occurrence. They also provide important details about a chain of events that could have led to an error or problem.

A log level helps in defining the severity level of a log message. For NEF, the log level of a microservice can be set to any one of the following valid values:

- **TRACE:** A log level that describes events, as a step by step execution of code. This can be ignored during the standard operation, but may be useful during extended debugging sessions.
- **DEBUG:** A log level used for events during software debugging when more granular information is needed.
- **INFO:** A standard log level indicating that something has happened, an application has entered a certain state, etc.
- **WARN:** A log level indicates that something unexpected has happened in the application, a problem, or a situation that might disturb one of the processes. But this does not mean that the application has failed. The WARN level should be used in situations that are unexpected, but the code can continue to work.
- **ERROR:** A log level that should be used when an application hits an issue preventing one or more functionalities from functioning.

Using this information, the logs can be filtered based on the system requirements. For instance, if you want to filter the critical information about your system from the informational log messages, set a filter to view messages with only WARN log level in Kibana.

The following table provides log level details that may be helpful to handle different NEF debugging issues:

Table 2-1 Log Levels

Scenarios	Pod	Logs to be searched	Log Level
Registration with NRF Successful	nrf-client-service	Register completed successfully / "nfServiceStatus":"REGISTERED"	INFO
Heartbeat message log	nrf-client-service	Update completed successfully	INFO
NRF configurations reloading	nrf-client-service	NRF client config reloaded	INFO
Check for exiting NF Instance Entry	nrf-client-service	No registered NF instance exists	WARN

Table 2-1 (Cont.) Log Levels

Scenarios	Pod	Logs to be searched	Log Level
Started Application	nrf-client-service	Successful application start	INFO
Started Application	ocnef-monitoringevents	Successful application start	INFO
NRF Client Config Initialized	nrf-client-service	Initialize NRF client configuration	INFO
FQDN/BASEURL/ livenessProbeUrl Improper	nrf-client-service	response=<503,java.net. UnknownHostException	WARN
nudr-drservice liveness probe failure	nrf-client-service	NFService liveness probe failed	WARN
SQL Exception during start up	ocnef-monitoringevents	java.sql.SQLException	WARN
DB connection pool Established	ocnef-monitoringevents	HikariPool-1 - Start completed	INFO
Error Code Mapping configurations loaded	ocnef-monitoringevents	Loaded Error Code Mapping Configuration	INFO
Error Code Mapping configurations loaded	ocnef-monitoringevents	Loaded Error Reason Mapping Configuration	INFO
Error Code Mapping configurations loaded	ocnef-monitoringevents	Loaded Error Title Mapping Configuration	INFO
Error Code Mapping configurations loaded	ocnef-monitoringevents	Loaded Error Type Mapping Configuration	INFO
Check if Ports successfully listening	ocnef-monitoringevents	Netty started on port	INFO
Check for message received	ocnef-monitoringevents	Entering SubscriptionsApiControll er	DEBUG
Check for message exception	ocnef-monitoringevents	Exception when executing	DEBUG
URI Pattern not supported	ocnef-monitoringevents	None match pattern found for URL	WARN
Check if Ports successfully listening	nrf-client-service	Undertow started on port(s)	INFO
Pod exit	ocnef-monitoringevents	HikariPool-1 - Shutdown completed	INFO
DB username/ DB password invalid	ocnef-monitoringevents	Access denied for user	WARN
Registration with NRF failed	nrf-client-service	Register failed	ERROR
De registration with NRF successful	nrf-client-service	Deregister completed successfully	INFO
De registration with NRF failed	nrf-client-service	Deregister failed	ERROR
NF Profile update failed	nrf-client-service	Update failed	ERROR

2.2 Collecting Logs

This section describes the steps to collect logs from PODs and containers. Perform the following steps:

1. Run the following command to get the PODs details:

   ```
   kubectl -n <namespace_name> get pods
   ```

2. Collect the logs from the specific pods or containers:

   ```
   kubectl logs <podname> -n <namespace> -c <containername>
   ```

3. Store the log in a file using the following command:

   ```
   kubectl logs <podname> -n <namespace>  <filename>
   ```

4. (Optional) You can also use the following commands for the log stream with file redirection starting with last 100 lines of log:

   ```
   kubectl logs <podname> -n <namespace> -f --tail <number of lines> >
   <filename>
   ```

For more information on how to collect the logs, see *Oracle Communication Cloud Native Core, Data Collector Guide*.

2.3 Understanding Logs

This chapter explains the logs you need to look into to handle different NEF debugging issues.

For more information on how to collect the logs, see *Oracle Communication Cloud Native Core, Data Collector Guide*.

This section provides log level attribute details for following services:

- monitoringevents
- fivegcagent
- expgw-apirouter
- expgw-afmgr

Sample Logs

Sample log statement monitoringevents:


```
{
"instant":
{
"epochSecond":1645710158,
"nanoOfSecond":717564000
},
"thread":"reactor-http-nio-2",
"level":"DEBUG",
"loggerName":"com.oracle.cne.nef.efc.svc.me.api.SubscriptionsApiController",
"message":"processSubscriptionPost: IN",
"endOfBatch":false,
"loggerFqcn":"org.apache.logging.log4j.spi.AbstractLogger",
"contextMap":{},"threadId":49,
"threadPriority":5,
"messageTimestamp":"2022-02-24T19:12:38.717+0530",
```

```
"application":"${env:CONFIGMAP_NAME}",
"microservice":"${env:CONFIGMAP_NAME}",
"vendor":"oracle",
"namespace":"${env:K8S_NAMESPACE}",
"node":"${env:K8S_NODE}",
"pod":"${env:POD_NAME}",
"instanceType":"prod"
}
```


Sample log statement fivegcagent:


```
{
"instant":
{
"epochSecond":1645710159,
"nanoOfSecond":347655000},
"thread":"reactor-http-nio-4",
"level":"DEBUG",
"loggerName":"com.oracle.cne.nef.efc.common.fivegcagent.api.FivegcAgentApiCont
roller",
"message":"Entering FivegcAgentApiController::processReceiverMessage",
"endOfBatch":false,
"loggerFqcn":"org.apache.logging.log4j.spi.AbstractLogger",
"contextMap":{},
"threadId":65,
"threadPriority":5,
"messageTimestamp":"2022-02-24T19:12:39.347+0530",
"application":"${env:CONFIGMAP_NAME}",
"microservice":"${env:CONFIGMAP_NAME}",
"vendor":"oracle",
"namespace":"${env:K8S_NAMESPACE}",
"node":"${env:K8S_NODE}",
"pod":"${env:POD_NAME}",
"instanceType":"prod"
}
```


Sample log statement expgw-apirouter:


```
{
"instant":
{
"epochSecond":1645679731,
"nanoOfSecond":795110229
},
"thread":"main",
"level":"WARN",
"loggerName":"io.opentracing.contrib.spring.tracer.configuration.TracerAutoCon
figuration",
"message":"Tracer bean is not configured! Switching to NoopTracer",
"endOfBatch":false,
"loggerFqcn":"org.apache.commons.logging.LogAdapter$Log4jLog",
"contextMap":{},
"threadId":1,
"threadPriority":5,
"messageTimestamp":"2022-02-24T05:15:31.795+0000",
```

```
"application":"ocnef-expgw-apirouter",
"microservice":"ocnef-expgw-apirouter",
"vendor":"oracle",
"namespace":"nef",
"node":"prowl-k8s-node-10",
"pod":"ocnef-expgw-apirouter-664b9bc79-5rv5n",
"instanceType":"prod"
}
```


Sample log statement expgw-afmgr:


```
{
"instant":
{
"epochSecond":1645710925,
"nanoOfSecond":734244575
},
"thread":"main",
"level":"INFO",
"loggerName":"com.zaxxer.hikari.HikariDataSource",
"message":"HikariPool-1 - Starting...",
"endOfBatch":false,
"loggerFqcn":"org.apache.logging.slf4j.Log4jLogger",
"contextMap":{},
"threadId":1,"threadPriority":5,
"messageTimestamp":"2022-02-24T13:55:25.734+0000",
"application":"ocnef-expgw-afmgr",
"microservice":"ocnef-expgw-afmgr",
"vendor":"oracle",
"namespace":"nef",
"node":"prowl-k8s-node-1",
"pod":"ocnef-expgw-afmgr-85b7fc5cb5-x9lfm",
"instanceType":"prod"
}
```


Table 2-2 Log Attribute Details

Log Attribute	Details	Sample Value	Data Type
instant	Epoch time **Note**: It is group of two values epochSecond and nanoOfSecond	{"epochSecond":1604655402," nanoOfSecond":946649000}	Object
thread	Logging Thread Name	"XNIO-1 task-1"	String
level	Log Level of the printed log	"WARN"	String
loggerName	Class or Module which printed the log	`"com.oracle.cne.nef.efc.svc.me.api.SubscriptionsApiController"`	String
message	Message related to the log providing brief details **Note**: Indicates that no NFProfiles found for mentioned search query	`"processSubscriptionPost : IN"`	String

Table 2-2 (Cont.) Log Attribute Details

Log Attribute	Details	Sample Value	Data Type
endOfBatch	Log4j2 Internal Default from log4j2: false	false	boolean
loggerFqcn	Log4j2 Internal Fully Qualified class name of logger module	org.apache.logging.log4j.spi.AbstractLogger	String
threadId	Thread Id generated internally by Log4j2	47	Integer
messageTimestamp	Timestamp of log from application container. Format: yyyy-MM-dd'T'HH:mm:ss.SSSZ	"2020-11-06T09:36:42.946+0000"	String

This section provides log level attribute details for ingress gateway service:
Sample log statement ingressgateway:

```
{
"thread":"ingress-h2c-epoll-3",
"level":"DEBUG",
"loggerName":"ocpm.cne.gateway.filters.PreGatewayFilter",
"message":"Exiting PreGatewayFilter",
"endOfBatch":false,
"loggerFqcn":"org.apache.logging.log4j.spi.AbstractLogger",
"instant":{"epochSecond":1604650229,
"nanoOfSecond":4993000},
"contextMap":
{
"hostname":"ocnrf-ingressgateway-69f6544b8d-cdbgx",
"ingressTxId":"ingress-tx-1087436877",
"ocLogId":"1604650229002_72_ocnrf-ingressgateway-69f6544b8d-cdbgx"
},
"threadId":72,
"threadPriority":5,
"messageTimestamp":"2020-11-06 08:10:29.004",
"ocLogId":"1604650229002_72_ocnrf-ingressgateway-69f6544b8d-cdbgx",
"pod":"ocnrf-ingressgateway-69f6544b8d-cdbgx",
"processId":"1",
"instanceType":"prod",
"ingressTxId":"ingress-tx-1087436877"
}
```

Table 2-3 Log Attribute Details for ingressgateway

Log Attribute	Details	Sample Value	Data Type	Notes
thread	Logging Thread Name	"ingress-h2c-epoll-3"	String	
level	Log Level of the log printed	"DEBUG"	String	

Table 2-3 (Cont.) Log Attribute Details for ingressgateway

Log Attribute	Details	Sample Value	Data Type	Notes
loggerName	Class/Module which printed the log	"ocpm.cne.gateway.filters.PreGatewayFilter"	String	
message	Message related to the log providing brief details	"Exiting PreGatewayFilter"	String	Indicates that the method PreGatewayFilter is being exited.
endOfBatch	Log4j2 Internal Default from log4j2: false	false	boolean	
loggerFqcn	Log4j2 Internal Fully Qualified class name of logger module	org.apache.logging.log4j.spi.AbstractLogger	String	
instant	Epoch timestamp	{"epochSecond":1604650229,"nanoOfSecond":4993000}	Object	It is group of two values epochSecond and nanoOfSecond
contextMap	contents of log4j ThreadContext map	{"hostname":"ocnrf-ingressgateway-69f6544b8d-cdbgx", "ingressTxId":"ingress-tx-1087436877", "ocLogId":"1604650229002_72_ocnrf-ingressgateway-69f6544b8d-cdbgx"}	Object	
threadId	Thread Id generated internally by Log4j2	72	Integer	
threadPriority	Thread Priority set internally by Log4j2	5	Integer	
messageTimestamp	Timestamp of log from application container. Format: yyyy-MM-dd'T'HH:mm:ss.SSSZ	"2020-11-06 08:10:29.004"	String	
ocLogId	End to End Log Identifier across the NRF microservices.	"1604650229002_72_ocnrf-ingressgateway-69f6544b8d-cdbgx"	String	Helps to correlate the logs across the microservices in NRF application
pod	Pod Name	"ocnrf-ingressgateway-69f6544b8d-cdbgx"	String	
processId	Process ID internally assigned	"1"	String	
instanceType	Instance type	"prod"	String	
ingressTxId	Transaction id that is added to log4j ThreadContext map and is unique to every transaction	"ingress-tx-1087436877"	String	

This section provides log level attribute details for the appinfo service:

Sample log statement appinfo:

```
{
"name": "gunicorn.access",
"message": "127.0.0.1 - - [29/Oct/2020:18:39:35 +0000]\"GET /v1/liveness HTTP/
1.1\" 200 16 \"-\"\"kube-probe/1.17+\"",
"level": "INFO",
"filename":"glogging.py",
"lineno": 344,
"module": "glogging",
"func":"access",
"thread":"MainThread",
"created": "2020-10-29T18:39:35.799448"
}
```

Table 2-4 Log Attribute Details for egressgateway

Log Attribute	Details	Sample Value	Data Type
name	Logger name	"gunicorn.access"	String
message	Message related to the log providing brief details	"message": "127.0.0.1 - - [29/Oct/2020:18:39:35 +0000] \"GET /v1/liveness HTTP/1.1\" 200 16 \"-\" \"kube-probe/1.17+ \""	String
level	Log Level of the log printed	"INFO"	String
filename	File name which generates the log	"glogging.py"	String
lineno	Line number	344	Integer
module	Python module name	"glogging"	String
func	Function Name	"func"	String
thread	Thread Name	"MainThread"	String
messageTimestamp	Timestamp of log from application container. Format: yyyy-MM-dd'T'HH:mm:ss.SSSZ	"2020-10-29T18:39:35.799448"	String

Common Useful log attributes

The below log attributes will be available only through Kibana. These attribute names are part of Kubernetes labels which are added in each NEF POD.

Table 2-5 Common useful log attributes

Log Attribute	Details	Sample Value	Data Type	Notes
application	NF Type	"ocnef"	String	complete attribute name in Kibana: kubernetes.labels.application

Table 2-5 (Cont.) Common useful log attributes

Log Attribute	Details	Sample Value	Data Type	Notes
microservice	Kubernetes service name Format: <Helm-Release-Name>-<Microservice-Name>	(assuming the <Helm-Release-Name> is "nefnorth" "nefnorth-monitoringevents" "nefnorth-fivegcagent" "nefnorth-expgw-afmgr" "nefnorth-expgw-apirouter" "nefnorth-appinfo" "nefnorth-ext-ingress-gateway" "nefnorth-ext-egress-gateway"	String	
engVersion	Engineering version	"1.9.0"	String	
mktgVersion	Marketing version	"1.9.1.0.0"	String	
vendor	Vendor Name	"Oracle"	Integer	

3 Using Debug Tool

The Debug Tool provides third-party troubleshooting tools for debugging the runtime issues in a lab environment. The following tools are available to debug NEF issues:

- tcpdump
- ip
- netstat
- curl
- ping
- nmap
- dig

Prerequisites

This section explains the prerequisites for using the debug tool.

> **Note**
>
> - For CNE 23.2.0 and later versions, follow Step a of Configuration in CNE.
> - For CNE versions prior to 23.2.0, follow Step b of Configuration in CNE.

1. **Configuration in CNE**
   Perform the following configurations in the Bastion Host. You need admin privileges to perform these configurations.
   a. When NEF is installed on CNE version 23.2.0 or above

      > **Note**
      >
      > - In CNE version 23.2.0 or above, the default CNE Kyverno policy, disallow-capabilities, do not allow NET_ADMIN and NET_RAW capabilities that are required for debug tool.
      > - To run Debug tool on CNE 23.2.0 and above, the user must modify the existing Kyverno policy, disallow-capabilities, as below.

      #### Adding a Namespace to an Empty Resource

      i. Run the following command to verify if the current disallow-capabilities cluster policy has namespace in it.
      Example:

      ```
      $ kubectl get clusterpolicies disallow-capabilities -oyaml
      ```

Sample output:

```
apiVersion: kyverno.io/v1
kind: ClusterPolicy
...
...
spec:
  rules:
  -exclude:
      any:
      -resources:{}
```

ii. If there are no namespaces, then patch the policy using the following command to add <namespace> under resources:

```
$ kubectl patch clusterpolicy disallow-capabilities --type=json \
  -p='[{"op": "add", "path": "/spec/rules/0/exclude/any/0/
resources", "value": {"namespaces":["<namespace>"]} }]'
```

Example:

```
$ kubectl patch clusterpolicy disallow-capabilities --type=json \
  -p='[{"op": "add", "path": "/spec/rules/0/exclude/any/0/
resources", "value": {"namespaces":["ocnef"]} }]'
```

Sample output:

```
apiVersion: kyverno.io/v1
kind: ClusterPolicy
...
...
spec:
  rules:
  -exclude:
      resources:
        namespaces:
        -ocnef
```

iii. If in case it is needed to remove the namespace added in the above step, use the following command:

```
$ kubectl patch clusterpolicy disallow-capabilities --type=json \
  -p='[{"op": "replace", "path": "/spec/rules/0/exclude/any/0/
resources", "value": {} }]'
```

Sample output:

```
apiVersion: kyverno.io/v1
kind: ClusterPolicy
...
...
spec:
  rules:
  -exclude:
```

```
      any:
      -resources:{}
```

Adding a Namespace to an Existing Namespace List

i. Run the following command to verify if the current disallow-capabilities cluster policy has namespaces in it.
Example:

```
$ kubectl get clusterpolicies disallow-capabilities -oyaml
```

Sample output:

```
apiVersion: kyverno.io/v1
kind: ClusterPolicy
...
...
spec:
  rules:
  -exclude:
      any:
      -resources:
          namespaces:
          -namespace1
          -namespace2
          -namespace3
```

ii. If there are namespaces already added, then patch the policy using the following command to add <namespace> to the existing list:

```
$ kubectl patch clusterpolicy disallow-capabilities --type=json \
  -p='[{"op": "add", "path": "/spec/rules/0/exclude/any/0/resources/
namespaces/-", "value": "<namespace>" }]'
```

Example:

```
$ kubectl patch clusterpolicy disallow-capabilities --type=json \
  -p='[{"op": "add", "path": "/spec/rules/0/exclude/any/0/resources/
namespaces/-", "value": "ocnef" }]'
```

Sample output:

```
apiVersion: kyverno.io/v1
kind: ClusterPolicy
...
...
spec:
  rules:
  -exclude:
      resources:
        namespaces:
        -namespace1
        -namespace2
```

```
        -namespace3
        -ocnef
```

iii. If in case it is needed to remove the namespace added in the above step, use the following command:

```
$ kubectl patch clusterpolicy disallow-capabilities --type=json \
  -p='[{"op": "remove", "path": "/spec/rules/0/exclude/any/0/
resources/namespaces/<index>"}]'
```

Example:

```
$ kubectl patch clusterpolicy disallow-capabilities --type=json \
  -p='[{"op": "remove", "path": "/spec/rules/0/exclude/any/0/
resources/namespaces/3"}]'
```

Sample output:

```
apiVersion: kyverno.io/v1
kind: ClusterPolicy
...
...
spec:
  rules:
  -exclude:
      resources:
        namespaces:
        -namespace1
        -namespace2
        -namespace3
```

> **Note**
>
> While removing the namespace, provide the index value for namespace within the array. The index starts from '0'.

b. When NEF is installed on CNE version prior to 23.2.0
PodSecurityPolicy (PSP) Creation

Create a PSP by running the following command from the bastion host. The parameters **readOnlyRootFileSystem, allowPrivilegeEscalation, allowedCapabilities** are required by debug container.

> **Note**
>
> Other parameters are mandatory for PSP creation and can be customized as per the CNE environment. The default values are recommended.

```
$ kubectl apply -f - <<EOF

apiVersion: policy/v1beta1
```

```
kind: PodSecurityPolicy
metadata:
  name: debug-tool-psp
spec:
  readOnlyRootFilesystem: false
  allowPrivilegeEscalation: true
  allowedCapabilities:
  - NET_ADMIN
  - NET_RAW
  fsGroup:
    ranges:
    - max: 65535
      min: 1
    rule: MustRunAs
  runAsUser:
    rule: MustRunAsNonRoot
  seLinux:
    rule: RunAsAny
  supplementalGroups:
    rule: RunAsAny
  volumes:
  - configMap
  - downwardAPI
  - emptyDir
  - persistentVolumeClaim
  - projected
  - secret
EOF
```

Role Creation

Run the following command to create a role for the PSP:

```
kubectl apply -f - <<EOF
apiVersion: rbac.authorization.k8s.io/v1
kind: Role
metadata:
  name: debug-tool-role
  namespace: cncc
rules:
- apiGroups:
  - policy
  resources:
  - podsecuritypolicies
  verbs:
  - use
  resourceNames:
  - debug-tool-psp
EOF
```

RoleBinding Creation

Run the following command to associate the service account for the NEF namespace with the role created for the PSP:

```
$ kubectl apply -f - <<EOF
apiVersion: rbac.authorization.k8s.io/v1
kind: RoleBinding
metadata:
  name: debug-tool-rolebinding
  namespace: ocnef
roleRef:
  apiGroup: rbac.authorization.k8s.io
  kind: Role
  name: debug-tool-role
subjects:
- kind: Group
  apiGroup: rbac.authorization.k8s.io
  name: system:serviceaccounts
EOF
```

2. **Configuration in NF specific Helm**
   Following updates must be performed in `custom-values.yaml` file.
   a. Log in to the NEF server.
   b. Open the custom-values file:

```
$ vim <custom-values file>
```

   c. Under global configuration, add the following:

```
# Allowed Values: DISABLED, ENABLED
extraContainers: ENABLED
debugToolContainerMemoryLimit: 4Gi
extraContainersVolumesTpl: |
 - name: debug-tools-dir
   emptyDir:
     medium: Memory
     sizeLimit: {{ .Values.global.debugToolContainerMemoryLimit |
quote }}
extraContainersTpl: |
  - command:
      - /bin/sleep
      - infinity
    image: {{ .Values.global.dockerRegistry }}/ocdebug-tools:23.3.1
    imagePullPolicy: Always
    name: {{ printf "%s-tools-%s"(include "getprefix".) (include
"getsuffix".) | trunc 63| trimPrefix "-"| trimSuffix "-"  }}
    resources:
      requests:
        ephemeral-storage: "512Mi"
        cpu: "0.5"
        memory: {{ .Values.global.debugToolContainerMemoryLimit |
quote }}
      limits:
        ephemeral-storage: "512Mi"
        cpu: "0.5"
        memory: {{ .Values.global.debugToolContainerMemoryLimit |
```

```
quote }}
    securityContext:
      allowPrivilegeEscalation: true
      capabilities:
        drop:
        - ALL
        add:
        - NET_RAW
        - NET_ADMIN
      runAsUser: 7000
    volumeMounts:
    - mountPath: /tmp/tools
      name: debug-tools-dir
```

Note

- Debug Tool Container comes up with the default user ID - 7000. If the operator wants to override this default value, it can be done using the `runAsUser` field, otherwise, the field can be skipped.
  Default value: uid=7000(debugtool) gid=7000(debugtool) groups=7000(debugtool)
- In case you want to customize the container name, replace the `name` field in the above values.yaml with the following:

  ```
  name: {{ printf "%s-tools-%s" (include "getprefix" .)
  (include "getsuffix" .) | trunc 63 | trimPrefix "-" |
  trimSuffix "-"  }}
  ```

  This will ensure that the container name is prefixed and suffixed with the necessary values.

For more information on how to customize parameters in the custom yaml value files, see *Oracle Communications Cloud Native Core, Network Exposure Function Installation, Upgrade and Fault Recovery Guide*.

d. Under service specific configurations for which debugging is required, add the following:

```
# Allowed Values: DISABLED, ENABLED, USE_GLOBAL_VALUE
extraContainers: USE_GLOBAL_VALUE
```

> **Note**
>
> - At the global level, `extraContainers` flag can be used to enable/disable injecting extra containers globally. This ensures that all the services that use this global value have extra containers enabled/disabled using a single flag.
> - At the service level, `extraContainers` flag determines whether to use the extra container configuration from the global level or enable/disable injecting extra containers for the specific service.

Run the Debug Tool

To run Debug Tool, perform the following steps:

1. Run the following command to retrieve the POD details:

   ```
   $ kubectl get pods -n <k8s namespace>
   ```

   Example:

   ```
   $ kubectl get pods -n ocnef
   ```

2. Run the following command to enter into Debug Tool Container:

   ```
   $ kubectl exec -it <pod name> -c <debug_container name> -n <namespace> bash
   ```

   Example:

   ```
   $ kubectl exec -it ocnef-nfaccesstoken-49fb96494c-k8w9q -c tools -n ocnef bash
   ```

3. Run the debug tools:

   ```
   bash -4.2$ <debug_tools>
   ```

   Example:

   ```
   bash -4.2$ tcpdump
   ```

4. Copy the output files from container to host:

   ```
   $ kubectl cp -c <debug_container name> <pod name>:<file location in
   container> -n <namespace> <destination location>
   ```

   Example:

   ```
   $ kubectl cp -c tools ocnef-nfaccesstoken-49fb96494c-k8w9q:/tmp/tools/
   capture.pcap -n ocnef /tmp/tools/
   ```

Tools Tested in Debug Container

Following is the list of debug tools that are tested.

Table 3-1 tcpdump

Options Tested	Description	Output	Capabilities
-D	Print the list of the network interfaces available on the system and on which *tcpdump* can capture packets.	`tcpdump -D` 1. eth02. 2. nflog (Linux netfilter log (NFLOG) interface) 3. nfqueue (Linux netfilter queue (NFQUEUE) interface) 4. any (Pseudo-device that captures on all interfaces) 5. lo [Loopback]	NET_ADMIN, NET_RAW
-i	Listen on *interface*	`tcpdump -i eth0` tcpdump: verbose output suppressed, use -v or -vv for full protocol decodelistening on eth0, link-type EN10MB (Ethernet), capture size 262144 bytes12:10:37.381199 IP cncc-core-ingress-gateway-7ffc49bb7f-2kkhc.46519 > kubernetes.default.svc.cluster.local.https: Flags [P.], seq 1986927241:1986927276, ack 1334332290, win 626, options [nop,nop,TS val 849591834 ecr 849561833], length 3512:10:37.381952 IP cncc-core-ingress-gateway-7ffc49bb7f-2kkhc.45868 > kube-dns.kube-system.svc.cluster.local.domain: 62870+ PTR? 1.0.96.10.in-addr.arpa. (40)	NET_ADMIN, NET_RAW
-w	Write the raw packets to file rather than parsing and printing them out.	`tcpdump -w capture.pcap -i eth0`	NET_ADMIN, NET_RAW
-r	Read packets from *file* (which was created with the **-w** option).	`tcpdump -r capture.pcap` reading from file /tmp/capture.pcap, link-type EN10MB (Ethernet)12:13:07.381019 IP cncc-core-ingress-gateway-7ffc49bb7f-2kkhc.46519 > kubernetes.default.svc.cluster.local.https: Flags [P.], seq 1986927416:1986927451, ack 1334332445, win 626, options [nop,nop,TS val 849741834 ecr 849711834], length 3512:13:07.381194 IP kubernetes.default.svc.cluster.local.https > cncc-core-ingress-gateway-7ffc49bb7f-2kkhc.46519: Flags [P.], seq 1:32, ack 35, win 247, options [nop,nop,TS val 849741834 ecr 849741834], length 3112:13:07.381207 IP cncc-core-ingress-gateway-7ffc49bb7f-2kkhc.46519 > kubernetes.default.svc.cluster.local.https: Flags [.], ack 32, win 626, options [nop,nop,TS val 849741834 ecr 849741834], length 0	NET_ADMIN, NET_RAW

Table 3-2 ip

Options Tested	Description	Output	Capabilities
addr show	Look at protocol addresses.	`ip addr show` 1: lo: <LOOPBACK,UP,LOWER_UP> mtu 65536 qdisc noqueue state UNKNOWN group defaultlink/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00inet 127.0.0.1/8 scope host lovalid_lft forever preferred_lft forever2: tunl0@NONE: <NOARP> mtu 1480 qdisc noop state DOWN group defaultlink/ipip 0.0.0.0 brd 0.0.0.04: eth0@if190: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1440 qdisc noqueue state UP group defaultlink/ether aa:5a:27:8d:74:6f brd ff:ff:ff:ff:ff:ff link-netnsid 0inet 192.168.219.112/32 scope global eth0valid_lft forever preferred_lft forever	--
route show	List routes.	`ip route show` default via 169.254.1.1 dev eth0 169.254.1.1 dev eth0 scope link	--
addrlabel list	List address labels	`ip addrlabel list` prefix ::1/128 label 0 prefix ::/96 label 3 prefix ::ffff:0.0.0.0/96 label 4 prefix 2001::/32 label 6 prefix 2001:10::/28 label 7 prefix 3ffe::/16 label 12 prefix 2002::/16 label 2 prefix fec0::/10 label 11 prefix fc00::/7 label 5 prefix ::/0 label 1	--

Table 3-3 netstat

Options Tested	Description	Output	Capabilities
-a	Show both listening and non-listening (for TCP, this means established connections) sockets.	`netstat -a` Active Internet connections (servers and established)Proto Recv-Q Send-Q Local Address Foreign Address Statetcp 0 0 0.0.0.0:tproxy 0.0.0.0:* LISTENtcp 0 0 0.0.0.0:websm 0.0.0.0:* LISTENtcp 0 0 cncc-core-ingress:websm 10-178-254-194.ku:47292 TIME_WAITtcp 0 0 cncc-core-ingress:46519 kubernetes.defaul:https ESTABLISHEDtcp 0 0 cncc-core-ingress:websm 10-178-254-194.ku:47240 TIME_WAITtcp 0 0 cncc-core-ingress:websm 10-178-254-194.ku:47347 TIME_WAITudp 0 0 localhost:59351 localhost:ambit-lm ESTABLISHEDActive UNIX domain sockets (servers and established)Proto RefCnt Flags Type State I-Node Pathunix 2 [ ] STREAM CONNECTED 576064861	--

Table 3-3 (Cont.) netstat

Options Tested	Description	Output	Capabilities
-l	Show only listening sockets.	`netstat -l` Active Internet connections (only servers)Proto Recv-Q Send-Q Local Address Foreign Address Statetcp 0 0 0.0.0.0:tproxy 0.0.0.0:* LISTENtcp 0 0 0.0.0.0:websm 0.0.0.0:* LISTENActive UNIX domain sockets (only servers)Proto RefCnt Flags Type State I-Node Path	--
-s	Display summary statistics for each protocol.	`netstat -s` Ip:4070 total packets received0 forwarded0 incoming packets discarded4070 incoming packets delivered4315 requests sent outIcmp:0 ICMP messages received0 input ICMP message failed.ICMP input histogram:2 ICMP messages sent0 ICMP messages failedICMP output histogram:destination unreachable: 2	--
-i	Display a table of all network interfaces.	`netstat -i` Kernel Interface tableIface MTU RX-OK RX-ERR RX-DRP RX-OVR TX-OK TX-ERR TX-DRP TX-OVR Flgeth0 1440 4131 0 0 0 4355 0 0 0 BMRUlo 65536 0 0 0 0 0 0 0 0 LRU	--

Table 3-4 curl

Options Tested	Description	Output	Capabilities
-o	Write output to <file> instead of stdout.	`curl -o file.txt http://abc.com/file.txt`	--
-x	Use the specified HTTP proxy.	`curl -x proxy.com:8080 -o http://abc.com/file.txt`	--

Table 3-5 ping

Options Tested	Description	Output	Capabilities
<ip>	Run a ping test to see whether the target host is reachable or not.	`ping 10.178.254.194`	NET_ADMIN, NET_RAW
-c	Stop after sending 'c' number of ECHO_REQUEST packets.	`ping -c 5 10.178.254.194`	NET_ADMIN, NET_RAW
-f (with non zero interval)	Flood ping. For every ECHO_REQUEST sent, a period "." is printed, while for every ECHO_REPLY received a backspace is printed.	`ping -f -i 2 10.178.254.194`	NET_ADMIN, NET_RAW

Table 3-6 nmap

Options Tested	Description	Output	Capabilities

<ip>	Scan for Live hosts, Operating systems, packet filters and open ports running on remote hosts.	`nmap 10.178.254.194`
		
		`Starting Nmap 6.40 ( http://nmap.org ) at 2020-09-29 05:54 UTCNmap scan report for`
		
		`10-178-254-194.kubernetes.default.svc.cluster .local (10.178.254.194)Host is up (0.00046s latency).Not shown: 995 closed portsPORT STATE SERVICE22/tcp open ssh179/tcp open bgp6666/tcp open irc6667/tcp open irc30000/tcp open unknownNmap done: 1 IP address (1 host up) scanned in 0.04 seconds`
			--

Table 3-6 (Cont.) nmap

Options Tested	Description	Output	Capabilities

-v	Increase verbosity level	`nmap -v 10.178.254.194`
		
		`Starting Nmap 6.40 ( http://nmap.org ) at 2020-09-29 05:55 UTC`
		`Initiating Ping Scan at 05:55`
		`Scanning 10.178.254.194 [2 ports]`
		`Completed Ping Scan at 05:55, 0.00s elapsed (1 total hosts)`
		`Initiating Parallel DNS resolution of 1 host. at 05:55`
		`Completed Parallel DNS resolution of 1 host. at 05:55, 0.00s elapsed`
		`Initiating Connect Scan at 05:55`
		`Scanning 10-178-254-194.kubernetes.default.svc.cluster.local (10.178.254.194) [1000 ports]`
		`Discovered open port 22/tcp on 10.178.254.194`
		`Discovered open port 30000/tcp on 10.178.254.194`
		`Discovered open port 6667/tcp on 10.178.254.194`
		`Discovered open port 6666/tcp on 10.178.254.194`
		`Discovered open port 179/tcp on 10.178.254.194`
		`Completed Connect Scan at 05:55, 0.02s elapsed (1000 total ports)`
		`Nmap scan report for 10-178-254-194.kubernetes.default.svc.cluster.local (10.178.254.194)`
		`Host is up (0.00039s latency).`
		`Not shown: 995 closed ports`
		`PORT STATE SERVICE`
		`22/tcp open ssh`
		`179/tcp open bgp`
		`6666/tcp open irc`
		`6667/tcp open irc`
		`30000/tcp open unknown`
		
		`Read data files from: /usr/bin/../share/nmap`
		`Nmap done: 1 IP address (1 host up) scanned in 0.04 seconds`
			--

Table 3-6 (Cont.) nmap

Options Tested	Description	Output	Capabilities

-iL	Scan all the listed IP addresses in a file. Sample file	`nmap -iL sample.txt`
		
		`Starting Nmap 6.40 ( http://nmap.org ) at 2020-09-29 05:57 UTC`
		`Nmap scan report for localhost (127.0.0.1)`
		`Host is up (0.00036s latency).`
		`Other addresses for localhost (not scanned): 127.0.0.1`
		`Not shown: 998 closed ports`
		`PORT STATE SERVICE`
		`8081/tcp open blackice-icecap`
		`9090/tcp open zeus-admin`
		
		`Nmap scan report for 10-178-254-194.kubernetes.default.svc.cluster.local (10.178.254.194)`
		`Host is up (0.00040s latency).`
		`Not shown: 995 closed ports`
		`PORT STATE SERVICE`
		`22/tcp open ssh`
		`179/tcp open bgp`
		`6666/tcp open irc`
		`6667/tcp open irc`
		`30000/tcp open unknown`
		
		`Nmap done: 2 IP addresses (2 hosts up) scanned in 0.06 seconds`
			--

Table 3-7 dig

Options Tested	Description	Output	Capabilities
<ip>	It performs DNS lookups and displays the answers that are returned from the name server(s) that were queried.	`dig 10.178.254.194` **Note**: The IP should be reachable from inside the container.	--
-x	Query DNS Reverse Look-up.	`dig -x 10.178.254.194`	--

3.1 Debug Tool Configuration Parameters

Following are the parameters used to configure debug tool.

CNE Parameters

Table 3-8 CNE Parameters

Parameter	Description
apiVersion	APIVersion defines the version schema of this representation of an object.
kind	Kind is a string value representing the REST resource this object represents.
metadata	Standard object's metadata.
metadata.name	Name must be unique within a namespace.
spec	spec defines the policy enforced.
spec.readOnlyRootFilesystem	Controls whether the containers run with a read-only root filesystem (that is, no writable layer).
spec.allowPrivilegeEscalation	Gates whether or not a user is allowed to set the security context of a container to allowPrivilegeEscalation=true.
spec.allowedCapabilities	Provides a list of capabilities that are allowed to be added to a container.
spec.fsGroup	Controls the supplemental group applied to some volumes. RunAsAny allows any fsGroup ID to be specified.
spec.runAsUser	Controls which user ID the containers are run with. RunAsAny allows any runAsUser to be specified.
spec.seLinux	RunAsAny allows any seLinuxOptions to be specified.
spec.supplementalGroups	Controls which group IDs containers add. RunAsAny allows any supplementalGroups to be specified.
spec.volumes	Provides a list of allowed volume types. The allowable values correspond to the volume sources that are defined when creating a volume.

Role Creation Parameters

Table 3-9 Role Creation

Parameter	Description
apiVersion	APIVersion defines the versioned schema of this representation of an object.
kind	Kind is a string value representing the REST resource this object represents.
metadata	Standard object's metadata.
metadata.name	Name must be unique within a namespace.
metadata.namespace	Namespace defines the space within which each name must be unique.
rules	Rules holds all the PolicyRules for this Role
apiGroups	APIGroups is the name of the APIGroup that contains the resources.
rules.resources	Resources is a list of resources this rule applies to.
rules.verbs	Verbs is a list of Verbs that apply to ALL the ResourceKinds and AttributeRestrictions contained in this rule.
rules.resourceNames	ResourceNames is an optional allowed list of names that the rule applies to.

Table 3-10 Role Binding Creation

Parameter	Description
apiVersion	APIVersion defines the versioned schema of this representation of an object.
kind	Kind is a string value representing the REST resource this object represents.
metadata	Standard object's metadata.
metadata.name	Name must be unique within a namespace.
metadata.namespace	Namespace defines the space within which each name must be unique.
roleRef	RoleRef can reference a Role in the current namespace or a ClusterRole in the global namespace.
roleRef.apiGroup	APIGroup is the group for the resource being referenced
roleRef.kind	Kind is the type of resource being referenced
roleRef.name	Name is the name of resource being referenced
subjects	Subjects holds references to the objects the role applies to.
subjects.kind	Kind of object being referenced. Values defined by this API group are "User", "Group", and "ServiceAccount".
subjects.apiGroup	APIGroup holds the API group of the referenced subject.
subjects.name	Name of the object being referenced.

Debug Tool Configuration Parameters

Table 3-11 Debug Tool Configuration Parameters

Parameter	Description
extraContainers	Specifies the spawns debug container along with application container in the pod.
debugToolContainerMemoryLimit	Indicates the memory assigned for the debug tool container.
extraContainersVolumesTpl	Specifies the extra container template for the debug tool volume.
extraContainersVolumesTpl.name	Indicates the name of the volume for debug tool logs storage.
extraContainersVolumesTpl.emptyDir.medium	Indicates the location where `emptyDir` volume is stored.
extraContainersVolumesTpl.emptyDir.sizeLimit	Indicates the `emptyDir` volume size.
command	String array used for container command.
image	Docker image name
imagePullPolicy	Image Pull Policy
name	Name of the container
resources	Compute Resources required by this container
resources.limits	Limits describes the maximum amount of compute resources allowed
resources.requests	Requests describes the minimum amount of compute resources required
resources.limits.cpu	CPU limits
resources.limits.memory	Memory limits
resources.limits.ephemeral-storage	Ephemeral Storage limits

Table 3-11 (Cont.) Debug Tool Configuration Parameters

Parameter	Description
resources.requests.cpu	CPU requests
resources.requests.memory	Memory requests
resources.requests.ephemeral-storage	Ephemeral Storage requests
securityContext	Security options the container should run with.
securityContext.allowPrivilegeEscalation	AllowPrivilegeEscalation controls whether a process can gain more privileges than its parent process. This directly controls if the no_new_privs flag will be set on the container process
securityContext.capabilities	The capabilities to add/drop when running containers. Defaults to the default set of capabilities granted by the container runtime.
securityContext.capabilities.drop	Removed capabilities
secuirtyContext.capabilities.add	Added capabilities
securityContext.runAsUser	The UID to run the entrypoint of the container process.
volumeMounts.mountPath	Indicates the path for volume mount.
volumeMounts.name	Indicates the name of the directory for debug tool logs storage.

4 Troubleshooting NEF

This chapter provides information to troubleshoot the common errors which can be encountered during the preinstallation, installation, upgrade, and rollback procedures of NEF.

4.1 Generic Checklist

The following sections provide a generic checklist for troubleshooting tips.

Deployment related tips

Perform the following checks after the deployment:

- Are NEF deployment, pods, and services created?
  Are NEF deployment, pods, and services running and available?

  Run the following the command:

  ```
  # kubectl -n <namespace> get deployments,pods,svc
  ```

  Inspect the output, check the following columns:

  – AVAILABLE of deployment
  – READY, STATUS, and RESTARTS of a pod
  – PORT(S) of service
- Is the correct image used?
  Is the correct environment variables set in the deployment?

  Run the following command:

  ```
  # kubectl -n <namespace> get deployment <deployment-name> -o yaml
  ```

  Inspect the output, check the environment and image.

  ```
  # kubectl -n nef-svc get deployment ocnef-monitoringevents -o yaml
  apiVersion: extensions/v1beta1
  kind: Deployment
  metadata:
    annotations:
      deployment.kubernetes.io/revision: "1"
      kubectl.kubernetes.io/last-applied-configuration: |
        {"apiVersion":"apps/v1","kind":"Deployment","metadata":
  {"annotations":{},"name":"ocnef-monitoringevents","namespace":"nef-
  svc"},"spec":{"replicas":1,"selector":{"matchLabels":{"app":"ocnef-
  monitoringevents"}},"template":{"metadata":{"labels":{"app":"ocnef-
  monitoringevents"}},"spec":{"containers":[{"env":
  [{"name":"MYSQL_HOST","value":"mysql"},
  {"name":"MYSQL_PORT","value":"3306"},
  {"name":"MYSQL_DATABASE","value":"nefdb"},
  ```

```
{"name":"NEF_SVC_ENDPOINT","value":"ocnef-monitoringevents"}],"image":"cne-
repo:5000/ocnef-
monitoringevents:latest","imagePullPolicy":"Always","name":"ocnef-
monitoringevents","ports":[{"containerPort":8080,"name":"server"}]}]}}}}
  creationTimestamp: 2018-08-27T15:45:59Z
  generation: 1
  name: ocnef-monitoringevents
  namespace: nef-svc
  resourceVersion: "2336498"
  selfLink: /apis/extensions/v1beta1/namespaces/nef-svc/deployments/ocnef-
monitoringevents
  uid: 4b82fe89-aa10-11e8-95fd-fa163f20f9e2
spec:
  progressDeadlineSeconds: 600
  replicas: 1
  revisionHistoryLimit: 10
  selector:
    matchLabels:
      app: ocnef-monitoringevents
  strategy:
    rollingUpdate:
      maxSurge: 25%
      maxUnavailable: 25%
    type: RollingUpdate
  template:
    metadata:
      creationTimestamp: null
      labels:
        app: ocnef-monitoringevents
    spec:
      containers:
      - env:
        - name: MYSQL_HOST
          value: mysql
        - name: MYSQL_PORT
          value: "3306"
        - name: MYSQL_DATABASE
          value: nefdb
        - name: NRF_SVC_ENDPOINT
          value: ocnef-monitoringevents
        image: cne-repo:5000/ocnef-monitoringevents:latest
        imagePullPolicy: Always
        name: ocnef-monitoringevents
        ports:
        - containerPort: 8080
          name: server
          protocol: TCP
        resources: {}
        terminationMessagePath: /dev/termination-log
        terminationMessagePolicy: File
      dnsPolicy: ClusterFirst
      restartPolicy: Always
      schedulerName: default-scheduler
      securityContext: {}
      terminationGracePeriodSeconds: 30
status:
```

```
  availableReplicas: 1
  conditions:
  - lastTransitionTime: 2018-08-27T15:46:01Z
    lastUpdateTime: 2018-08-27T15:46:01Z
    message: Deployment has minimum availability.
    reason: MinimumReplicasAvailable
    status: "True"
    type: Available
  - lastTransitionTime: 2018-08-27T15:45:59Z
    lastUpdateTime: 2018-08-27T15:46:01Z
    message: ReplicaSet "ocnef-monitoringevents-7898d657d9" has
successfully progressed.
    reason: NewReplicaSetAvailable
    status: "True"
    type: Progressing
  observedGeneration: 1
  readyReplicas: 1
  replicas: 1
  updatedReplicas: 1
```

- Check if the microservices can access each other via REST interface.
  Run the following command:

  ```
  # kubectl -n <namespace> exec <pod name> -- curl <uri>
  ```

  Example:

  ```
  # kubectl -n nef-svc exec ocnef-fivegcagent-44f4d8f5d5-6q92i -- curl
  http://ocnef-monitoringevents:8080/3gpp-monitoring-event/v1/anyAfID1000/
  subscriptions
  ```

  > **Note**
  >
  > These commands are in their simple form and display the logs only if there is a single nef<registration> and nf<subscription> pod deployed.

Application related tips

Run the following command to check the application logs and look for exceptions:

```
# kubectl -n <namespace> logs -f <pod name>
```

You can use '-f' to follow the logs or 'grep' for specific pattern in the log output.

Example:

```
# kubectl -n nef-svc logs -f $(kubectl -n nef-svc get pods -o name|cut -d'/' -
f2|grep nfr)
# kubectl -n nef-svc logs -f $(kubectl -n nef-svc get pods -o name|cut -d'/' -
f2|grep nfs)
```

> **Note**
>
> These commands are in their simple form and display the logs only if there is 1 nef<registration> and nf<subscription> pod deployed.

4.2 Deployment Related Issue

This section describes the most common deployment related issues and their resolution steps. It is recommended to perform the resolution steps provided in this guide. If the issue still persists, then contact My Oracle Support.

4.2.1 Installation

4.2.1.1 Helm Install Failure

This section describes the various scenarios in which `helm install` might fail. Following are some of the scenarios:

4.2.1.1.1 Incorrect image name in ocnef-custom-values file

Problem

`helm install` might fail if an incorrect image name is provided in the *ocnef-custom-values.yaml* file.

Error Code/Error Message

When `kubectl get pods -n <ocnef_namespace>` is performed, the status of the pods might be ImagePullBackOff or ErrImagePull.

For example:

```
$ kubectl get pods -n ocnef
```

```
nefats-ocats-nef-8bd489d58-jd7ld               1/1     Running   0          47h
ocats-ocats-nef-67cf948f67-k59cn               1/1     Running   2          6d22h
ocnef-config-server-75bd4fc7f8-ttgbx           1/1     Running   0          4h41m
ocnef-expgw-afmgr-67dff6c6fd-tblvq             2/2     Running   0          4h41m
ocnef-expgw-apimgr-5665864dc4-bq9qj            1/1     Running   0          4h41m
ocnef-expgw-apirouter-5dc68f4c69-jdh9q         2/2     Running   0          4h41m
ocnef-expgw-eventmgr-67c5fbdb9c-zg6ll          1/1     Running   0          4h41m
ocnef-ext-egress-gateway-f569449d4-xd7gs       1/1     Running   0          4h41m
ocnef-ext-ingress-gateway-69f989878b-2tvdh     1/1     Running   0          4h41m
ocnef-fivegc-egress-gateway-6f84b8685c-xp292   1/1     Running   0          4h41m
ocnef-fivegc-ingress-gateway-757566b6d5-bjrqm  1/1     Running   0          4h41m
ocnef-fivegcagent-667d87696d-pqfd7             1/1     Running   0          4h41m
ocnef-monitoringevents-87cdb4b67-qfpn2         1/1     Running   0          4h41m
ocnef-nfdb-5ff78cf4d6-qm2mn                    1/1     Running   0          47h
ocnef-ocnef-ccfclient-7fd9c5c4bc-jc9tz         1/1     Running   0          4h41m
ocnef-ocnef-expiry-auditor-6c97cf49f7-r47kb    1/1     Running   0          4h41m
ocnefsim-ocstub-nef-af-74df6f7b4f-rp54q        1/1     Running   0          46h
ocnefsim-ocstub-nef-gmlc-5d456ffddb-2xzx8      1/1     Running   0          46h
ocnefsim-ocstub-nef-nrf-67fbd5bdf6-lqbvp       1/1     Running   0          46h
ocnefsim-ocstub-nef-udm-9d86d96c7-wjhf5        1/1     Running   0          46h
```

Solution

Perform the following steps to verify and correct the image name:

1. Check *ocnef-custom-values.yaml* file has the release specific image name and tags.

   ```
   vi ocnef-custom-values-<release-number>
   ```

   For ocnef images details, see "**Customizing NEF**" in *Oracle Communications Cloud Native Core, Network Exposure Function Installation, Upgrade, and Fault Recovery Guide*.
2. Edit *ocnef-custom-values* file in case the release specific image name and tags must be modified.
3. Save the file.
4. Run the following command to delete the deployment:

   ```
   helm delete --purge <release_namespace>
   ```

   Sample command:

   ```
   helm delete --purge ocnef
   ```

5. To verify the deletion, see the "Verifying Uninstallation" section in *Oracle Communications Cloud Native Core, Network Exposure Function Installation, Upgrade, and Fault Recovery Guide*.
6. Run `helm install` command. For helm install command, see the "Customizing NEF" section in *Oracle Communications Cloud Native Core, Network Exposure Function Installation, Upgrade, and Fault Recovery Guide*.
7. Run `kubectl get pods -n <ocnef_namespace>` to verify if all the pods are in **Running** state.
   For example:

   ```
   $ kubectl get pods -n ocnef
   ```

   ```
   NAME                                              READY   STATUS    RESTARTS   AGE
   nefats-ocats-nef-8bd489d58-jd7ld                  1/1     Running   0          47h
   ocats-ocats-nef-67cf948f67-k59cn                  1/1     Running   2          6d22h
   ocnef-config-server-75bd4fc7f8-ttgbx              1/1     Running   0          4h41m
   ocnef-expgw-afmgr-67dff6c6fd-tblvq                2/2     Running   0          4h41m
   ocnef-expgw-apimgr-5665864dc4-bq9qj               1/1     Running   0          4h41m
   ocnef-expgw-apirouter-5dc68f4c69-jdh9q            2/2     Running   0          4h41m
   ocnef-expgw-eventmgr-67c5fbdb9c-zg6ll             1/1     Running   0          4h41m
   ocnef-ext-egress-gateway-f569449d4-xd7gs          1/1     Running   0          4h41m
   ocnef-ext-ingress-gateway-69f989878b-2tvdh        1/1     Running   0          4h41m
   ocnef-fivegc-egress-gateway-6f84b8685c-xp292      1/1     Running   0          4h41m
   ocnef-fivegc-ingress-gateway-757566b6d5-bjrqm     1/1     Running   0          4h41m
   ocnef-fivegcagent-667d87696d-pqfd7                1/1     Running   0          4h41m
   ocnef-monitoringevents-87cdb4b67-qfpn2            1/1     Running   0          4h41m
   ocnef-nfdb-5ff78cf4d6-qm2mn                       1/1     Running   0          47h
   ocnef-ocnef-ccfclient-7fd9c5c4bc-jc9tz            1/1     Running   0          4h41m
   ocnef-ocnef-expiry-auditor-6c97cf49f7-r47kb       1/1     Running   0          4h41m
   ocnefsim-ocstub-nef-af-74df6f7b4f-rp54q           1/1     Running   0          46h
   ocnefsim-ocstub-nef-gmlc-5d456ffddb-2xzx8         1/1     Running   0          46h
   ocnefsim-ocstub-nef-nrf-67fbd5bdf6-lqbvp          1/1     Running   0          46h
   ocnefsim-ocstub-nef-udm-9d86d96c7-wjhf5           1/1     Running   0          46h
   ```

4.2.1.1.2 Docker registry is configured incorrectly

Problem

`helm install` might fail if the docker registry is not configured in all primary and secondary nodes.

Error Code/Error Message

When `kubectl get pods -n <ocnef_namespace>` is performed, the status of the pods might be ImagePullBackOff or ErrImagePull.

For example:

```
$ kubectl get pods -n ocnef
```

Solution

Configure docker registry on all primary and secondary nodes. For more information on configuring the docker registry, see *Oracle Communications Cloud Native Core, Cloud Native Environment Installation, Upgrade, and Fault Recovery Guide*.

4.2.1.1.3 Continuous Restart of Pods

Problem

`helm install` might fail if the MySQL primary and secondary hosts are not configured properly in `ocnef-custom-values.yaml`.

Error Code/Error Message

When `kubectl get pods -n <ocnef_namespace>` is performed, the pods restart count increases continuously.

For example:

```
$ kubectl get pods -n ocnef
```

Solution

MySQL servers(s) may not be configured properly according to the pre-installation steps. For configuring MySQL servers, see the "Configuring Database, Creating Users, and Granting Permissions" section in *Oracle Communications Cloud Native Core, Network Exposure Function Installation, Upgrade, and Fault Recovery Guide*.

4.2.1.2 Pod Creation Failure

A pod creation can fail due to various reasons. Some of the possible scenarios are as follows:

Verifying Pod Image Correctness

To verify pod image:

- Check whether any of the pods is in the **ImagePullBackOff** state.
- Check if the image name used for all the pods are correct. Verify the image names and versions from the values in the NEF custom-values.yaml file. For more information about the custom value file, see *Oracle Communications Cloud Native Core, Network Exposure Function Installation, Upgrade, and Fault Recovery Guide*.
- After updating the custom-values.yaml file, run the following command for helm upgrade:
  `helm upgrade <helm chart> [--version <OCNEF version>] --name <release> --namespace <ocnefnamespace> -f <ocnef_values.yaml>`

Verifying Resource Allocation Failure

To verify any resource allocation failure:

- Run the following command to verify whether any pod is in the pending state.

```
kubectl describe <nef-drservice pod id> --n <ocnef-namespace>
```

- Verify whether any warning on insufficient CPU exists in the describe output of the respective pod. If it exists, it means there are insufficient CPUs for the pods to start. Address this hardware issue.
- Run the following helm upgrade command after updating the values.yaml file.

Verifying Resource Allocation Issues on Webscale Environment

Webscale environment has openshift container installed. There can be cases where,

- Pods does not scale after you run the installation command and the helm install command fails with timeout error. In this case, check for preinstall hooks failure. Run the **oc get job** command to create the jobs. Describe the job for which the pods are not getting scaled and check if there are **quota limit exceeded** errors with CPU or memory.
- Any of the actual microservice pods do not scale post the hooks completion. In this case, run the **oc get rs** command to get the list of replicaset created for the NF deployment. Then, describe the replicaset for which the pods are not getting scaled and check for **resource quota limit exceeded** errors with CPU or memory.
- Helm install command times-out after all the microservice pods are scaled as expected with the expected number of replicas. In this case, check for post install hooks failure. Run the **oc get job** command to get the post install jobs and do a describe on the job for which the pods are not getting scaled and check if there are **quota limit exceeded** errors with CPU or memory.
- Resource quota exceed beyond limits.

4.2.1.3 Pod Startup Failure

Follow the guidelines shared below to debug the pod startup failure liveness check issues:

- If dr-service, diameter-proxy, and diam-gateway services are stuck in the Init state, then the reason could be that config-server is not yet up. A sample log on these services is as follows:

  ```
  "Config Server is Not yet Up, Wait For config server to be up."
  ```

  To resolve this, you must either check for the reason of config-server not being up or if the config-server is not required, then disable it.
- If the notify and on-demand migration service is stuck in the Init state, then the reason could be the dr-service is not yet up. A sample log on these services is as follows:

  ```
  "DR Service is Not yet Up, Wait For dr service to be up."
  ```

  To resolve this, check for failures on dr-service.

4.2.1.4 NRF Registration Failure

The NEF registration with NRF may fail due to various reasons. Some of the possible scenarios are as follows:

- Confirm whether registration was successful from the nrf-client-service pod.
- Check the ocnef-nrf-client-nfmanagement logs. If the log has **"OCNEF is Deregistration"** then:

  - Check if all the services mentioned under allorudr/slf (depending on NEF mode) in the custom-values.yaml file has same spelling as that of service name and are enabled.
  - Once all services are up, NEF must register with NRF.
- If you see a log for **SERVICE_UNAVAILABLE(503)**, check if the primary and secondary NRF configurations (primaryNrfApiRoot/secondaryNrfApiRoot) are correct and they are UP and Running.

4.2.1.5 Custom Value File Parse Failure

This section explains troubleshooting procedure in case of failure while parsing `ocnef-custom-values.yaml` file.

Problem

Not able to parse *ocnef-custom-values-x.x.x.yaml*, while running `helm install`.

Error Code/Error Message

Error: failed to parse ocnef-custom-values-x.x.x.yaml: error converting YAML to JSON: yaml

Symptom

While creating the *ocnef-custom-values-x.x.x.yaml* file, if the aforementioned error is received, it means that the file is not created properly. The tree structure may not have been followed or there may also be tab spaces in the file.

Solution

Following the procedure as mentioned:

1. Download the latest NRF templates zip file from MOS. For more information, see the "Downloading NEF Package" section in *Oracle Communications Cloud Native Core, Network Exposure Function Installation, Upgrade, and Fault Recovery Guide*.
2. Follow the steps mentioned in the "Installation Tasks" section in *Oracle Communications Cloud Native Core, Network Exposure Function Installation, Upgrade, and Fault Recovery Guide*.

4.2.2 Post Installation

4.2.2.1 Helm Test Error Scenario

Following are the error scenarios that may be identified using helm test.

1. Run the following command to get the Helm Test pod name:

   ```
   kubectl get pods -n <deployment-namespace>
   ```

2. When a helm test is performed, a new helm test pod is created. Check for the Helm Test pod that is in an error state.
3. Get the logs using the following command:

   ```
   kubectl logs <podname> -n <namespace>
   ```

Example:

```
kubectl get <helm_test_pod> -n ocnef
```

For further assistance, collect the logs and contact MOS.

4.3 Database Related Issues

This section describes the most common database related issues and their resolution steps. It is recommended to perform the resolution steps provided in this guide. If the issue still persists, then contact My Oracle Support.

4.3.1 MySQL DB Access Failure

Problem

Keyword - wait-for-db

Tags - "config-server" "database" "readiness" "init" "SQLException" "access denied"

Because of the database accessibility issues from the NEF service, pods stay in the init state.

For some pods, if they come up, they will be kept on getting the exception : " Cannot connect to database server java.sql.SQLException"

Reasons:

1. MySQL host IP address OR MySQL-service name[in case of occne-infra] is not correctly given.
2. Few MySQL nodes are probably down.
3. Username/Password given in the secrets are not created in the database OR not having proper grant/access to service databases.
4. **MOST LIKELY** - Databases are not created correctly with the same name mentioned in the NEF-custom-value file while installing NEF.

Resolution Steps

To resolve this issue, perform the following steps:

1. Check if the database IP is proper and pingable from worker nodes of the Kubernetes cluster. Update the database IP and service accordingly. If required, you can use floating IP as well. If the database connectivity issue is there, then please update the proper IP address.
   In the case of the OCCNE-infra, Instead of mentioning IP address for MySQL connection, please use FQDN for mysql-connectivity-service to connect to the database.
2. Manually log in to MySQL via the same database IP mentioned in a custom-value file. In case of MySQL service name, describe the service by command :

   ```
   kubectl describe svc <mysql-servicename> -n <namespace>
   ```

   and login to the MySQL database with all sets of IPs described in the MySQL service, If any SQL node is down, it will lead to an intermittent DB query failure issue. So make sure that you can log in to MySQL from all the Nodes mentioned in the IP list of MySQL-service describe command.

Make sure that all the MySQL nodes are up and running before installing NEF.

3. Check the existing user list into the database using SQL query: `select user from mysql.user;`
   Check if all the mentioned users in the custom-value of NEF installation are present in the database.

   > **Note**
   >
   > Create the user with proper password as mentioned in the secret file of the NEF.

4. Check the grants of all the users mentioned into the custom-value file by SQL query: "show grants for <username>;"
   If username/password issue is there, then please correctly create the user with the required password and provide grants as per the *Oracle Communications Cloud Native Core, Network Exposure Function Installation, Upgrade, and Fault Recovery Guide*.
5. Check the databases are created with the same name mentioned in the custom-value file for the services.

   > **Note**
   >
   > Create the database as per the custom-value file.

6. Check if problematic pods are getting created on any one unique worker node. If yes, then may be the cause of the error can be the worker node. Try draining the problematic worker node and allow pods to move to another node.

4.4 Service Related Issues

This section describes the most common service related issues and their resolution steps. It is recommended to perform the resolution steps provided in this guide. If the issue still persists, then contact My Oracle Support.

4.4.1 Errors from Egress Gateway

If the traffic is not routed through Egress Gateway, then check the following:

- Check whether the Egress gateway parameters are configured correctly through the NEF custom values.
- Check whether Egress pod is running from Kubectl. To check, run the following command: `kubectl get pods -n <Release.name>`
- To enable the outgoing traffic using HTTPS, set the **enableOutgoingHttps** parameter as **true**.

4.4.2 Debugging Errors from Ingress Gateway

The possible errors that you may encounter from Ingress Gateway are:

- **Check for 500 Error:** If the request fails with 500 status code without Problem Details information, it means that the flow ended in `ocnef-ingressgateway` pod without route. You can confirm the same in the errors or exception section of the `ocnef-ingressgateway` pod logs. Y

- **Check for 503 Error:** If the request fails with 503 status code with "SERVICE_UNAVAILABLE" in Problem Details, then it means that the `ocnef-expgw-apirouter` pod is not reachable due to some reason.

4.5 Upgrade or Rollback Failure

When Oracle Communications Network Exposure Function (NEF) upgrade or rollback fails, perform the following procedure.

1. Check the pre or post upgrade logs or rollback hook logs in Kibana as applicable. You can filter upgrade or rollback logs using the following filters:
   - For upgrade: hookName = "pre-upgrade" or hookName = "post-upgrade"
   - For rollback: hookName = "pre-rollback" or hookName = "post-rollback"

```
{
   "instant":{
      "epochSecond":1669292396,
      "nanoOfSecond":918939800
   },
   "thread":"main",
   "level":"INFO",
   "loggerName":"com.oracle.utils.SqlUtils",
   "message":"Executing the SQL query: ALTER TABLE
`ocnef11`.`ocnef_me_subscription` \nADD CONSTRAINT
`me_subscription_id_to_owner_site_id`\n FOREIGN KEY (`owner_site_id`)\n
REFERENCES `ocnef_site_instance_model` (`site_instance_ref_id`)\n ON
DELETE NO ACTION\n ON UPDATE NO ACTION;\n",
   "endOfBatch":false,
   "loggerFqcn":"org.apache.logging.slf4j.Log4jLogger",
   "threadId":1,
   "threadPriority":5,
   "messageTimestamp":"2022-11-24T17:49:56.918+0530",
   "hookName":"pre-upgrade"
}
```

2. Check the pod logs in Kibana to analyze the cause of failure.
3. After detecting the cause of failure, do the following:
   - For upgrade failure:
     - If the cause of upgrade failure is database or network connectivity issue, contact your system administrator. When the issue is resolved, rerun the upgrade command.
     - If the cause of failure occurs during the preupgrade phase, do not perform the roll back.
     - If the upgrade failure occurs during the postupgrade phase, for example, post upgrade hook failure due to target release pod not moving to ready state, then perform a rollback.
   - For rollback failure: If the cause of rollback failure is database or network connectivity issue, contact your system administrator. When the issue is resolved, rerun the rollback command.
   - If the issue persists, contact My Oracle Support.

5
NEF Alerts

This chapter includes information about the following NEF alerts:

- System Level Alerts
- Application Level Alerts

Note

- The performance and capacity of the NEF system may vary based on the call model, feature or interface configuration, and underlying CNE and hardware environment.
- Due to unavailability of metric and/or MQL queries, the following alerts are not supported for OCI:
  - OcnefNfStatusUnavailable
  - OcnefPodsRestart
  - OcnefIngressGatewayServiceDown
  - OcnefApiRouterServiceDown
  - OcnefFiveGcAgentServiceDown
  - OcnefMonitoringEventServiceDown
  - OcnefCCFClientServiceDown
  - OcnefExpiryAuditorServiceDown
  - OcnefQOSServiceDown
  - OcnefTIServiceDown
  - OcnefDTServiceDown
  - OcnefEgressGatewayServiceDown
  - OcnefMemoryUsageCrossedMinorThreshold
  - OcnefMemoryUsageCrossedMajorThreshold
  - FiveGcInvalidConfiguration
  - OcnefAllSiteStatus
  - OcnefDBReplicationStatus

5.1 System Level Alerts

This section lists the system level alerts for NEF.

5.1.1 OcnefNfStatusUnavailable

Table 5-1 OcnefNfStatusUnavailable

Field	Details
Description	'NEF services unavailable'
Summary	"kubernetes_namespace: {{$labels.kubernetes_namespace}}, timestamp: {{ with query \"time()\" }}{{ . | first | value | humanizeTimestamp }}{{ end }} : All NEF services are unavailable."
Severity	Critical
Condition	All the NEF services are unavailable, either because the NEF is getting deployed or purged.
Metric Used	'up' **Note**: This is a Prometheus metric used for instance availability monitoring. If this metric is not available, use a similar metric as exposed by the monitoring system.

Recommended Actions	The alert is cleared automatically when the NEF services restart.
	Steps:
	1. Check for service-specific alerts which may be causing the issues with service exposure.
	2. Run the following command to check the pod status:
	`$ kubectl get po -n <namespace>`
	a. Run the following command to analyze the error condition of the pod that is not in the *Running* state:
	`$ kubectl describe pod <pod name not in Running state> -n <namespace>`
	Where <pod name not in Running state> indicates the pod that is not in the *Running* state.
	3. Refer to the application logs on Kibana and check for database related failures such as connectivity and invalid secrets. The logs can be filtered based on the services.
	4. Check for helm status to make sure there are no errors:
	`$ helm status <helm release name of the desired NF> -n <namespace>`
	If it is not in “STATUS : DEPLOYED”, then capture logs and event again.
	5. In case the issue persists, capture all the outputs for the above steps and contact My Oracle Support.
	Note: Use CNC NF Data Collector tool for capturing logs. For more information on the Data Collector tool, see *Oracle Communications Cloud Native Core, Network Function Data Collector User Guide*.

5.1.2 OcnefPodsRestart

Table 5-2 OcnefPodsRestart

Field	Details
Description	'Pod <*Pod Name*> has restarted.
Summary	"kubernetes_namespace: {{$labels.namespace}}, podname: {{$labels.pod}}, timestamp: {{ with query \"time()\" }}{{ . | first | value | humanizeTimestamp }}{{ end }} : A Pod has restarted"
Severity	Major
Condition	A pod belonging to any of the NEF services has restarted.
Metric Used	kube_pod_container_status_restarts_total

Recommended Actions	The alert is cleared automatically if the specific pod is up.
	Steps:
	1. Refer to the application logs on Kibana and filter based on pod name, check for database related failures such as connectivity and Kubernetes secrets.
	2. To check the orchestration logs for liveness or readiness probe failures, do the following:
	a. Run the following command to check the pod status:
	`$ kubectl get po -n <namespace>`
	b. Run the following command to analyze the error condition of the pod that is not in the *Running* state:
	`$ kubectl describe pod <pod name not in Running state> -n <namespace>`
	Where <pod name not in Running state> indicates the pod that is not in the *Running* state.
	3. Check the database status. For more information, see *Oracle Communications Cloud Native Core, cnDBTier User Guide*.
	4. In case the issue persists, capture all the outputs for the above steps and contact My Oracle Support.
	Note: Use CNC NF Data Collector tool for capturing logs. For more information on the Data Collector tool, see *Oracle Communications Cloud Native Core, Network Function Data Collector User Guide*.

5.1.3 OcnefTotalExternalIngressTrafficRateAboveMinorThreshold

Table 5-3 OcnefTotalExternalIngressTrafficRateAboveMinorThreshold

Field	Details
Description	OCNEF External Ingress traffic rate is above the configured minor threshold i.e. 800 TPS (current value is: {{ $value }})

Table 5-3 (Cont.) OcnefTotalExternalIngressTrafficRateAboveMinorThreshold

Field	Details
Summary	"timestamp: {{ with query "time()" }}{{ . | first | value | humanizeTimestamp }}{{ end }}: Traffic rate is above 80 percent of max TPS (1000)"
Severity	Minor
Condition	The total NEF External Ingress traffic rate has crossed the configured minor threshold of 800 TPS. Default value of this alert trigger point in NefAlertrules alert file is 80 % of 1000 (maximum ingress request rate).
OID	1.3.6.1.4.1.323.5.3.39.1.2.7003
Metric Used	oc_ingressgateway_http_requests_total
Recommended Actions	The alert is cleared either when the total External Ingress traffic rate falls below the minor threshold or when the total traffic rate crosses the major threshold, in which case the OcnefTotalExternalIngressTrafficRateAboveMajorThreshold alert is raised. **Note**: The threshold is configurable in the NefAlertrules alert file. Reassess why the NEF is receiving additional traffic. If this alert is unexpected, contact My Oracle Support. **Steps**: 1. Refer Grafana to determine which service is receiving high traffic. 2. Refer **Ingress gateway** section in Grafana to determine the increase in 4xx and 5xx error codes. 3. Check **Ingress gateway** logs on Kibana to determine the reason for the errors.

5.1.4 OcnefTotalFivegcIngressTrafficRateAboveMinorThreshold

Table 5-4 OcnefTotalFivegcIngressTrafficRateAboveMinorThreshold

Field	Details
Description	OCNEF Fivegc Ingress traffic rate is above the configured minor threshold i.e. 800 TPS (current value is: {{ $value }})
Summary	"timestamp: {{ with query "time()" }}{{ . | first | value | humanizeTimestamp }}{{ end }}: Traffic rate is above 80 percent of max TPS (1000)"
Severity	Minor
Condition	The total NEF Fivegc Ingress traffic rate has crossed the configured minor threshold of 800 TPS. Default value of this alert trigger point in NefAlertrules alert file is 80 % of 1000 (maximum ingress request rate).
OID	1.3.6.1.4.1.323.5.3.39.1.2.7004
Metric Used	oc_ingressgateway_http_requests_total

Table 5-4 (Cont.) OcnefTotalFivegcIngressTrafficRateAboveMinorThreshold

Field	Details
Recommended Actions	The alert is cleared either when the total Fivegc Ingress traffic rate falls below the minor threshold or when the total traffic rate crosses the major threshold, in which case the OcnefTotalFivegcIngressTrafficRateAboveMajorThreshold alert is raised. **Note**: The threshold is configurable in the NefAlertrules alert file. Reassess why the NEF is receiving additional traffic. If this alert is unexpected, contact My Oracle Support. **Steps**: 1. Refer Grafana to determine which service is receiving high traffic. 2. Refer **Ingress gateway** section in Grafana to determine the increase in 4xx and 5xx error codes. 3. Check **Ingress gateway** logs on Kibana to determine the reason for the errors.

5.1.5 OcnefTotalExternalIngressTrafficRateAboveMajorThreshold

Table 5-5 OcnefTotalExternalIngressTrafficRateAboveMajorThreshold

Field	Details
Description	OCNEF External Ingress traffic rate is above the configured major threshold i.e. 900 TPS (current value is: {{ $value }})
Summary	"timestamp: {{ with query "time()" }}{{ . | first | value | humanizeTimestamp }}{{ end }}: Traffic rate is above 90 percent of max TPS (1000)"
Severity	Major
Condition	The total NEF External Ingress traffic rate has crossed the configured major threshold of 900 TPS. Default value of this alert trigger point in NefAlertrules alert file is 90 % of 1000 (maximum ingress request rate).
OID	1.3.6.1.4.1.323.5.3.39.1.2.7005
Metric Used	oc_ingressgateway_http_requests_total
Recommended Actions	The alert is cleared either when the total External Ingress traffic rate falls below the major threshold or when the total traffic rate crosses the critical threshold, in which case the OcnefTotalExternalIngressTrafficRateAboveCriticalThreshold alert is raised. **Note**: The threshold is configurable in the NefAlertrules alert file. Reassess why the NEF is receiving additional traffic. If this alert is unexpected, contact My Oracle Support. **Steps**: 1. Refer Grafana to determine which service is receiving high traffic. 2. Refer **Ingress gateway** section in Grafana to determine the increase in 4xx and 5xx error codes. 3. Check **Ingress gateway** logs on Kibana to determine the reason for the errors.

5.1.6 OcnefTotalFivegcIngressTrafficRateAboveMajorThreshold

Table 5-6 OcnefTotalFivegcIngressTrafficRateAboveMajorThreshold

Field	Details
Description	OCNEF Fivegc Ingress traffic rate is above the configured major threshold i.e. 900 TPS (current value is: {{ $value }})
Summary	"timestamp: {{ with query "time()" }}{{ . | first | value | humanizeTimestamp }}{{ end }}: Traffic rate is above 90 percent of max TPS (1000)"
Severity	Major
Condition	The total NEF Fivegc Ingress traffic rate has crossed the configured major threshold of 900 TPS. Default value of this alert trigger point in NefAlertrules alert file is 90 % of 1000 (maximum ingress request rate).
OID	1.3.6.1.4.1.323.5.3.39.1.2.7006
Metric Used	oc_ingressgateway_http_requests_total
Recommended Actions	The alert is cleared either when the total Fivegc Ingress traffic rate falls below the major threshold or when the total traffic rate crosses the critical threshold, in which case the OcnefTotalFivegcIngressTrafficRateAboveCriticalThreshold alert is raised. **Note**: The threshold is configurable in the NefAlertrules alert file. Reassess why the NEF is receiving additional traffic. If this alert is unexpected, contact My Oracle Support. **Steps**: 1. Refer Grafana to determine which service is receiving high traffic. 2. Refer **Ingress gateway** section in Grafana to determine the increase in 4xx and 5xx error codes. 3. Check **Ingress gateway** logs on Kibana to determine the reason for the errors.

5.1.7 OcnefTotalExternalIngressTrafficRateAboveCriticalThreshold

Table 5-7 OcnefTotalExternalIngressTrafficRateAboveCriticalThreshold

Field	Details
Description	OCNEF External Ingress traffic rate is above the configured critical threshold i.e. 950 TPS (current value is: {{ $value }})
Summary	"timestamp: {{ with query "time()" }}{{ . | first | value | humanizeTimestamp }}{{ end }}: Traffic Rate is above 95 percent of max TPS (1000)"
Severity	Critical
Condition	The total NEF External Ingress traffic rate has crossed the configured critical threshold of 950 TPS. Default value of this alert trigger point in NefAlertrules alert file is 95 % of 1000 (maximum ingress request rate).
OID	1.3.6.1.4.1.323.5.3.39.1.2.7007
Metric Used	oc_ingressgateway_http_requests_total

Table 5-7 (Cont.) OcnefTotalExternalIngressTrafficRateAboveCriticalThreshold

Field	Details
Recommended Actions	The alert is cleared either when the total External Ingress traffic rate falls below the critical threshold. **Note**: The threshold is configurable in the NefAlertrules alert file. Reassess why the NEF is receiving additional traffic. If this alert is unexpected, contact My Oracle Support. **Steps**: 1. Refer Grafana to determine which service is receiving high traffic. 2. Refer **Ingress gateway** section in Grafana to determine the increase in 4xx and 5xx error codes. 3. Check **Ingress gateway** logs on Kibana to determine the reason for the errors.

5.1.8 OcnefTotalFivegcIngressTrafficRateAboveCriticalThreshold

Table 5-8 OcnefTotalFivegcIngressTrafficRateAboveCriticalThreshold

Field	Details
Description	OCNEF Fivegc Ingress traffic rate is above the configured critical threshold i.e. 950 TPS (current value is: {{ $value }})
Summary	"timestamp: {{ with query "time()" }}{{ . | first | value | humanizeTimestamp }}{{ end }}: Traffic Rate is above 95 percent of max TPS (1000)"
Severity	Critical
Condition	The total NEF Fivegc Ingress traffic rate has crossed the configured critical threshold of 950 TPS. Default value of this alert trigger point in NefAlertrules alert file is 95 % of 1000 (maximum ingress request rate).
OID	1.3.6.1.4.1.323.5.3.39.1.2.7008
Metric Used	oc_ingressgateway_http_requests_total
Recommended Actions	The alert is cleared either when the total Fivegc Ingress traffic rate falls below the critical threshold. **Note**: The threshold is configurable in the NefAlertrules alert file. Reassess why the NEF is receiving additional traffic. If this alert is unexpected, contact My Oracle Support. **Steps**: 1. Refer Grafana to determine which service is receiving high traffic. 2. Refer **Ingress gateway** section in Grafana to determine the increase in 4xx and 5xx error codes. 3. Check **Ingress gateway** logs on Kibana to determine the reason for the errors.

5.1.9 OcnefExternalIngressTransactionErrorRateAboveZeroPointOnePercent

Table 5-9 OcnefExternalIngressTransactionErrorRateAboveZeroPointOnePercent

Field	Details
Description	External Ingress transaction Error rate is above 0.1 percent(current value is {{ $value }})
Summary	"timestamp: {{ with query "time()" }}{{ . | first | value | humanizeTimestamp }}{{ end }}: Transaction Error rate detected above 0.1 percent of total transactions"
Severity	Warning
Condition	The number of failed external ingress transactions is above 0.1 percent of the total transactions.
OID	1.3.6.1.4.1.323.5.3.39.1.2.7009
Metric Used	oc_ingressgateway_http_responses_total
Recommended Actions	The alert is cleared when the number of failure external ingress transactions is below 0.1 percent of the total transactions or when the number of failed transactions crosses the 1% threshold, in which case the OcnefExternalIngressTransactionErrorRateAbove1Percent is raised. **Steps**: 1. Check the service specific metrics to understand the specific service request errors. 2. Check metrics per service, per method: 3. If guidance is required, contact My Oracle Support.

5.1.10 OcnefFivegcIngressTransactionErrorRateAboveZeroPointOnePercent

Table 5-10 OcnefFivegcIngressTransactionErrorRateAboveZeroPointOnePercent

Field	Details
Description	Fivegc Ingress transaction error rate is above 0.1 percent of total transactions (current value is {{ $value }})
Summary	"timestamp: {{ with query "time()" }}{{ . | first | value | humanizeTimestamp }}{{ end }}: Transaction error rate detected above 0.1 percent of total transactions"
Severity	Warning
Condition	The number of failed Fivegc ingress transactions is above 0.1 percent of the total transactions.
OID	1.3.6.1.4.1.323.5.3.39.1.2.7010
Metric Used	oc_ingressgateway_http_responses_total

Table 5-10 (Cont.) OcnefFivegcIngressTransactionErrorRateAboveZeroPointOnePercent

Field	Details
Recommended Actions	The alert is cleared when the number of failure Fivegc ingress transactions is below 0.1 percent of the total transactions or when the number of failed transactions crosses the 1% threshold, in which case the OcnefFivegcIngressTransactionErrorRateAbove1Percent is raised. **Steps**: 1. Check the service specific metrics to understand the specific service request errors. 2. Check metrics per service, per method: 3. If guidance is required, contact My Oracle Support.

5.1.11 OcnefExternalIngressTransactionErrorRateAbove1Percent

Table 5-11 OcnefExternalIngressTransactionErrorRateAbove1Percent

Field	Details
Description	External Ingress transaction error rate is above 1 percent of total transactions (current value is {{ $value }})
Summary	"timestamp: {{ with query "time()" }}{{ . | first | value | humanizeTimestamp }}{{ end }}: Transaction error rate detected above 1 percent of total transactions"
Severity	Warning
Condition	The number of failed External Ingress transactions is above 1 percent of the total transactions.
OID	1.3.6.1.4.1.323.5.3.39.1.2.7011
Metric Used	oc_ingressgateway_http_responses_total
Recommended Actions	The alert is cleared when the number of failure External Ingress transactions is below 1 percent of the total transactions or when the number of failed transactions crosses the 10% threshold, in which case the OcnefExternalIngressTransactionErrorRateAbove10Percent is raised. **Steps**: 1. Check the service specific metrics to understand the specific service request errors. 2. Check metrics per service, per method: 3. If guidance is required, contact My Oracle Support.

5.1.12 OcnefFivegcIngressTransactionErrorRateAbove1Percent

Table 5-12 OcnefFivegcIngressTransactionErrorRateAbove1Percent

Field	Details
Description	Fivegc Ingress transaction error rate is above 1 percent of total Fivegc Ingress transactions (current value is {{ $value }})

Table 5-12 (Cont.) OcnefFivegcIngressTransactionErrorRateAbove1Percent

Field	Details
Summary	"timestamp: {{ with query "time()" }}{{ . | first | value | humanizeTimestamp }}{{ end }}: Transaction error rate detected above 1 percent of total transactions"
Severity	Warning
Condition	The number of failed Fivegc Ingress transactions is above 1 percent of the total transactions.
OID	1.3.6.1.4.1.323.5.3.39.1.2.7012
Metric Used	oc_ingressgateway_http_responses_total
Recommended Actions	The alert is cleared when the number of failure Fivegc Ingress transactions is below 1 percent of the total transactions or when the number of failed transactions crosses the 10% threshold, in which case the OcnefFivegcIngressTransactionErrorRateAbove10Percent is raised. **Steps**: 1. Check the service specific metrics to understand the specific service request errors. 2. Check metrics per service, per method: 3. If guidance is required, contact My Oracle Support.

5.1.13 OcnefExternalIngressTransactionErrorRateAbove10Percent

Table 5-13 OcnefExternalIngressTransactionErrorRateAbove10Percent

Field	Details
Description	External Ingress transaction error rate is above 10 percent of total External Ingress transactions (current value is {{ $value }})
Summary	"timestamp: {{ with query "time()" }}{{ . | first | value | humanizeTimestamp }}{{ end }}: Transaction error rate detected above 10 percent of total transactions"
Severity	Minor
Condition	The number of failed External Ingress transactions is above 10 percent of the total transactions.
OID	1.3.6.1.4.1.323.5.3.39.1.2.7013
Metric Used	oc_ingressgateway_http_responses_total
Recommended Actions	The alert is cleared when the number of failure External Ingress transactions is below 10 percent of the total transactions or when the number of failed transactions crosses the 25% threshold, in which case the OcnefExternalIngressTransactionErrorRateAbove25Percent is raised. **Steps**: 1. Check the service specific metrics to understand the specific service request errors. 2. Check metrics per service, per method: 3. If guidance is required, contact My Oracle Support.

5.1.14 OcnefFivegcIngressTransactionErrorRateAbove10Percent

Table 5-14 OcnefFivegcIngressTransactionErrorRateAbove10Percent

Field	Details
Description	Fivegc Ingress transaction error rate is above 10 percent of total Fivegc Ingress transactions (current value is {{ $value }})
Summary	"timestamp: {{ with query "time()" }}{{ . | first | value | humanizeTimestamp }}{{ end }}: Transaction error rate detected above 10 percent of total transactions"
Severity	Minor
Condition	The number of failed Fivegc Ingress transactions is above 10 percent of the total transactions.
OID	1.3.6.1.4.1.323.5.3.39.1.2.7014
Metric Used	oc_ingressgateway_http_responses_total
Recommended Actions	The alert is cleared when the number of failure Fivegc Ingress transactions is below 10 percent of the total transactions or when the number of failed transactions crosses the 25% threshold, in which case the OcnefFivegcIngressTransactionErrorRateAbove25Percent is raised. **Steps**: 1. Check the service specific metrics to understand the specific service request errors. 2. Check metrics per service, per method: 3. If guidance is required, contact My Oracle Support.

5.1.15 OcnefExternalIngressTransactionErrorRateAbove25Percent

Table 5-15 OcnefExternalIngressTransactionErrorRateAbove25Percent

Field	Details
Description	External Ingress transaction error rate detected above 25 percent of total External Ingress transactions (current value is {{ $value }})
Summary	"timestamp: {{ with query "time()" }}{{ . | first | value | humanizeTimestamp }}{{ end }}: Transaction error rate detected above 25 percent of total transactions"
Severity	Major
Condition	The number of failed External Ingress transactions is above 25 percent of the total transactions.
OID	1.3.6.1.4.1.323.5.3.39.1.2.7015
Metric Used	oc_ingressgateway_http_responses_total

Table 5-15 (Cont.) OcnefExternalIngressTransactionErrorRateAbove25Percent

Field	Details
Recommended Actions	The alert is cleared when the number of failure External Ingress transactions is below 25 percent of the total transactions or when the number of failed transactions crosses the 50% threshold, in which case the OcnefExternalIngressTransactionErrorRateAbove50Percent is raised. **Steps**: 1. Check the service specific metrics to understand the specific service request errors. 2. Check metrics per service, per method: 3. If guidance is required, contact My Oracle Support.

5.1.16 OcnefFivegcIngressTransactionErrorRateAbove25Percent

Table 5-16 OcnefFivegcIngressTransactionErrorRateAbove25Percent

Field	Details
Description	Fivegc Ingress transaction error rate detected above 25 percent of total Fivegc Ingress transactions (current value is {{ $value }})
Summary	"timestamp: {{ with query "time()" }}{{ . | first | value | humanizeTimestamp }}{{ end }}: Transaction error rate detected above 25 percent of total transactions"
Severity	Major
Condition	The number of failed Fivegc Ingress transactions is above 25 percent of the total transactions.
OID	1.3.6.1.4.1.323.5.3.39.1.2.7016
Metric Used	oc_ingressgateway_http_responses_total
Recommended Actions	The alert is cleared when the number of failure Fivegc Ingress transactions is below 25 percent of the total transactions or when the number of failed transactions crosses the 50% threshold, in which case the OcnefFivegcIngressTransactionErrorRateAbove50Percent is raised. **Steps**: 1. Check the service specific metrics to understand the specific service request errors. 2. Check metrics per service, per method: 3. If guidance is required, contact My Oracle Support.

5.1.17 OcnefExternalIngressTransactionErrorRateAbove50Percent

Table 5-17 OcnefExternalIngressTransactionErrorRateAbove50Percent

Field	Details
Description	External Ingress transaction error rate detected above 50 percent of total External Ingress transactions (current value is {{ $value }})

Table 5-17 (Cont.) OcnefExternalIngressTransactionErrorRateAbove50Percent

Field	Details
Summary	"timestamp: {{ with query "time()" }}{{ . | first | value | humanizeTimestamp }}{{ end }}: Transaction error rate detected above 50 percent of total transactions"
Severity	Critical
Condition	The number of failed External Ingress transactions is above 50 percent of the total transactions.
OID	1.3.6.1.4.1.323.5.3.39.1.2.7017
Metric Used	oc_ingressgateway_http_responses_total
Recommended Actions	The alert is cleared when the number of failure External Ingress transactions is below 50 percent of the total transactions. **Steps**: 1. Check the service specific metrics to understand the specific service request errors. 2. Check metrics per service, per method: 3. If guidance is required, contact My Oracle Support.

5.1.18 OcnefFivegcIngressTransactionErrorRateAbove50Percent

Table 5-18 OcnefFivegcIngressTransactionErrorRateAbove50Percent

Field	Details
Description	Fivegc Ingress transaction error rate detected above 50 percent of total Fivegc Ingress transactions (current value is {{ $value }})
Summary	"timestamp: {{ with query "time()" }}{{ . | first | value | humanizeTimestamp }}{{ end }}: Transaction error rate detected above 50 percent of total transactions"
Severity	Critical
Condition	The number of failed Fivegc Ingress transactions is above 50 percent of the total transactions.
OID	1.3.6.1.4.1.323.5.3.39.1.2.7018
Metric Used	oc_ingressgateway_http_responses_total
Recommended Actions	The alert is cleared when the number of failure Fivegc Ingress transactions is below 50 percent of the total transactions. **Steps**: 1. Check the service specific metrics to understand the specific service request errors. 2. Check metrics per service, per method: 3. If guidance is required, contact My Oracle Support.

5.1.19 OcnefEgressGatewayServiceDown

Table 5-19 OcnefEgressGatewayServiceDown

Field	Details
Description	"NEF Egress-Gateway service {{$labels.app_kubernetes_io_name}} is down"
Summary	"kubernetes_namespace: {{$labels.kubernetes_namespace}}, podname: {{$labels.kubernetes_pod_name}}, timestamp: {{ with query \"time()\" }}{{ . | first | value | humanizeTimestamp }}{{ end }} : Egress-Gateway service down"
Severity	Critical
Condition	None of the pods of the Egress Gateway microservice is available.
Metric Used	'up' **Note**: This is a Prometheus metric used for instance availability monitoring. If this metric is not available, use a similar metric as exposed by the monitoring system.

Recommended Actions	The alert is cleared when the Egress Gateway service is available.
	Note: The threshold is configurable in the NefAlertrules alert file.
	Steps:
	1. To check the orchestration logs of Egress Gateway service and check for liveness or readiness probe failures, do the following:
	a. Run the following command to check the pod status:
	`$ kubectl get po -n <namespace>`
	b. Run the following command to analyze the error condition of the pod that is not in the *Running* state:
	`$ kubectl describe pod <pod name not in Running state> -n <namespace>`
	Where <pod name not in Running state> indicates the pod that is not in the *Running* state.
	2. Refer to the application logs on Kibana and filter based on Egress Gateway service names. Check for ERROR WARNING logs related to thread exceptions.
	3. Depending on the failure reason, take the resolution steps.
	4. In case the issue persists, capture all the outputs for the above steps and contact My Oracle Support.
	Note: Use CNC NF Data Collector tool for capturing logs. For more information on how to collect logs using Data Collector tool, see *Oracle Communications Cloud Native Core, Network Function Data Collector User Guide*.

5.1.20 OcnefMemoryUsageCrossedMinorThreshold

Table 5-20 OcnefMemoryUsageCrossedMinorThreshold

Field	Details
Description	"NEF Memory Usage for pod {{ $labels.pod }} has crossed the configured minor threshold (50%) (value={{ $value }}) of its limit."
Summary	"namespace: {{$labels.namespace}}, podname: {{$labels.pod}}, timestamp: {{ with query \"time()\" }}{{ . | first | value | humanizeTimestamp }}{{ end }} : Memory Usage of pod exceeded 50% of its limit."
Severity	Minor
Condition	A pod has reached the configured minor threshold (50%) of its memory resource limits.
OID	1.3.6.1.4.1.323.5.3.39.1.2.7020
Metric Used	'container_memory_usage_bytes''container_spec_memory_limit_bytes' **Note**: This is a Kubernetes metric used for instance availability monitoring. If the metric is not available, use the similar metric as exposed by the monitoring system.
Recommended Actions	The alert gets cleared when the memory utilization falls below the Minor Threshold or crosses the major threshold, in which case OcnefMemoryUsageCrossedMajorThreshold alert is raised. **Note**: The threshold is configurable in the NefAlertrules alert file. In case the issue persists, capture all the outputs for the above steps and contact My Oracle Support. **Note**: Use CNC NF Data Collector tool for capturing logs. For more information on how to collect logs using Data Collector tool, see *Oracle Communications Cloud Native Core, Network Function Data Collector User Guide*.

5.1.21 OcnefMemoryUsageCrossedMajorThreshold

Table 5-21 OcnefMemoryUsageCrossedMajorThreshold

Field	Details
Description	"NEF Memory Usage for pod {{ $labels.pod }} has crossed the configured major threshold (60%) (value = {{ $value }}) of its limit."
Summary	"namespace: {{$labels.namespace}}, podname: {{$labels.pod}}, timestamp: {{ with query \"time()\" }}{{ . | first | value | humanizeTimestamp }}{{ end }} : Memory Usage of pod exceeded 60% of its limit."
Severity	Major
Condition	A pod has reached the configured major threshold (60%) of its memory resource limits.
OID	1.3.6.1.4.1.323.5.3.39.1.2.7021
Metric Used	'container_memory_usage_bytes' 'container_spec_memory_limit_bytes' **Note**: This is a Kubernetes metric used for instance availability monitoring. If the metric is not available, use the similar metric as exposed by the monitoring system.

Table 5-21 (Cont.) OcnefMemoryUsageCrossedMajorThreshold

Field	Details
Recommended Actions	The alert gets cleared when the memory utilization falls below the Major Threshold or crosses the critical threshold, in which case OcnefMemoryUsageCrossedCriticalThreshold alert is raised. **Note**: The threshold is configurable in the NefAlertrules alert file. In case the issue persists, capture all the outputs for the above steps and contact My Oracle Support. **Note**: Use CNC NF Data Collector tool for capturing logs. For more information on how to collect logs using Data Collector tool, see *Oracle Communications Cloud Native Core, Network Function Data Collector User Guide*.

5.1.22 OcnefMemoryUsageCrossedCriticalThreshold

Table 5-22 OcnefMemoryUsageCrossedCriticalThreshold

Field	Details
Description	"NEF Memory Usage for pod {{ $labels.pod }} has crossed the configured major threshold (70%) (value = {{ $value }}) of its limit."
Summary	"namespace: {{$labels.namespace}}, podname: {{$labels.pod}}, timestamp: {{ with query \"time()\" }}{{ . | first | value | humanizeTimestamp }}{{ end }} : Memory Usage of pod exceeded 70% of its limit."
Severity	Critical
Condition	A pod has reached the configured critical threshold (70%) of its memory resource limits.
OID	1.3.6.1.4.1.323.5.3.39.1.2.7022
Metric Used	'container_memory_usage_bytes' 'container_spec_memory_limit_bytes' **Note**: This is a Kubernetes metric used for instance availability monitoring. If the metric is not available, use a similar metric as exposed by the monitoring system.
Recommended Actions	The alert gets cleared when the memory utilization falls below the Critical threshold. **Note**: The threshold is configurable in the NefAlertrules alert file. In case the issue persists, capture all the outputs for the above steps and contact My Oracle Support. **Note**: Use CNC NF Data Collector tool for capturing logs. For more information on how to collect logs using Data Collector tool, see *Oracle Communications Cloud Native Core, Network Function Data Collector User Guide*.

5.1.23 OcnefIngressGatewayServiceDown

Table 5-23 OcnefIngressGatewayServiceDown

Field	Details
Description	"NEF Ingress-Gateway service {{$labels.app_kubernetes_io_name}} is down"
Summary	"kubernetes_namespace: {{$labels.kubernetes_namespace}}, timestamp: {{ with query \"time()\" }}{{ . | first | value | humanizeTimestamp }}{{ end }} : Ingress-gateway service down"
Severity	Critical
Condition	None of the pods of the Ingress-Gateway microservice is available.
Metric Used	'up' **Note**: This is a Prometheus metric used for instance availability monitoring. If this metric is not available, use a similar metric as exposed by the monitoring system.

Recommended Actions	The alert is cleared when the Ingress Gateway service is available.
	Steps:
	1. To check the orchestration logs of Ingress Gateway service and check for liveness or readiness probe failures, do the following:
	a. Run the following command to check the pod status:
	`$ kubectl get po -n <namespace>`
	b. Run the following command to analyze the error condition of the pod that is not in the *Running* state:
	`$ kubectl describe pod <pod name not in Running state> -n <namespace>`
	Where <pod name not in Running state> indicates the pod that is not in the *Running* state.
	2. Refer to the application logs on Kibana and filter based on Ingress Gateway service names. Check for ERROR WARNING logs related to thread exceptions.
	3. Depending on the failure reason, take the resolution steps.
	4. In case the issue persists, capture all the outputs for the above steps and contact My Oracle Support.
	Note: Use CNC NF Data Collector tool for capturing logs. For more information on how to collect logs using Data Collector tool, see *Oracle Communications Cloud Native Core, Network Function Data Collector User Guide*.

5.1.24 OcnefApiRouterServiceDown

Table 5-24 OcnefApiRouterServiceDown

Field	Details
Description	"NEF API Router service {{$labels.app_kubernetes_io_name}} is down"
Summary	"namespace: {{$labels.namespace}}, podname: {{$labels.kubernetes_pod_name}}, timestamp: {{ with query "time()" }}{{ . | first | value | humanizeTimestamp }}{{ end }} : ApiRouter service down"
Severity	Critical
Condition	The API Router service is down.
Metric Used	'up' **Note:** This is a Prometheus metric used for instance availability monitoring. If this metric is not available, use a similar metric as exposed by the monitoring system.

Recommended Actions	The alert is cleared when the NEF API Router service is available.
	Steps:
	1. To check the orchestration logs of `ocnef_expgw_apirouter` service and check for liveness or readiness probe failures, do the following:
	a. Run the following command to check the pod status:
	`$ kubectl get pod -n <namespace>`
	b. Run the following command to analyze the error condition of the pod that is not in the **Running** state:
	`$ kubectl describe pod <pod name not in Running state> -n <namespace>`
	Where <pod name not in Running state> indicates the pod that is not in the Running state.
	2. Refer the application logs on Kibana and filter based on `ocnef_expgw_apirouter` service names. Check for ERROR WARNING logs related to thread exceptions.
	3. Check the DB status. For more information on how to check the DB status, see Oracle Communications Cloud Native Core, cnDBTier User Guide.
	4. Depending on the failure reason, take the resolution steps.
	5. In case the issue persists, capture all the outputs for the above steps and contact My Oracle Support.
	Note: Use CNC NF Data Collector tool for capturing logs. For more information on how to collect logs using Data Collector tool, see Oracle Communications Cloud Native Core, Network Function Data Collector User Guide.

5.1.25 OcnefFiveGcAgentServiceDown

Table 5-25 OcnefFiveGcAgentServiceDown

Field	Details
Description	"NEF FiveGc Agent service down {{$labels.app_kubernetes_io_name}} is down"
Summary	"kubernetes_namespace: {{$labels.kubernetes_namespace}}, timestamp: {{ with query \"time()\" }}{{ . | first | value | humanizeTimestamp }}{{ end }} : FiveGc Agent service down"
Severity	Critical
Condition	The 5GC Agent service is down.
Metric Used	'up' **Note:** This is a Prometheus metric used for instance availability monitoring. If this metric is not available, use a similar metric as exposed by the monitoring system.

Recommended Actions	The alert is cleared when the NEF 5GC Agent service is available.
	Steps:
	1. To check the orchestration logs of `5gcagent` service and check for liveness or readiness probe failures, do the following:
	a. Run the following command to check the pod status:
	`$ kubectl get pod -n <namespace>`
	b. Run the following command to analyze the error condition of the pod that is not in the **Running** state:
	`$ kubectl describe pod <pod name not in Running state> -n <namespace>`
	Where <pod name not in Running state> indicates the pod that is not in the Running state.
	2. Refer the application logs on Kibana and filter based on `5gcagent` service names. Check for ERROR WARNING logs related to thread exceptions.
	3. Check the DB status. For more information on how to check the DB status, see Oracle Communications Cloud Native Core, cnDBTier User Guide.
	4. Depending on the failure reason, take the resolution steps.
	5. In case the issue persists, capture all the outputs for the above steps and contact My Oracle Support.
	Note: Use CNC NF Data Collector tool for capturing logs. For more information on how to collect logs using Data Collector tool, see Oracle Communications Cloud Native Core, Network Function Data Collector User Guide.

5.1.26 OcnefMonitoringEventServiceDown

Table 5-26 OcnefMonitoringEventServiceDown

Field	Details
Description	"NEF MonitoringEvent service {{$labels.app_kubernetes_io_name}} is down"
Summary	"kubernetes_namespace: {{$labels.kubernetes_namespace}}, timestamp: {{ with query \"time()\" }}{{ . | first | value | humanizeTimestamp }}{{ end }} : MonitoringEvent service down"
Severity	Critical
Condition	The Monitoring Event (ME) service is down.
Metric Used	'up' **Note:** This is a Prometheus metric used for instance availability monitoring. If this metric is not available, use a similar metric as exposed by the monitoring system.

Recommended Actions	The alert is cleared when the NEF Monitoring Event (ME) service is available.
	Steps:
	1. To check the orchestration logs of `ocnef_monitoring_events` service and check for liveness or readiness probe failures, do the following:
	a. Run the following command to check the pod status:
	`$ kubectl get pod -n <namespace>`
	b. Run the following command to analyze the error condition of the pod that is not in the **Running** state:
	`$ kubectl describe pod <pod name not in Running state> -n <namespace>`
	Where <pod name not in Running state> indicates the pod that is not in the Running state.
	2. Refer the application logs on Kibana and filter based on `ocnef_monitoring_events` service names. Check for ERROR WARNING logs related to thread exceptions.
	3. Check the DB status. For more information on how to check the DB status, see Oracle Communications Cloud Native Core, cnDBTier User Guide.
	4. Depending on the failure reason, take the resolution steps.
	5. In case the issue persists, capture all the outputs for the above steps and contact My Oracle Support.
	Note: Use CNC NF Data Collector tool for capturing logs. For more information on how to collect logs using Data Collector tool, see Oracle Communications Cloud Native Core, Network Function Data Collector User Guide.

5.1.27 OcnefCCFClientServiceDown

Table 5-27 OcnefCCFClientServiceDown

Field	Details
Description	"NEF CCFClient service {{$labels.app_kubernetes_io_name}} is down"
Summary	"kubernetes_namespace: {{$labels.kubernetes_namespace}}, timestamp: {{ with query \"time()\" }}{{ . | first | value | humanizeTimestamp }}{{ end }} : CCFClient service down"
Severity	Critical
Condition	The CCF Client service is down.
Metric Used	'up' **Note:** This is a Prometheus metric used for instance availability monitoring. If this metric is not available, use a similar metric as exposed by the monitoring system.

Recommended Actions	The alert is cleared when the NEF CCF Client service is available.
	Steps:
	1. To check the orchestration logs of `ocnef_ccfclient` service and check for liveness or readiness probe failures, do the following:
	a. Run the following command to check the pod status:
	`$ kubectl get pod -n <namespace>`
	b. Run the following command to analyze the error condition of the pod that is not in the **Running** state:
	`$ kubectl describe pod <pod name not in Running state> -n <namespace>`
	Where <pod name not in Running state> indicates the pod that is not in the Running state.
	2. Refer the application logs on Kibana and filter based on `ocnef_ccfclient` service names. Check for ERROR WARNING logs related to thread exceptions.
	3. Check the DB status. For more information on how to check the DB status, see Oracle Communications Cloud Native Core, cnDBTier User Guide.
	4. Depending on the failure reason, take the resolution steps.
	5. In case the issue persists, capture all the outputs for the above steps and contact My Oracle Support.
	Note: Use CNC NF Data Collector tool for capturing logs. For more information on how to collect logs using Data Collector tool, see Oracle Communications Cloud Native Core, Network Function Data Collector User Guide.

5.1.28 OcnefExpiryAuditorServiceDown

Table 5-28 OcnefExpiryAuditorServiceDown

Field	Details
Description	"NEF Expiry Auditor service {{$labels.app_kubernetes_io_name}} is down"
Summary	"kubernetes_namespace: {{$labels.kubernetes_namespace}}, timestamp: {{ with query \"time()\" }}{{ . | first | value | humanizeTimestamp }}{{ end }} : Expiry Auditor service down"
Severity	Critical
Condition	The expiry auditor service is down.
Metric Used	'up' **Note:** This is a Prometheus metric used for instance availability monitoring. If this metric is not available, use a similar metric as exposed by the monitoring system.

Recommended Actions	The alert is cleared when the NEF Expiry Auditor service is available.
	Steps:
	1. To check the orchestration logs of `ocnef-expiry-auditor` service and check for liveness or readiness probe failures, do the following:
	a. Run the following command to check the pod status:
	`$ kubectl get pod -n <namespace>`
	b. Run the following command to analyze the error condition of the pod that is not in the **Running** state:
	`$ kubectl describe pod <pod name not in Running state> -n <namespace>`
	Where <pod name not in Running state> indicates the pod that is not in the Running state.
	2. Refer the application logs on Kibana and filter based on `ocnef-expiry-auditor` service names. Check for ERROR WARNING logs related to thread exceptions.
	3. Check the DB status. For more information on how to check the DB status, see Oracle Communications Cloud Native Core, cnDBTier User Guide.
	4. Depending on the failure reason, take the resolution steps.
	5. In case the issue persists, capture all the outputs for the above steps and contact My Oracle Support.
	Note: Use CNC NF Data Collector tool for capturing logs. For more information on how to collect logs using Data Collector tool, see Oracle Communications Cloud Native Core, Network Function Data Collector User Guide.

5.1.29 OcnefQOSServiceDown

Table 5-29 OcnefQOSServiceDown

Field	Details
Description	"NEF QOS service {{$labels.app_kubernetes_io_name}} is down"
Summary	namespace: {{$labels.namespace}}, podname: {{$labels.kubernetes_pod_name}}, timestamp: {{ with query "time()" }}{{ . | first | value | humanizeTimestamp }}{{ end }} : QOS service down
Severity	Critical
Condition	The QoS service is down.
Metric Used	'up' **Note:** This is a Prometheus metric used for instance availability monitoring. If this metric is not available, use a similar metric as exposed by the monitoring system.

Recommended Actions	The alert is cleared when the NEF QoS service is available.
	Steps:
	1. To check the orchestration logs of `ocnef-qualityofservice` service and check for liveness or readiness probe failures, do the following:
	a. Run the following command to check the pod status:
	`$ kubectl get pod -n <namespace>`
	b. Run the following command to analyze the error condition of the pod that is not in the **Running** state:
	`$ kubectl describe pod <pod name not in Running state> -n <namespace>`
	Where <pod name not in Running state> indicates the pod that is not in the Running state.
	2. Refer the application logs on Kibana and filter based on `ocnef-expiry-auditor` service names. Check for ERROR WARNING logs related to thread exceptions.
	3. Check the DB status. For more information on how to check the DB status, see Oracle Communications Cloud Native Core, cnDBTier User Guide.
	4. Depending on the failure reason, take the resolution steps.
	5. In case the issue persists, capture all the outputs for the above steps and contact My Oracle Support.
	Note: Use CNC NF Data Collector tool for capturing logs. For more information on how to collect logs using Data Collector tool, see Oracle Communications Cloud Native Core, Network Function Data Collector User Guide.

5.1.30 OcnefDTServiceDown

Table 5-30 OcnefDTServiceDown

Field	Details
Description	"OCNEF Device Trigger service {{$labels.app_kubernetes_io_name}} is down"
Summary	"namespace: {{$labels.namespace}}, podname: {{$labels.kubernetes_pod_name}}, timestamp: {{ with query "time()" }}{{ . | first | value | humanizeTimestamp }}{{ end }} : DT service down"
Severity	Critical
Condition	The Device Trigger service is down.
Metric Used	'up' **Note:** This is a Prometheus metric used for instance availability monitoring. If this metric is not available, use a similar metric as exposed by the monitoring system.

Recommended Actions	The alert is cleared when the NEF DT service is available.
	Steps:
	1. To check the orchestration logs of `ocnef-devicetrigger` service and check for liveness or readiness probe failures, do the following:
	a. Run the following command to check the pod status:
	`$ kubectl get pod -n <namespace>`
	b. Run the following command to analyze the error condition of the pod that is not in the **Running** state:
	`$ kubectl describe pod <pod name not in Running state> -n <namespace>`
	Where <pod name not in Running state> indicates the pod that is not in the Running state.
	2. Refer to application logs on Kibana and filter them based on `ocnef_expgw_apimgr` service name. Check for ERROR and WARNING logs related to thread exceptions.
	3. Check the DB status. For more information on how to check the DB status, see Oracle Communications Cloud Native Core, cnDBTier User Guide.
	4. Depending on the failure reason, take the resolution steps.
	5. In case the issue persists, capture all the outputs for the above steps and contact My Oracle Support.
	Note: Use CNC NF Data Collector tool for capturing logs. For more information on how to collect logs using Data Collector tool, see Oracle Communications Cloud Native Core, Network Function Data Collector User Guide.

5.2 Application Level Alerts

This section lists the application level alerts for NEF.

5.2.1 MENotificationFailureRateCrossedThreshold

Table 5-31 MENotificationFailureRateCrossedThreshold

Field	Details
Description	"Failure Rate of Delete ME Notifications Is Crossing the Threshold (10%)"
Summary	"namespace: {{$labels.kubernetes_namespace}}, timestamp: {{ with query \"time()\" }}{{ . | first | value | humanizeTimestamp }}{{ end }} : Failure Rate Of delete ME Subscriptions requests is above 10 percent of total requests."
Severity	Error
Condition	The failure rate of the DELETE Monitoring Event notification requests is reaching the threshold value.
Metric Used	ocnef_me_af_resp_total
Recommended Actions	The alert is cleared when the failure rate of DELETE ME notification requests is below the threshold. Steps: 1. Check for pod logs on Kibana for ERROR WARN logs. 2. In case the issue persists, capture all the outputs for the above steps and contact My Oracle Support.

5.2.2 MEDeleteSubscriptionFailureRateCrossedThreshold

Table 5-32 MEDeleteSubscriptionFailureRateCrossedThreshold

Field	Details
Description	"Failure Rate of Delete ME Subscriptions Is Crossing the Threshold (10%)"
Summary	"namespace: {{$labels.kubernetes_namespace}}, timestamp: {{ with query \"time()\" }}{{ . | first | value | humanizeTimestamp }}{{ end }} : Failure Rate Of delete ME Subscriptions requests is above 10 percent of total requests."
Severity	Error
Condition	The failure rate of the Monitoring Event subscription deletion requests is reaching the threshold value.
Metric Used	ocnef_me_af_resp_total
Recommended Actions	The alert is cleared when the failure rate of Monitoring Event subscription deletion requests is below the threshold. Steps: 1. Check for pod logs on Kibana for ERROR WARN logs. 2. In case the issue persists, capture all the outputs for the above steps and contact My Oracle Support.

5.2.3 MEAddSubscriptionFailureRateCrossedThreshold

Table 5-33 MEAddSubscriptionFailureRateCrossedThreshold

Field	Details
Description	"Failure Rate of ME Subscriptions Is Crossing the Threshold (10%)"
Summary	"namespace: {{$labels.namespace}}, timestamp: {{ with query "time()" }} {{ . | first | value | humanizeTimestamp }}{{ end }} : Failure Rate Of ME Subscriptions requests is above 10 percent of total requests."
Severity	Error
Condition	The failure rate of the Monitoring Event subscription requests is reaching the threshold value.
Metric Used	ocnef_me_af_resp_total
Recommended Actions	The alert is cleared when the failure rate of Monitoring Event subscription requests is below the threshold. Steps: 1. Check for pod logs on Kibana for ERROR WARN logs. 2. In case the issue persists, capture all the outputs for the above steps and contact My Oracle Support.

5.2.4 AEFApiRouterOauthValidationFailureRateCrossedThreshold

Table 5-34 AEFApiRouterOauthValidationFailureRateCrossedThreshold

Field	Details
Description	"Failure Rate of API Router Oauth Validation Is Crossing the Threshold (10%)"
Summary	"{{$labels.namespace}}, timestamp: {{ with query "time()" }}{{ . | first | value | humanizeTimestamp }}{{ end }} : Failure Rate Of Oauth Validation is above 10 percent of total requests."
Severity	Error
Condition	The failure rate of the OAuth validations at API Router is reaching the threshold value.
OID	1.3.6.1.4.1.323.5.3.39.1.2.7032
Metric Used	ocnef_aef_apirouter_resp_total
Recommended Actions	The alert is cleared when the failure rate of OAuth validations at API Router is below the threshold. Steps: 1. Check for pod logs on Kibana for ERROR WARN logs. 2. In case the issue persists, capture all the outputs for the above steps and contact My Oracle Support.

5.2.5 FiveGcInvalidConfiguration

Table 5-35 FiveGcInvalidConfiguration

Field	Details
Description	"Invalid Configuration For Five GC Service"
Summary	"namespace: {{$labels.kubernetes_namespace}}, timestamp: {{ with query \"time()\" }}{{ . | first | value | humanizeTimestamp }}{{ end }} : Invalid Configuration For Five GC Service."
Severity	Error
Condition	Invalid configuration of the 5GCAgent service.
Metric Used	ocnef_5gc_invalid_config
Recommended Actions	The alert is cleared when the 5GCAgent service configuration are valid. Steps: 1. Check for pod logs on Kibana for ERROR WARN logs. 2. In case the issue persists, capture all the outputs for the above steps and contact My Oracle Support.

5.2.6 MENotificationFailureRateCrossedThreshold

Table 5-36 MENotificationFailureRateCrossedThreshold

Field	Details
Description	"Failure Rate of Delete ME Notifications Is Crossing the Threshold (10%)"
Summary	"kubernetes_namespace: {{$labels.kubernetes_namespace}}, timestamp: {{ with query \"time()\" }}{{ . | first | value | humanizeTimestamp }}{{ end }} : Failure Rate Of delete ME Subscriptions requests is above 10 percent of total requests."
Severity	Error
Condition	The failure rate of the DELETE Monitoring Event notification requests is reaching the threshold value.
Metric Used	ocnef_efc_svc_me_notification_resp_total
Recommended Actions	The alert is cleared when the failure rate of DELETE ME notification requests is below the threshold. Steps: 1. Check for pod logs on Kibana for ERROR WARN logs. 2. In case the issue persists, capture all the outputs for the above steps and contact My Oracle Support.

5.2.7 MENotificationFailureRateCrossedThreshold

Table 5-37 MENotificationFailureRateCrossedThreshold

Field	Details
Description	"Failure Rate of Delete ME Notifications Is Crossing the Threshold (10%)"
Summary	"kubernetes_namespace: {{$labels.kubernetes_namespace}}, timestamp: {{ with query \"time()\" }}{{ . | first | value | humanizeTimestamp }}{{ end }} : Failure Rate Of delete ME Subscriptions requests is above 10 percent of total requests."
Severity	Error
Condition	The failure rate of the DELETE Monitoring Event notification requests is reaching the threshold value.
Metric Used	ocnef_efc_svc_me_notification_resp_total
Recommended Actions	The alert is cleared when the failure rate of DELETE ME notification requests is below the threshold. Steps: 1. Check for pod logs on Kibana for ERROR WARN logs. 2. In case the issue persists, capture all the outputs for the above steps and contact My Oracle Support.

5.2.8 MENotificationFailureRateCrossedThreshold

Table 5-38 MENotificationFailureRateCrossedThreshold

Field	Details
Description	"Failure Rate of Delete ME Notifications Is Crossing the Threshold (10%)"
Summary	"kubernetes_namespace: {{$labels.kubernetes_namespace}}, timestamp: {{ with query \"time()\" }}{{ . | first | value | humanizeTimestamp }}{{ end }} : Failure Rate Of delete ME Subscriptions requests is above 10 percent of total requests."
Severity	Error
Condition	The failure rate of the DELETE Monitoring Event notification requests is reaching the threshold value.
Metric Used	ocnef_efc_svc_me_notification_resp_total
Recommended Actions	The alert is cleared when the failure rate of DELETE ME notification requests is below the threshold. Steps: 1. Check for pod logs on Kibana for ERROR WARN logs. 2. In case the issue persists, capture all the outputs for the above steps and contact My Oracle Support.

5.2.9 OcnefAllSiteStatus

Table 5-39 OcnefAllSiteStatus

Field	Details
Description	"Alert for any NEF sites status if SUSPENDED in Georedundant setup"
Summary	namespace: {{$labels.namespace}}, timestamp: {{ with query "time()" }} {{ . | first | value | humanizeTimestamp }}{{ end }} : Alert for any NEF sites status if SUSPENDED in Georedundant setup
Severity	Error
Condition	An NEF site of a georedundant deployment is in Suspended state.
OID	1.3.6.1.4.1.323.5.3.39.1.2.7040
Metric Used	ocnef_all_site_status
Recommended Actions	The alert is cleared when all the sites in a georedundant deployment are UP. 1. Check for pod logs on Kibana for ERROR WARN logs. 2. In case the issue persists, capture all the outputs for the above steps and contact My Oracle Support.

5.2.10 OcnefDBReplicationStatus

Table 5-40 OcnefDBReplicationStatus

Field	Details
Description	"Alert for NEF sites status if DB Replication down in Georedundant setup"
Summary	namespace: {{$labels.namespace}}, timestamp: {{ with query "time()" }} {{ . | first | value | humanizeTimestamp }}{{ end }} : Alert for NEF sites status if DB Replication down in Georedundant setup
Severity	Error
Condition	The database replication channel status between the given site and the georedundant site(s) is inactive. The alert is raised per replication channel. The alarm is raised or cleared only if the georedundancy feature is enabled.
OID	1.3.6.1.4.1.323.5.3.39.1.2.7041
Metric Used	ocnef_db_replication_status
Recommended Actions	The alert is cleared when the database channel replication status between the given site and the georedundant site(s) is UP. For more information on how to check the database replication status, see *Oracle Communications Cloud Native Core, cnDBTier User Guide*.

5.2.11 CHFAddChargingDataRequestFailureRateCrossedErrorThreshold

Table 5-41 CHFAddChargingDataRequestFailureRateCrossedErrorThreshold

Field	Details
Description	Failure rate of CHF Create Charging Data request is crossing the threshold (10%).
Summary	namespace: {{$labels.namespace}}, timestamp: {{ with query "time()" }} {{ . | first | value | humanizeTimestamp }}{{ end }}
Severity	error
Condition	Failure rate of CHF Create Charging Data request is above 10 percent of total requests.
Metric Used	ocnef_chf_qos_resp_total
Recommended Actions	The alert is cleared when the failure rate of CHF Requests is below the failure threshold. Steps: 1. Check for pod logs on Kibana for ERROR WARN logs. 2. In case the issue persists, capture all the outputs for the above steps and contact My Oracle Support.

5.2.12 CHFAddChargingDataRequestFailureRateCrossedCriticalThreshold

Table 5-42 CHFAddChargingDataRequestFailureRateCrossedCriticalThreshold

Field	Details
Description	Failure rate of CHF Create Charging Data request is crossing the threshold (25%).
Summary	namespace: {{$labels.namespace}}, timestamp: {{ with query "time()" }} {{ . | first | value | humanizeTimestamp }}{{ end }}
Severity	critical
Condition	Failure rate of CHF Create Charging Data request is above 25 percent of total requests.
Metric Used	ocnef_chf_qos_resp_total
Recommended Actions	The alert is cleared when the failure rate of CHF Requests is below the failure threshold. Steps: 1. Check for pod logs on Kibana for ERROR WARN logs. 2. In case the issue persists, capture all the outputs for the above steps and contact My Oracle Support.

5.2.13 CHFAddChargingDataRequestFailureRateCrossedMinorThreshold

Table 5-43 CHFAddChargingDataRequestFailureRateCrossedMinorThreshold

Field	Details
Description	Failure rate of CHF Create Charging Data request is crossing the threshold (5%).
Summary	namespace: {{$labels.namespace}}, timestamp: {{ with query "time()" }} {{ . | first | value | humanizeTimestamp }}{{ end }}
Severity	minor
Condition	Failure rate of CHF Create Charging Data request is above 5 percent of total requests.
Metric Used	ocnef_chf_qos_resp_total
Recommended Actions	The alert is cleared when the failure rate of CHF Requests is below the failure threshold. Steps: 1. Check for pod logs on Kibana for ERROR WARN logs. 2. In case the issue persists, capture all the outputs for the above steps and contact My Oracle Support.

5.2.14 MeEPCAddSubscriptionFailureRateCrossedThreshold

Table 5-44 MeEPCAddSubscriptionFailureRateCrossedThreshold

Field	Details
Description	Failure rate of ME subscriptions to EPC crossing the threshold (10%).
Summary	namespace: {{$labels.namespace}}, timestamp: {{ with query "time()" }} {{ . | first | value | humanizeTimestamp }}{{ end }}
Severity	error
Condition	Failure Rate of ME EPC Subscriptions requests is above 10 percent of total requests.
Metric Used	ocnef_me_epc_sub_total
Recommended Actions	The alert is cleared when the failure rate of CHF Requests is below the failure threshold. Steps: 1. Check for pod logs on Kibana for ERROR WARN logs. 2. In case the issue persists, capture all the outputs for the above steps and contact My Oracle Support.

5.2.15 DiameterGwT6InvocationFailureRateCrossedThreshold

Table 5-45 DiameterGwT6InvocationFailureRateCrossedThreshold

Field	Details
Description	Failure rate of Diameter Gateway T6x Invocation requests crossing the threshold (10%).

Table 5-45 (Cont.) DiameterGwT6InvocationFailureRateCrossedThreshold

Field	Details
Summary	namespace: {{$labels.namespace}}, timestamp: {{ with query "time()" }} {{ . | first | value | humanizeTimestamp }}{{ end }}
Severity	error
Condition	Failure rate of Diameter Gateway T6x Invocation requests crossing the threshold (10%).
Metric Used	ocnef_diamgw_diam_resp_total
Recommended Actions	The alert is cleared when the failure rate of CHF Requests is below the failure threshold. Steps: 1. Check for pod logs on Kibana for ERROR WARN logs. 2. In case the issue persists, capture all the outputs for the above steps and contact My Oracle Support.

5.2.16 DiameterGwT4InvocationFailureRateCrossedThreshold

Table 5-46 DiameterGwT4InvocationFailureRateCrossedThreshold

Field	Details
Description	Failure rate of Diameter Gateway T4 Invocation requests crossing the threshold (10%).
Summary	namespace: {{$labels.namespace}}, timestamp: {{ with query "time()" }} {{ . | first | value | humanizeTimestamp }}{{ end }}
Severity	error
Condition	Failure rate of Diameter Gateway T4 Invocation requests crossing the threshold (10%).
Metric Used	ocnef_diamgw_diam_resp_total
Recommended Actions	The alert is cleared when the failure rate of CHF Requests is below the failure threshold. Steps: 1. Check for pod logs on Kibana for ERROR WARN logs. 2. In case the issue persists, capture all the outputs for the above steps and contact My Oracle Support.

5.2.17 DiameterGwRxInvocationFailureRateCrossedThreshold

Table 5-47 DiameterGwRxInvocationFailureRateCrossedThreshold

Field	Details
Description	Failure rate of Diameter Gateway Rx Invocation requests crossing the threshold (10%).
Summary	namespace: {{$labels.namespace}}, timestamp: {{ with query "time()" }} {{ . | first | value | humanizeTimestamp }}{{ end }}
Severity	error

Table 5-47 (Cont.) DiameterGwRxInvocationFailureRateCrossedThreshold

Field	Details
Condition	Failure rate of Diameter Gateway Rx Invocation requests crossing the threshold (10%).
Metric Used	ocnef_diamgw_diam_resp_total
Recommended Actions	The alert is cleared when the failure rate of CHF Requests is below the failure threshold. Steps: 1. Check for pod logs on Kibana for ERROR WARN logs. 2. In case the issue persists, capture all the outputs for the above steps and contact My Oracle Support.

5.2.18 DiameterGwSgdT4InvocationFailureRateCrossedThreshold

Table 5-48 DiameterGwSgdT4InvocationFailureRateCrossedThreshold

Field	Details
Description	Failure rate of Diameter Gateway SgdT4 Invocation requests crossing the threshold (10%).
Summary	namespace: {{$labels.namespace}}, timestamp: {{ with query "time()" }} {{ . | first | value | humanizeTimestamp }}{{ end }}
Severity	error
Condition	Failure rate of Diameter Gateway SgdT4 Invocation requests crossing the threshold (10%).
Metric Used	ocnef_diamgw_diam_resp_total
Recommended Actions	The alert is cleared when the failure rate of CHF Requests is below the failure threshold. Steps: 1. Check for pod logs on Kibana for ERROR WARN logs. 2. In case the issue persists, capture all the outputs for the above steps and contact My Oracle Support.

5.2.19 DiameterGwT6NotificationFailureRateCrossedThreshold

Table 5-49 DiameterGwT6NotificationFailureRateCrossedThreshold

Field	Details
Description	Failure rate of Diameter Gateway T6x Notification requests crossing the threshold (10%).
Summary	namespace: {{$labels.namespace}}, timestamp: {{ with query "time()" }} {{ . | first | value | humanizeTimestamp }}{{ end }}
Severity	error
Condition	Failure rate of Diameter Gateway T6x Notification requests crossing the threshold (10%).
Metric Used	ocnef_diamgw_diam_resp_total

Table 5-49 (Cont.) DiameterGwT6NotificationFailureRateCrossedThreshold

Field	Details
Recommended Actions	The alert is cleared when the failure rate of CHF Requests is below the failure threshold. Steps: 1. Check for pod logs on Kibana for ERROR WARN logs. 2. In case the issue persists, capture all the outputs for the above steps and contact My Oracle Support.

5.2.20 DiameterGwT4NotificationFailureRateCrossedThreshold

Table 5-50 DiameterGwT4NotificationFailureRateCrossedThreshold

Field	Details
Description	Failure rate of Diameter Gateway T4 Notification requests crossing the threshold (10%).
Summary	namespace: {{$labels.namespace}}, timestamp: {{ with query "time()" }} {{ . | first | value | humanizeTimestamp }}{{ end }}
Severity	error
Condition	Failure rate of Diameter Gateway T4 Notification requests crossing the threshold (10%).
Metric Used	ocnef_diamgw_diam_resp_total
Recommended Actions	The alert is cleared when the failure rate of CHF Requests is below the failure threshold. Steps: 1. Check for pod logs on Kibana for ERROR WARN logs. 2. In case the issue persists, capture all the outputs for the above steps and contact My Oracle Support.

5.2.21 DiameterGwRxNotificationFailureRateCrossedThreshold

Table 5-51 DiameterGwRxNotificationFailureRateCrossedThreshold

Field	Details
Description	Failure rate of Diameter Gateway Rx Notification requests crossing the threshold (10%).
Summary	namespace: {{$labels.namespace}}, timestamp: {{ with query "time()" }} {{ . | first | value | humanizeTimestamp }}{{ end }}
Severity	error
Condition	Failure rate of Diameter Gateway Rx Notification requests crossing the threshold (10%).
Metric Used	ocnef_diamgw_diam_resp_total

Table 5-51 (Cont.) DiameterGwRxNotificationFailureRateCrossedThreshold

Field	Details
Recommended Actions	The alert is cleared when the failure rate of CHF Requests is below the failure threshold. Steps: 1. Check for pod logs on Kibana for ERROR WARN logs. 2. In case the issue persists, capture all the outputs for the above steps and contact My Oracle Support.

5.2.22 DiameterGwSgdT4NotificationFailureRateCrossedThreshold

Table 5-52 DiameterGwSgdT4NotificationFailureRateCrossedThreshold

Field	Details
Description	Failure rate of Diameter Gateway SgdT4 Notification requests crossing the threshold (10%).
Summary	namespace: {{$labels.namespace}}, timestamp: {{ with query "time()" }} {{ . | first | value | humanizeTimestamp }}{{ end }}
Severity	error
Condition	Failure rate of Diameter Gateway SgdT4 Notification requests crossing the threshold (10%).
Metric Used	ocnef_diamgw_diam_resp_total
Recommended Actions	The alert is cleared when the failure rate of CHF Requests is below the failure threshold. Steps: 1. Check for pod logs on Kibana for ERROR WARN logs. 2. In case the issue persists, capture all the outputs for the above steps and contact My Oracle Support.

5.2.23 DiameterGwT6TranslationFailureRateCrossedThreshold

Table 5-53 DiameterGwT6TranslationFailureRateCrossedThreshold

Field	Details
Description	Failure Rate of T6x Translations In Diameter GW is above 10 percent of total requests.
Summary	namespace: {{$labels.namespace}}, timestamp: {{ with query "time()" }} {{ . | first | value | humanizeTimestamp }}{{ end }}
Severity	error
Condition	Failure Rate of T6x Translations In Diameter GW is above 10 percent of total requests.
Metric Used	ocnef_diamgw_translator_request_total

Table 5-53 (Cont.) DiameterGwT6TranslationFailureRateCrossedThreshold

Field	Details
Recommended Actions	The alert is cleared when the failure rate of CHF Requests is below the failure threshold. Steps: 1. Check for pod logs on Kibana for ERROR WARN logs. 2. In case the issue persists, capture all the outputs for the above steps and contact My Oracle Support.

5.2.24 DiameterGwT4TranslationFailureRateCrossedThreshold

Table 5-54 DiameterGwT4TranslationFailureRateCrossedThreshold

Field	Details
Description	Failure Rate of T4 Translations In Diameter GW is above 10 percent of total requests.
Summary	namespace: {{$labels.namespace}}, timestamp: {{ with query "time()" }} {{ . | first | value | humanizeTimestamp }}{{ end }}
Severity	error
Condition	Failure Rate of T4 Translations In Diameter GW is above 10 percent of total requests.
Metric Used	ocnef_diamgw_translator_request_total
Recommended Actions	The alert is cleared when the failure rate of CHF Requests is below the failure threshold. Steps: 1. Check for pod logs on Kibana for ERROR WARN logs. 2. In case the issue persists, capture all the outputs for the above steps and contact My Oracle Support.

5.2.25 DiameterGwRxTranslationFailureRateCrossedThreshold

Table 5-55 DiameterGwRxTranslationFailureRateCrossedThreshold

Field	Details
Description	Failure Rate of Rx Translations In Diameter GW is above 10 percent of total requests.
Summary	namespace: {{$labels.namespace}}, timestamp: {{ with query "time()" }} {{ . | first | value | humanizeTimestamp }}{{ end }}
Severity	error
Condition	Failure Rate of Rx Translations In Diameter GW is above 10 percent of total requests.
Metric Used	ocnef_diamgw_translator_request_total

Table 5-55 (Cont.) DiameterGwRxTranslationFailureRateCrossedThreshold

Field	Details
Recommended Actions	The alert is cleared when the failure rate of CHF Requests is below the failure threshold. Steps: 1. Check for pod logs on Kibana for ERROR WARN logs. 2. In case the issue persists, capture all the outputs for the above steps and contact My Oracle Support.

5.2.26 DiameterGwSgdT4TranslationFailureRateCrossedThreshold

Table 5-56 DiameterGwSgdT4TranslationFailureRateCrossedThreshold

Field	Details
Description	Failure Rate of SgdT4 Translations in Diameter GW is above 10 percent of total requests.
Summary	namespace: {{$labels.namespace}}, timestamp: {{ with query "time()" }} {{ . | first | value | humanizeTimestamp }}{{ end }}
Severity	error
Condition	Failure Rate of SgdT4 Translations in Diameter GW is above 10 percent of total requests.
Metric Used	ocnef_diamgw_translator_request_total
Recommended Actions	The alert is cleared when the failure rate of CHF Requests is below the failure threshold. Steps: 1. Check for pod logs on Kibana for ERROR WARN logs. 2. In case the issue persists, capture all the outputs for the above steps and contact My Oracle Support.

5.2.27 MSISDNLessMoSMSRequestFailureRateCrossedCriticalThreshold

Table 5-57 MSISDNLessMoSMSRequestFailureRateCrossedCriticalThreshold

Field	Details
Description	Failure rate of MSISDNLess MO SMS notification request is crossing the threshold (25%).
Summary	namespace: {{$labels.namespace}}, timestamp: {{ with query "time()" }} {{ . | first | value | humanizeTimestamp }}{{ end }}
Severity	Critical
Condition	Failure rate Of MSISDNLess MO SMS request is above 25 percent of total requests.
Metric Used	ocnef_msisdnless_mo_sms_diamgw_notify_resp_total

Table 5-57 (Cont.) MSISDNLessMoSMSRequestFailureRateCrossedCriticalThreshold

Field	Details
Recommended Actions	The alert is cleared when the failure rate of notification requests is below the failure threshold. Steps: 1. Check for pod logs on Kibana for ERROR WARN logs. 2. In case the issue persists, capture all the outputs for the above steps and contact My Oracle Support.

5.2.28 MSISDNLessMoSMSRequestFailureRateCrossedMajorThreshold

Table 5-58 MSISDNLessMoSMSRequestFailureRateCrossedMajorThreshold

Field	Details
Description	Failure rate of MSISDNLess MO SMS notification request is crossing the threshold (10%).
Summary	namespace: {{$labels.namespace}}, timestamp: {{ with query "time()" }} {{ . | first | value | humanizeTimestamp }}{{ end }}
Severity	Major
Condition	Failure rate of MSISDNLess MO SMS request is above 10 percent of total requests.
Metric Used	ocnef_msisdnless_mo_sms_diamgw_notify_resp_total
Recommended Actions	The alert is cleared when the failure rate of notification requests is below the failure threshold. Steps: 1. Check for pod logs on Kibana for ERROR WARN logs. 2. In case the issue persists, capture all the outputs for the above steps and contact My Oracle Support.

5.2.29 MSISDNLessMoSMSRequestFailureRateCrossedMinorThreshold

Table 5-59 MSISDNLessMoSMSRequestFailureRateCrossedMinorThreshold

Field	Details
Description	Failure rate of MSISDNLess MO SMS notification request is crossing the threshold (5%).
Summary	namespace: {{$labels.namespace}}, timestamp: {{ with query "time()" }} {{ . | first | value | humanizeTimestamp }}{{ end }}
Severity	Minor
Condition	Failure rate of MSISDNLess MO SMS request is above 5 percent of total requests.
Metric Used	ocnef_msisdnless_mo_sms_diamgw_notify_resp_total

Table 5-59 (Cont.) MSISDNLessMoSMSRequestFailureRateCrossedMinorThreshold

Field	Details
Recommended Actions	The alert is cleared when the failure rate of notification requests is below the failure threshold. Steps: 1. Check for pod logs on Kibana for ERROR WARN logs. 2. In case the issue persists, capture all the outputs for the above steps and contact My Oracle Support.

5.2.30 MSISDNLessMoSMSShortCodeConfigMatchFailure

Table 5-60 MSISDNLessMoSMSShortCodeConfigMatchFailure

Field	Details
Description	Failure when shortcode configured doesn't match the shortcode from incoming request.
Summary	namespace: {{$labels.namespace}}, timestamp: {{ with query "time()" }} {{ . | first | value | humanizeTimestamp }}{{ end }}
Severity	Error
Condition	Failure when shortcode configured doesn't match the shortcode from SMSSC.
Metric Used	ocnef_diamgw_http_resp_total
Recommended Actions	The alert is cleared when the failure rate of notification requests is below the failure threshold. Steps: 1. Check for pod logs on Kibana for ERROR WARN logs. 2. In case the issue persists, capture all the outputs for the above steps and contact My Oracle Support.

5.2.31 QOSAddSubscriptionFailureRateCrossedThreshold

Table 5-61 QOSAddSubscriptionFailureRateCrossedThreshold

Field	Details
Description	Failure rate of QoS subscriptions is crossing the threshold (10%).
Summary	namespace: {{$labels.namespace}}, timestamp: {{ with query "time()" }} {{ . | first | value | humanizeTimestamp }}{{ end }}
Severity	Error
Condition	Failure rate of QoS subscription requests is above 10 percent of total requests.
Metric Used	ocnef_qos_af_resp_total

Table 5-61 (Cont.) QOSAddSubscriptionFailureRateCrossedThreshold

Field	Details
Recommended Actions	The alert is cleared when the failure rate of subscription requests is below the failure threshold. Steps: 1. Check for pod logs on Kibana for ERROR WARN logs. 2. In case the issue persists, capture all the outputs for the above steps and contact My Oracle Support.

5.2.32 QOSDeleteSubscriptionFailureRateCrossedThreshold

Table 5-62 QOSDeleteSubscriptionFailureRateCrossedThreshold

Field	Details
Description	Failure rate of delete QoS subscriptions is crossing the threshold (10%).
Summary	namespace: {{$labels.namespace}}, timestamp: {{ with query "time()" }} {{ . | first | value | humanizeTimestamp }}{{ end }}
Severity	Error
Condition	Failure rate of delete QoS subscriptions requests is above 10 percent of total requests.
Metric Used	ocnef_qos_af_resp_total
Recommended Actions	The alert is cleared when the failure rate of subscription requests is below the failure threshold. Steps: 1. Check for pod logs on Kibana for ERROR WARN logs. 2. In case the issue persists, capture all the outputs for the above steps and contact My Oracle Support.

5.2.33 QOSNotificationFailureRateCrossedThreshold

Table 5-63 QOSNotificationFailureRateCrossedThreshold

Field	Details
Description	Failure rate of QoS notifications is crossing the threshold (10%).
Summary	namespace: {{$labels.namespace}}, timestamp: {{ with query "time()" }} {{ . | first | value | humanizeTimestamp }}{{ end }}
Severity	Error
Condition	Failure rate of QoS notifications requests is above 10 percent of total requests.
Metric Used	ocnef_qos_5g_resp_total

Table 5-63 (Cont.) QOSNotificationFailureRateCrossedThreshold

Field	Details
Recommended Actions	The alert is cleared when the failure rate of notification requests is below the failure threshold. Steps: 1. Check for pod logs on Kibana for ERROR WARN logs. 2. In case the issue persists, capture all the outputs for the above steps and contact My Oracle Support.

6
CAPIF Alerts

This chapter includes information about the following CAPIF alerts:

- System Level Alerts
- Application Level Alerts

> **Note**
>
> - The performance and capacity of the CAPIF system may vary based on the call model, feature or interface configuration, and underlying CNE and hardware environment.
> - Due to unavailability of metric and/or MQL queries, the following alerts are not supported for OCI:
>   - OccapifNfStatusUnavailable
>   - OccapifPodsRestart
>   - OccapifEgressGatewayServiceDown
>   - OccapifIngressGatewayServiceDown
>   - OccapifAfManagerServiceDown
>   - OccapifAPIManagerServiceDown
>   - OccapifEventManagerServiceDown

6.1 System Level Alerts

This section lists the system level alerts for CAPIF.

6.1.1 OccapifNfStatusUnavailable

Table 6-1 OccapifNfStatusUnavailable

Field	Details
Description	CAPIF services unavailable'
Summary	"namespace: {{$labels.namespace}}, timestamp: {{ with query \"time() \" }}{{ . | first | value | humanizeTimestamp }}{{ end }} : All OCCAPIF services are unavailable."
Severity	Critical
Condition	All the CAPIF services are unavailable.

Table 6-1 (Cont.) OccapifNfStatusUnavailable

Field	Details
Metric Used	'up' **Note**: This is a Prometheus metric used for instance availability monitoring. If this metric is not available, use a similar metric as exposed by the monitoring system.

Recommended Actions	The alert is cleared automatically when the CAPIF services restart.
	Steps:
	1. Check for service-specific alerts which may be causing the issues with service exposure.
	2. Run the following command to check the pod status:
	`$ kubectl get po -n <namespace>`
	a. Run the following command to analyze the error condition of the pod that is not in the *Running* state:
	`$ kubectl describe pod <pod name not in Running state> -n <namespace>`
	Where <pod name not in Running state> indicates the pod that is not in the *Running* state.
	3. Refer to the application logs on Kibana and check for database related failures such as connectivity and invalid secrets. The logs can be filtered based on the services.
	4. Check for helm status to make sure there are no errors:
	`$ helm status <helm release name of the desired NF> -n <namespace>`
	If it is not in “STATUS : DEPLOYED”, then capture logs and event again.
	5. In case the issue persists, capture all the outputs for the above steps and contact My Oracle Support.
	Note: Use CNC NF Data Collector tool for capturing logs. For more information on the Data Collector tool, see *Oracle Communications Cloud Native Core, Network Function Data Collector User Guide*.

6.1.2 OccapifPodsRestart

Table 6-2 OccapifPodsRestart

Field	Details
Description	'Pod *<Pod Name>* has restarted.

Table 6-2 (Cont.) OccapifPodsRestart

Field	Details
Summary	"namespace: {{$labels.namespace}}, podname: {{$labels.pod}}, timestamp: {{ with query \"time()\" }}{{ . | first | value | humanizeTimestamp }}{{ end }} : A Pod has restarted"
Severity	Major
Condition	A pod belonging to any of the CAPIF services has restarted.
Metric Used	kube_pod_container_status_restarts_total

Recommended Actions	The alert is cleared automatically if the specific pod is up.
	Steps:
	1. Refer to the application logs on Kibana and filter based on pod name, check for database related failures such as connectivity and Kubernetes secrets.
	2. To check the orchestration logs for liveness or readiness probe failures, do the following:
	a. Run the following command to check the pod status:
	`$ kubectl get po -n <namespace>`
	b. Run the following command to analyze the error condition of the pod that is not in the *Running* state:
	`$ kubectl describe pod <pod name not in Running state> -n <namespace>`
	Where <pod name not in Running state> indicates the pod that is not in the *Running* state.
	3. Check the database status. For more information, see *Oracle Communications Cloud Native Core, cnDBTier User Guide*.
	4. In case the issue persists, capture all the outputs for the above steps and contact My Oracle Support.
	Note: Use CNC NF Data Collector tool for capturing logs. For more information on the Data Collector tool, see *Oracle Communications Cloud Native Core, Network Function Data Collector User Guide*.

6.1.3 OccapifTotalExternalIngressTrafficRateAboveMinorThreshold

Table 6-3 OccapifTotalExternalIngressTrafficRateAboveMinorThreshold

Field	Details
Description	"OCCAPIF External Ingress traffic rate is above the configured minor threshold i.e. 800 TPS (current value is: {{ $value }})"
Summary	"timestamp: {{ with query \"time()\" }}{{ . | first | value | humanizeTimestamp }}{{ end }}: Traffic rate is above 80 percent of max TPS (1000)"
Severity	Minor

Table 6-3 (Cont.) OccapifTotalExternalIngressTrafficRateAboveMinorThreshold

Field	Details
Condition	The total CAPIF External Ingress traffic rate has crossed the configured minor threshold of 800 TPS. Default value of this alert trigger point in *Occapif Alert.yaml* is 80 % of 1000 (Maximum ingress request rate).
OID	1.3.6.1.4.1.323.5.3.39.1.3.5003
Metric Used	oc_ingressgateway_http_requests_total
Recommended Actions	The alert is cleared either when the External Ingress traffic rate goes above the minor threshold. **Note**: The threshold is configurable in the *Occapif Alert.yaml* alert file. Reassess why the CAPIF is receiving additional traffic. If this alert is unexpected, contact My Oracle Support. **Steps**: 1. Refer Grafana to determine which service is receiving high traffic. 2. Refer **Ingress gateway** section in Grafana to determine the increase in 4xx and 5xx error codes. 3. Check **Ingress gateway** logs on Kibana to determine the reason for the errors.

6.1.4 OccapifTotalNetworkIngressTrafficRateAboveMinorThreshold

Table 6-4 OccapifTotalNetworkIngressTrafficRateAboveMinorThreshold

Field	Details
Description	"OCCAPIF Network Ingress traffic rate is above the configured minor threshold i.e. 800 TPS (current value is: {{ $value }})"
Summary	"timestamp: {{ with query \"time()\" }}{{ . | first | value | humanizeTimestamp }}{{ end }}: Traffic rate is above 80 percent of max TPS (1000)"
Severity	Minor
Condition	The total CAPIF Network Ingress traffic rate has crossed the configured minor threshold of 800 TPS. Default value of this alert trigger point in *Occapif Alert.yaml* is 80% of 1000 (maximum ingress request rate).
OID	1.3.6.1.4.1.323.5.3.39.1.3.5004
Metric Used	oc_ingressgateway_http_requests_total

Table 6-4 (Cont.) OccapifTotalNetworkIngressTrafficRateAboveMinorThreshold

Field	Details
Recommended Actions	The alert is cleared either when the Network Ingress traffic rate goes above the minor threshold. **Note**: The threshold is configurable in the *Occapif Alert.yaml* alert file. Reassess why the CAPIF is receiving additional traffic. If this alert is unexpected, contact My Oracle Support. **Steps**: 1. Refer Grafana to determine which service is receiving high traffic. 2. Refer **Ingress gateway** section in Grafana to determine the increase in 4xx and 5xx error codes. 3. Check **Ingress gateway** logs on Kibana to determine the reason for the errors.

6.1.5 OccapifTotalExternalIngressTrafficRateAboveMajorThreshold

Table 6-5 OccapifTotalExternalIngressTrafficRateAboveMajorThreshold

Field	Details
Description	"OCCAPIF External Ingress traffic rate is above the configured major threshold i.e. 900 TPS (current value is: {{ $value }})"
Summary	"timestamp: {{ with query \"time()\" }}{{ . | first | value | humanizeTimestamp }}{{ end }}: Traffic rate is above 90 percent of max TPS (1000)"
Severity	Major
Condition	The total CAPIF External Ingress traffic rate has crossed the configured major threshold of 900 TPS. Default value of this alert trigger point in *Occapif Alert.yaml* is 90 % of 1000 (maximum ingress request rate).
OID	1.3.6.1.4.1.323.5.3.39.1.3.5005
Metric Used	oc_ingressgateway_http_requests_total
Recommended Actions	The alert is cleared either when the External Ingress traffic rate goes above the major threshold. **Note**: The threshold is configurable in the *Occapif Alert.yaml* alert file. Reassess why the CAPIF is receiving additional traffic. If this alert is unexpected, contact My Oracle Support. **Steps**: 1. Refer Grafana to determine which service is receiving high traffic. 2. Refer **Ingress gateway** section in Grafana to determine the increase in 4xx and 5xx error codes. 3. Check **Ingress gateway** logs on Kibana to determine the reason for the errors.

6.1.6 OccapifTotalNetworkIngressTrafficRateAboveMajorThreshold

Table 6-6 OccapifTotalNetworkIngressTrafficRateAboveMajorThreshold

Field	Details
Description	"OCCAPIF Network Ingress traffic rate is above the configured major threshold i.e. 900 TPS (current value is: {{ $value }})"
Summary	"timestamp: {{ with query \"time()\" }}{{ . | first | value | humanizeTimestamp }}{{ end }}: Traffic Rate is above 90 percent of max TPS (1000)"
Severity	Major
Condition	The total CAPIF Network Ingress traffic rate has crossed the configured major threshold of 900 TPS. Default value of this alert trigger point in *Occapif Alert.yaml* is 90 % of 1000 (maximum ingress request rate).
OID	1.3.6.1.4.1.323.5.3.39.1.3.5006
Metric Used	oc_ingressgateway_http_requests_total
Recommended Actions	The alert is cleared either when the Network Ingress traffic rate goes above the major threshold. **Note**: The threshold is configurable in the *Occapif Alert.yaml* alert file. Reassess why the CAPIF is receiving additional traffic. If this alert is unexpected, contact My Oracle Support. **Steps**: 1. Refer Grafana to determine which service is receiving high traffic. 2. Refer **Ingress gateway** section in Grafana to determine the increase in 4xx and 5xx error codes. 3. Check **Ingress gateway** logs on Kibana to determine the reason for the errors.

6.1.7 OccapifTotalExternalIngressTrafficRateAboveCriticalThreshold

Table 6-7 OccapifTotalExternalIngressTrafficRateAboveCriticalThreshold

Field	Details
Description	"OCCAPIF External Ingress traffic rate is above the configured critical threshold i.e. 950 TPS (current value is: {{ $value }})"
Summary	"timestamp: {{ with query \"time()\" }}{{ . | first | value | humanizeTimestamp }}{{ end }}: Traffic rate is above 95 percent of max TPS (1000)"
Severity	Critical
Condition	The total CAPIF External Ingress traffic rate has crossed the configured critical threshold of 950 TPS. Default value of this alert trigger point in *Occapif Alert.yaml* is 95 % of 1000 (maximum ingress request rate).
OID	1.3.6.1.4.1.323.5.3.39.1.3.5007
Metric Used	oc_ingressgateway_http_requests_total

Table 6-7 (Cont.) OccapifTotalExternalIngressTrafficRateAboveCriticalThreshold

Field	Details
Recommended Actions	The alert is cleared either when the External Ingress traffic rate goes above the critical threshold. **Note**: The threshold is configurable in the *Occapif Alert.yaml* alert file. Reassess why the CAPIF is receiving additional traffic. If this alert is unexpected, contact My Oracle Support. **Steps**: 1. Refer Grafana to determine which service is receiving high traffic. 2. Refer **Ingress gateway** section in Grafana to determine the increase in 4xx and 5xx error codes. 3. Check **Ingress gateway** logs on Kibana to determine the reason for the errors.

6.1.8 OccapifTotalNetworkIngressTrafficRateAboveCriticalThreshold

Table 6-8 OccapifTotalNetworkIngressTrafficRateAboveCriticalThreshold

Field	Details
Description	"OCCAPIF Network Ingress traffic rate is above the configured critical threshold i.e. 950 TPS (current value is: {{ $value }})"
Summary	"timestamp: {{ with query \"time()\" }}{{ . | first | value | humanizeTimestamp }}{{ end }}: Traffic rate is above 95 percent of max TPS (1000)"
Severity	Critical
Condition	The total CAPIF Network Ingress traffic rate has crossed the configured critical threshold of 950 TPS. Default value of this alert trigger point in *Occapif Alert.yaml* is 95 % of 1000 (Maximum ingress request rate).
OID	1.3.6.1.4.1.323.5.3.39.1.3.5008
Metric Used	oc_ingressgateway_http_requests_total
Recommended Actions	The alert is cleared either when the Network Ingress traffic rate goes above the critical threshold. **Note**: The threshold is configurable in the *Occapif Alert.yaml* alert file. Reassess why the CAPIF is receiving additional traffic. If this alert is unexpected, contact My Oracle Support. **Steps**: 1. Refer Grafana to determine which service is receiving high traffic. 2. Refer **Ingress gateway** section in Grafana to determine the increase in 4xx and 5xx error codes. 3. Check **Ingress gateway** logs on Kibana to determine the reason for the errors.

6.1.9 OccapifExternalIngressTransactionErrorRateAboveZeroPointOnePercent

Table 6-9 OccapifExternalIngressTransactionErrorRateAboveZeroPointOnePercent

Field	Details
Description	"OCCAPIF External Ingress transaction error rate is above 0.1 percent of total transactions (current value is {{ $value }})"
Summary	"timestamp: {{ with query \"time()\" }}{{ . | first | value | humanizeTimestamp }}{{ end }}: Transaction error rate detected above 0.1 percent of total transactions"
Severity	Warning
Condition	The number of failed External Ingress transactions is above 0.1 percent of the total transactions.
OID	1.3.6.1.4.1.323.5.3.39.1.3.5009
Metric Used	oc_ingressgateway_http_responses_total
Recommended Actions	The alert is cleared when the number of failure External Ingress transactions is below 0.1 percent of the total transactions or when the number of failed transactions crosses the 1% threshold, in which case the OccapifExternalIngressTransactionErrorRateAbove1Percent is raised. **Steps**: 1. Check the service specific metrics to understand the specific service request errors. 2. Check metrics per service, per method: 3. If guidance is required, contact My Oracle Support.

6.1.10 OccapifNetworkIngressTransactionErrorRateAboveZeroPointOnePercent

Table 6-10 OccapifNetworkIngressTransactionErrorRateAboveZeroPointOnePercent

Field	Details
Description	"OCCAPIF Network Ingress transaction error rate is above 0.1 percent of total transactions (current value is {{ $value }})"
Summary	"timestamp: {{ with query \"time()\" }}{{ . | first | value | humanizeTimestamp }}{{ end }}: Transaction error rate detected above 0.1 percent of total transactions"
Severity	Warning
Condition	The number of failed Network Ingress transactions is above 0.1 percent of the total transactions.
OID	1.3.6.1.4.1.323.5.3.39.1.3.5010
Metric Used	oc_ingressgateway_http_responses_total

Table 6-10 (Cont.) OccapifNetworkIngressTransactionErrorRateAboveZeroPointOnePercent

Field	Details
Recommended Actions	The alert is cleared when the number of failure Network Ingress transactions is below 0.1 percent of the total transactions or when the number of failed transactions crosses the 1% threshold, in which case the OccapifNetworkIngressTransactionErrorRateAbove1Percent is raised. **Steps**: 1. Check the service specific metrics to understand the specific service request errors. 2. Check metrics per service, per method: 3. If guidance is required, contact My Oracle Support.

6.1.11 OccapifExternalIngressTransactionErrorRateAbove1Percent

Table 6-11 OccapifExternalIngressTransactionErrorRateAbove1Percent

Field	Details
Description	"OCCAPIF External Ingress transaction error rate is above 1 percent of total transactions (current value is {{ $value }})"
Summary	"timestamp: {{ with query \"time()\" }}{{ . | first | value | humanizeTimestamp }}{{ end }}: Transaction error rate detected above 1 percent of total transactions"
Severity	Warning
Condition	The number of failed External Ingress transactions is above 1 percent of the total transactions.
OID	1.3.6.1.4.1.323.5.3.39.1.3.5011
Metric Used	oc_ingressgateway_http_responses_total
Recommended Actions	The alert is cleared when the number of failure External Ingress transactions is below 1 percent of the total transactions. **Steps**: 1. Check the service specific metrics to understand the specific service request errors. 2. Check metrics per service, per method: 3. If guidance is required, contact My Oracle Support.

6.1.12 OccapifNetworkIngressTransactionErrorRateAbove1Percent

Table 6-12 OccapifNetworkIngressTransactionErrorRateAbove1Percent

Field	Details
Description	"OCCAPIF Network Ingress transaction error rate is above 1 percent of total transactions (current value is {{ $value }})"

Table 6-12 (Cont.) OccapifNetworkIngressTransactionErrorRateAbove1Percent

Field	Details
Summary	"timestamp: {{ with query \"time()\" }}{{ . | first | value | humanizeTimestamp }}{{ end }}: Transaction error rate detected above 1 percent of total transactions"
Severity	Warning
Condition	The number of failed Network Ingress transactions is above 1 percent of the total transactions.
OID	1.3.6.1.4.1.323.5.3.39.1.3.5012
Metric Used	oc_ingressgateway_http_responses_total
Recommended Actions	The alert is cleared when the number of failure Network Ingress transactions is below 1 percent of the total transactions. **Steps**: 1. Check the service specific metrics to understand the specific service request errors. 2. Check metrics per service, per method: 3. If guidance is required, contact My Oracle Support.

6.1.13 OccapifExternalIngressTransactionErrorRateAbove10Percent

Table 6-13 OccapifExternalIngressTransactionErrorRateAbove10Percent

Field	Details
Description	"OCCAPIF External Ingress transaction error rate is above 10 percent of total transactions (current value is {{ $value }})"
Summary	"timestamp: {{ with query \"time()\" }}{{ . | first | value | humanizeTimestamp }}{{ end }}: Transaction error rate detected above 10 percent of total transactions"
Severity	Minor
Condition	The number of failed External Ingress transactions is above 10 percent of the total transactions.
OID	1.3.6.1.4.1.323.5.3.39.1.3.5013
Metric Used	oc_ingressgateway_http_responses_total
Recommended Actions	The alert is cleared when the number of failure External Ingress transactions is below 10 percent of the total transactions. **Steps**: 1. Check the service specific metrics to understand the specific service request errors. 2. Check metrics per service, per method: 3. If guidance is required, contact My Oracle Support.

6.1.14 OccapifNetworkIngressTransactionErrorRateAbove10Percent

Table 6-14 OccapifNetworkIngressTransactionErrorRateAbove10Percent

Field	Details
Description	"OCCAPIF Network Ingress transaction error rate is above 10 percent of total transactions (current value is {{ $value }})"
Summary	"timestamp: {{ with query \"time()\" }}{{ . | first | value | humanizeTimestamp }}{{ end }}: Transaction error rate detected above 10 percent of total transactions"
Severity	Minor
Condition	The number of failed Network Ingress transactions is above 10 percent of the total transactions.
OID	1.3.6.1.4.1.323.5.3.39.1.3.5014
Metric Used	oc_ingressgateway_http_responses_total
Recommended Actions	The alert is cleared when the number of failure Network Ingress transactions is below 10 percent of the total transactions. **Steps**: 1. Check the service specific metrics to understand the specific service request errors. 2. Check metrics per service, per method: 3. If guidance is required, contact My Oracle Support.

6.1.15 OccapifExternalIngressTransactionErrorRateAbove25Percent

Table 6-15 OccapifExternalIngressTransactionErrorRateAbove25Percent

Field	Details
Description	"OCCAPIF External Ingress transaction error rate detected above 25 percent of total transactions (current value is {{ $value }})"
Summary	"timestamp: {{ with query \"time()\" }}{{ . | first | value | humanizeTimestamp }}{{ end }}: Transaction error rate detected above 25 percent of total transactions"
Severity	Major
Condition	The number of failed External Ingress transactions is above 25 percent of the total transactions.
OID	1.3.6.1.4.1.323.5.3.39.1.3.5015
Metric Used	oc_ingressgateway_http_responses_total
Recommended Actions	The alert is cleared when the number of failure External Ingress transactions is below 25 percent of the total transactions. **Steps**: 1. Check the service specific metrics to understand the specific service request errors. 2. Check metrics per service, per method: 3. If guidance is required, contact My Oracle Support.

6.1.16 OccapifNetworkIngressTransactionErrorRateAbove25Percent

Table 6-16 OccapifNetworkIngressTransactionErrorRateAbove25Percent

Field	Details
Description	"OCCAPIF Network Ingress transaction error rate detected above 25 percent of total transactions (current value is {{ $value }})"
Summary	"timestamp: {{ with query \"time()\" }}{{ . | first | value | humanizeTimestamp }}{{ end }}: Transaction error rate detected above 25 percent of total transactions"
Severity	Major
Condition	The number of failed Network Ingress transactions is above 25 percent of the total transactions.
OID	1.3.6.1.4.1.323.5.3.39.1.3.5016
Metric Used	oc_ingressgateway_http_responses_total
Recommended Actions	The alert is cleared when the number of failure Network Ingress transactions is below 25 percent of the total transactions. **Steps**: 1. Check the service specific metrics to understand the specific service request errors. 2. Check metrics per service, per method: 3. If guidance is required, contact My Oracle Support.

6.1.17 OccapifExternalIngressTransactionErrorRateAbove50Percent

Table 6-17 OccapifExternalIngressTransactionErrorRateAbove50Percent

Field	Details
Description	"OCCAPIF External Ingress transaction error rate detected above 50 percent of total transactions (current value is {{ $value }})"
Summary	"timestamp: {{ with query \"time()\" }}{{ . | first | value | humanizeTimestamp }}{{ end }}: Transaction error rate detected above 50 percent of total transactions"
Severity	Critical
Condition	The number of failed External Ingress transactions is above 50 percent of the total transactions.
OID	1.3.6.1.4.1.323.5.3.39.1.3.5017
Metric Used	oc_ingressgateway_http_responses_total
Recommended Actions	The alert is cleared when the number of failure External Ingress transactions is below 50 percent of the total transactions. **Steps**: 1. Check the service specific metrics to understand the specific service request errors. 2. Check metrics per service, per method: 3. If guidance is required, contact My Oracle Support.

6.1.18 OccapifNetworkIngressTransactionErrorRateAbove50Percent

Table 6-18 OccapifNetworkIngressTransactionErrorRateAbove50Percent

Field	Details
Description	"OCCAPIF Network Ingress transaction error rate detected above 50 percent of total transactions (current value is {{ $value }})"
Summary	"timestamp: {{ with query \"time()\" }}{{ . | first | value | humanizeTimestamp }}{{ end }}: Transaction error rate detected above 50 percent of total transactions"
Severity	Critical
Condition	The number of failed Network Ingress transactions is above 50 percent of the total transactions.
OID	1.3.6.1.4.1.323.5.3.39.1.3.5018
Metric Used	oc_ingressgateway_http_responses_total
Recommended Actions	The alert is cleared when the number of failure Network Ingress transactions is below 50 percent of the total transactions. **Steps**: 1. Check the service specific metrics to understand the specific service request errors. 2. Check metrics per service, per method: 3. If guidance is required, contact My Oracle Support.

6.1.19 OccapifEgressGatewayServiceDown

Table 6-19 OccapifEgressGatewayServiceDown

Field	Details
Description	"CAPIF Egress-Gateway service {{$labels.app_kubernetes_io_name}} is down"
Summary	"kubernetes_namespace: {{$labels.kubernetes_namespace}}, podname: {{$labels.kubernetes_pod_name}}, timestamp: {{ with query \"time()\" }}{{ . | first | value | humanizeTimestamp }}{{ end }} : Egress-Gateway service down"
Severity	Critical
Condition	None of the pods of the Egress Gateway microservice is available.
Metric Used	'up' **Note**: This is a Prometheus metric used for instance availability monitoring. If this metric is not available, use a similar metric as exposed by the monitoring system.

Table 6-19 (Cont.) OccapifEgressGatewayServiceDown

Field	Details

Recommended Actions	The alert is cleared when the Egress Gateway service is available.
	Note: The threshold is configurable in the NefAlertrules alert file.
	Steps:
	1. To check the orchestration logs of Egress Gateway service and check for liveness or readiness probe failures, do the following:
	a. Run the following command to check the pod status:
	`$ kubectl get po -n <namespace>`
	b. Run the following command to analyze the error condition of the pod that is not in the *Running* state:
	`$ kubectl describe pod <pod name not in Running state> -n <namespace>`
	Where <pod name not in Running state> indicates the pod that is not in the *Running* state.
	2. Refer to the application logs on Kibana and filter based on Egress Gateway service names. Check for ERROR WARNING logs related to thread exceptions.
	3. Depending on the failure reason, take the resolution steps.
	4. In case the issue persists, capture all the outputs for the above steps and contact My Oracle Support.
	Note: Use CNC NF Data Collector tool for capturing logs. For more information on how to collect logs using Data Collector tool, see *Oracle Communications Cloud Native Core, Network Function Data Collector User Guide*.

6.1.20 OccapifMemoryUsageCrossedMinorThreshold

Table 6-20 OccapifMemoryUsageCrossedMinorThreshold

Field	Details
Description	"CAPIF Memory Usage for pod {{ $labels.pod }} has crossed the configured minor threshold (50%) (value={{ $value }}) of its limit."
Summary	"namespace: {{$labels.namespace}}, podname: {{$labels.pod}}, timestamp: {{ with query \"time()\" }}{{ . | first | value | humanizeTimestamp }}{{ end }} : Memory Usage of pod exceeded 50% of its limit."
Severity	Minor
Condition	A pod has reached the configured minor threshold (50%) of its memory resource limits.
Metric Used	'container_memory_usage_bytes''container_spec_memory_limit_bytes' **Note**: This is a Kubernetes metric used for instance availability monitoring. If the metric is not available, use the similar metric as exposed by the monitoring system.

Table 6-20 (Cont.) OccapifMemoryUsageCrossedMinorThreshold

Field	Details
Recommended Actions	The alert gets cleared when the memory utilization falls below the Minor Threshold or crosses the major threshold, in which case OccapifMemoryUsageCrossedMajorThreshold alert is raised. **Note**: The threshold is configurable in the NefAlertrules alert file. In case the issue persists, capture all the outputs for the above steps and contact My Oracle Support. **Note**: Use CNC NF Data Collector tool for capturing logs. For more information on how to collect logs using Data Collector tool, see *Oracle Communications Cloud Native Core, Network Function Data Collector User Guide*.

6.1.21 OccapifMemoryUsageCrossedMajorThreshold

Table 6-21 OccapifMemoryUsageCrossedMajorThreshold

Field	Details
Description	"CAPIF Memory Usage for pod {{ $labels.pod }} has crossed the configured major threshold (60%) (value = {{ $value }}) of its limit."
Summary	"namespace: {{$labels.namespace}}, podname: {{$labels.pod}}, timestamp: {{ with query \"time()\" }}{{ . | first | value | humanizeTimestamp }}{{ end }} : Memory Usage of pod exceeded 60% of its limit."
Severity	Major
Condition	A pod has reached the configured major threshold (60%) of its memory resource limits.
OID	1.3.6.1.4.1.323.5.3.39.1.3.5021
Metric Used	'container_memory_usage_bytes' 'container_spec_memory_limit_bytes' **Note**: This is a Kubernetes metric used for instance availability monitoring. If the metric is not available, use the similar metric as exposed by the monitoring system.
Recommended Actions	The alert gets cleared when the memory utilization falls below the Major Threshold or crosses the critical threshold, in which case OccapifMemoryUsageCrossedCriticalThreshold alert is raised. **Note**: The threshold is configurable in the NefAlertrules alert file. In case the issue persists, capture all the outputs for the above steps and contact My Oracle Support. **Note**: Use CNC NF Data Collector tool for capturing logs. For more information on how to collect logs using Data Collector tool, see *Oracle Communications Cloud Native Core, Network Function Data Collector User Guide*.

6.1.22 OccapifMemoryUsageCrossedCriticalThreshold

Table 6-22 OccapifMemoryUsageCrossedCriticalThreshold

Field	Details
Description	"CAPIF Memory Usage for pod {{ $labels.pod }} has crossed the configured major threshold (70%) (value = {{ $value }}) of its limit."
Summary	"namespace: {{$labels.namespace}}, podname: {{$labels.pod}}, timestamp: {{ with query \"time()\" }}{{ . | first | value | humanizeTimestamp }}{{ end }} : Memory Usage of pod exceeded 70% of its limit."
Severity	Critical
Condition	A pod has reached the configured critical threshold (70%) of its memory resource limits.
Metric Used	'container_memory_usage_bytes' 'container_spec_memory_limit_bytes' **Note**: This is a Kubernetes metric used for instance availability monitoring. If the metric is not available, use a similar metric as exposed by the monitoring system.
Recommended Actions	The alert gets cleared when the memory utilization falls below the Critical threshold. **Note**: The threshold is configurable in the NefAlertrules alert file. In case the issue persists, capture all the outputs for the above steps and contact My Oracle Support. **Note**: Use CNC NF Data Collector tool for capturing logs. For more information on how to collect logs using Data Collector tool, see *Oracle Communications Cloud Native Core, Network Function Data Collector User Guide*.

6.1.23 OccapifIngressGatewayServiceDown

Table 6-23 OccapifIngressGatewayServiceDown

Field	Details
Description	"CAPIF Ingress-Gateway service {{$labels.app_kubernetes_io_name}} is down"
Summary	"kubernetes_namespace: {{$labels.kubernetes_namespace}}, timestamp: {{ with query \"time()\" }}{{ . | first | value | humanizeTimestamp }}{{ end }} : Ingress-gateway service down"
Severity	Critical
Condition	None of the pods of the Ingress-Gateway microservice is available.
Metric Used	'up' **Note**: This is a Prometheus metric used for instance availability monitoring. If this metric is not available, use a similar metric as exposed by the monitoring system.

Table 6-23 (Cont.) OccapifIngressGatewayServiceDown

Field	Details

Recommended Actions	The alert is cleared when the Ingress Gateway service is available.
	Steps:
	1. To check the orchestration logs of Ingress Gateway service and check for liveness or readiness probe failures, do the following:
	a. Run the following command to check the pod status:
	`$ kubectl get po -n <namespace>`
	b. Run the following command to analyze the error condition of the pod that is not in the *Running* state:
	`$ kubectl describe pod <pod name not in Running state> -n <namespace>`
	Where <pod name not in Running state> indicates the pod that is not in the *Running* state.
	2. Refer to the application logs on Kibana and filter based on Ingress Gateway service names. Check for ERROR WARNING logs related to thread exceptions.
	3. Depending on the failure reason, take the resolution steps.
	4. In case the issue persists, capture all the outputs for the above steps and contact My Oracle Support.
	Note: Use CNC NF Data Collector tool for capturing logs. For more information on how to collect logs using Data Collector tool, see *Oracle Communications Cloud Native Core, Network Function Data Collector User Guide*.

6.1.24 OccapifApiManagerServiceDown

Table 6-24 OccapifApiManagerServiceDown

Field	Details
Description	"CAPIF API Manager service {{$labels.app_kubernetes_io_name}} is down"
Summary	"namespace: {{$labels.namespace}}, podname: {{$labels.kubernetes_pod_name}}, timestamp: {{ with query \"time()\" }} {{ . | first | value | humanizeTimestamp }}{{ end }} : AF Manager service down"
Severity	Critical
Condition	The API Manager service is down.
Metric Used	'up' **Note:** This is a Prometheus metric used for instance availability monitoring. If this metric is not available, use a similar metric as exposed by the monitoring system.

Table 6-24 (Cont.) OccapifApiManagerServiceDown

Field	Details

Recommended Actions	The alert is cleared when the CAPIF API Manager service is available.
	Steps:
	1. To check the orchestration logs of `occapif_apimgr` service and check for liveness or readiness probe failures, do the following:
	a. Run the following command to check the pod status:
	`$ kubectl get pod -n <namespace>`
	b. Run the following command to analyze the error condition of the pod that is not in the **Running** state:
	`$ kubectl describe pod <pod name not in Running state> -n <namespace>`
	Where <pod name not in Running state> indicates the pod that is not in the Running state.
	2. Refer the application logs on Kibana and filter based on `occapif_apimgr` service names. Check for ERROR WARNING logs related to thread exceptions.
	3. Check the DB status. For more information on how to check the DB status, see Oracle Communications Cloud Native Core, cnDBTier User Guide.
	4. Depending on the failure reason, take the resolution steps.
	5. In case the issue persists, capture all the outputs for the above steps and contact My Oracle Support.
	Note: Use CNC NF Data Collector tool for capturing logs. For more information on how to collect logs using Data Collector tool, see Oracle Communications Cloud Native Core, Network Function Data Collector User Guide.

6.1.25 OcapifAfManagerServiceDown

Table 6-25 OcapifAfManagerServiceDown

Field	Details
Description	"CAPIF AF Manager service {{$labels.app_kubernetes_io_name}} is down"
Summary	"kubernetes_namespace: {{$labels.kubernetes_namespace}}, timestamp: {{ with query \"time()\" }}{{ . | first | value | humanizeTimestamp }}{{ end }} : AF Manager service down"
Severity	Critical
Condition	The AF Manager service is down.
Metric Used	'up' **Note:** This is a Prometheus metric used for instance availability monitoring. If this metric is not available, use a similar metric as exposed by the monitoring system.

Table 6-25 (Cont.) OcapifAfManagerServiceDown

Field	Details

Recommended Actions	The alert is cleared when the CAPIF AF Manager service is available.
	Steps:
	1. To check the orchestration logs of `occapif_afmgr` service and check for liveness or readiness probe failures, do the following:
	a. Run the following command to check the pod status:
	`$ kubectl get pod -n <namespace>`
	b. Run the following command to analyze the error condition of the pod that is not in the **Running** state:
	`$ kubectl describe pod <pod name not in Running state> -n <namespace>`
	Where <pod name not in Running state> indicates the pod that is not in the Running state.
	2. Refer the application logs on Kibana and filter based on `occapif_afmgr` service names. Check for ERROR WARNING logs related to thread exceptions.
	3. Check the DB status. For more information on how to check the DB status, see Oracle Communications Cloud Native Core, cnDBTier User Guide.
	4. Depending on the failure reason, take the resolution steps.
	5. In case the issue persists, capture all the outputs for the above steps and contact My Oracle Support.
	Note: Use CNC NF Data Collector tool for capturing logs. For more information on how to collect logs using Data Collector tool, see Oracle Communications Cloud Native Core, Network Function Data Collector User Guide.

6.1.26 OccapifEventManagerServiceDown

Table 6-26 OccapifEventManagerServiceDown

Field	Details
Description	"CAPIF API Manager service {{$labels.app_kubernetes_io_name}} is down"
Summary	"kubernetes_namespace: {{$labels.kubernetes_namespace}}, timestamp: {{ with query \"time()\" }}{{ . | first | value | humanizeTimestamp }}{{ end }} : API Manager service down"
Severity	Critical
Condition	The Event Manager service is down.
Metric Used	'up' **Note:** This is a Prometheus metric used for instance availability monitoring. If this metric is not available, use a similar metric as exposed by the monitoring system.

Table 6-26 (Cont.) OccapifEventManagerServiceDown

Field	Details

Recommended Actions	The alert is cleared when the CAPIF Event Manager service is available.
	Steps:
	1. To check the orchestration logs of `occapif_eventmanager` service and check for liveness or readiness probe failures, do the following:
	a. Run the following command to check the pod status:
	`$ kubectl get pod -n <namespace>`
	b. Run the following command to analyze the error condition of the pod that is not in the **Running** state:
	`$ kubectl describe pod <pod name not in Running state> -n <namespace>`
	Where <pod name not in Running state> indicates the pod that is not in the Running state.
	2. Refer the application logs on Kibana and filter based on `ocnef_expgw_apimgr` service names. Check for ERROR WARNING logs related to thread exceptions.
	3. Check the DB status. For more information on how to check the DB status, see Oracle Communications Cloud Native Core, cnDBTier User Guide.
	4. Depending on the failure reason, take the resolution steps.
	5. In case the issue persists, capture all the outputs for the above steps and contact My Oracle Support.
	Note: Use CNC NF Data Collector tool for capturing logs. For more information on how to collect logs using Data Collector tool, see Oracle Communications Cloud Native Core, Network Function Data Collector User Guide.

6.2 Application Level Alerts

This section lists the application level alerts for CAPIF.

6.2.1 AfMgrOnboardingOauthValidationFailureRateCrossedThreshold

Table 6-27 AfMgrOnboardingOauthValidationFailureRateCrossedThreshold

Field	Details
Description	"Failure Rate of AI Onboarding Oauth Validation Is Crossing the Threshold (10%)"
Summary	"namespace: {{$labels.namespace}}, timestamp: {{ with query \"time() \" }}{{ . | first | value | humanizeTimestamp }}{{ end }} : Failure Rate Of Onboarding is above 10 percent of total requests."

Table 6-27 (Cont.) AfMgrOnboardingOauthValidationFailureRateCrossedThreshold

Field	Details
Severity	Error
Condition	The failure rate of API Invoker onboarding is reaching the threshold value.
Metric Used	occapif_afmgr_resp_total
Recommended Actions	The alert is cleared when the failure rate of API invoker onboarding is below the threshold. Steps: 1. Check for pod logs on Kibana for ERROR WARN logs. 2. In case the issue persists, capture all the outputs for the above steps and contact My Oracle Support.

7
HTTP Error Codes

NEF generates HTTP error codes in case of any request failure over Service Based Interface (SBI). The error descriptions contains the possible reason of the failure to facilitate proper troubleshooting.

The error description is provided in JSON format for the HTTP (POST/GET/DELETE) requests. The error code format is defined as follows:

OCNEF-<*Service Name*>-<*Error Description*>-<*HTTP Response Code*>-<*Problem Code*>

where,

- <*Service Name*>: Specifies the name of the service that generates the error code.
- <*Error Description*>: The reason or cause for the error.
- <*HTTP Response Code*>: The HTTP Status Code for the error.
- <*Problem Code*>: NEF defined 4 digit code for the error.

The following table describes each error code to help you understand why the error occurred and to provide a recovery procedure to help correct the condition that caused the error:

Table 7-1 HTTP Error Codes Supported on Service Based Interface (SBI)

Error Code	Description	Recovery Procedure
HTTP Response Code: 400 BAD_REQUEST		
`OCNEF-<svc_name>-<error>-400-1000`	The value is not supported the for field: <*fieldName*>.	Verify the parameter specifications for the request and provide a supported value to retry the operation.
`OCNEF-<service name>-<error>-400-1001`	The mandatory field is missing from the request body: <*fieldName*>.	Verify the request specifications and provide the missing mandatory value to retry the operation.
`OCNEF-<svc_name>-<error>-400-1002`	Validation failed for the field with validation details: <*fieldName*>	Update the field value as per the validation error and retry the operation.
`OCNEF-<svc_name>-<error>-400-1003`	Could not convert request JSON into valid object	Verify the JSON format of the request and retry the operation.
`OCNEF-<svc_name>-<error>-400-1004`	Internal header parameter missing: <*Parameter name*>	Contact My Oracle Support.
`OCNEF-<svc_name>-<error>-400-1005`	Duplicate Error : <*Field details*>	Provide a unique value and retry the operation.
`OCNEF-<svc_name>-<error>-400-1006`	Message translation error	Contact My Oracle Support in case the issue persists.
`OCNEF-<svc_name>-<error>-400-1007`	Communication profile not found	Contact My Oracle Support in case the issue persists.
`OCNEF-<svc_name>-<error>-400-1008`	Empty Security method	Contact My Oracle Support in case the issue persists.
HTTP Response Code: 404 NOT_FOUND		

Table 7-1 (Cont.) HTTP Error Codes Supported on Service Based Interface (SBI)

Error Code	Description	Recovery Procedure
OCNEF-<svc_name>-<error>-404-1501	The record not available in OCNEF. Possible errors: • Subscription not found • Invoker not found • Group Id not found	Verify the relevant value in the system. Contact My Oracle Support in case the issue persists.
HTTP Response Code: 401 UNAUTHORIZED		
OCNEF-<svc_name>-<error>-401-1601	Possible errors: • Invoker onboarding related errors • Invalid Token • Invalid Invoker name • Mismatch between global capif configuration parameter "capifApiPrefix" and global nef configuration parameter "capifDetails.apiPrefix"	Verify the relevant values in the system.
500 INTERNAL_SERVER_ERROR		
OCNEF-<svc_name>-<error>-500-1701	Connectivity issue between NEF and UDM/GMLC	Verify the connectivity between OCNEF and UDM/GMLC. Contact My Oracle Support in case the issue persists.
OCNEF-<svc_name>-<error>-500-1702	Connectivity issue between Monitoring Event (ME) and 5GCAgent services	Verify the connectivity between the ME service and 5GC Agent service. Contact My Oracle Support in case the issue persists.
OCNEF-<svc_name>-<error>-500-1703	Connectivity issue between EG and ME services	Verify the connectivity between the EG and ME service. Contact My Oracle Support in case the issue persists.
OCNEF-<svc_name>-<error>-500-1704	Model D JSON to Object conversion error	Contact My Oracle Support in case the issue persists.
OCNEF-<svc_name>-<error>-500-1705	ME Subscription Service JSON to object conversion error	Contact My Oracle Support in case the issue persists.
OCNEF-<svc_name>-<error>-500-1706	NEF to Core connectivity failure using NRF Discovery	Contact My Oracle Support in case the issue persists.
OCNEF-<svc_name>-<error>-500-1707	Invalid communication model configuration	Contact My Oracle Support in case the issue persists.
OCNEF-<svc_name>-<error>-500-1708	Error during the following conversions: • QOS Service JSON to object conversion • Object to JSON conversion	Contact My Oracle Support in case the issue persists.
HTTP Response Code: 503 SERVICE_UNAVAILABLE		
OCNEF-<svc_name>-<error>-503-1801	Connectivity issue between NEF and the 5G core network	Check the connectivity. Contact My Oracle Support in case the issue persists.
OCNEF-<svc_name>-<error>-503-1802	Connectivity issue between EG and ME services	Verify the connectivity between the EG and ME service. Contact My Oracle Support in case the issue persists.

Table 7-1 (Cont.) HTTP Error Codes Supported on Service Based Interface (SBI)

Error Code	Description	Recovery Procedure
HTTP Response Code: 501 NOT_IMPLEMENTED		
OCNEF-<svc_name>-<error>-501-1901	Method not supported : <*Method Name*>	The requested HTTP Method not supported in the current release.
Diameter Error Code: 503 SERVICE_UNAVAILABLE		
DIAMETER_INVALID_AVP_VALUE-5004	If scef-reference-id passed from MME is not found at ME end.	Verify if scef-reference-id passed from MME is the correct one belonging to ME subscription.
DIAMETER_UNABLE_TO_COMPLY-5012	If there is Code Exception or Database is down or there is a Network Issue.	Contact My Oracle Support in case the issue persists.