

Oracle® Communications

Cloud Native Core, Binding Support Function

Troubleshooting Guide



Release 26.1.201
G56171-01
August 2026

ORACLE®

Copyright © 2021, 2026, Oracle and/or its affiliates.

This software and related documentation are provided under a license agreement containing restrictions on use and disclosure and are protected by intellectual property laws. Except as expressly permitted in your license agreement or allowed by law, you may not use, copy, reproduce, translate, broadcast, modify, license, transmit, distribute, exhibit, perform, publish, or display any part, in any form, or by any means. Reverse engineering, disassembly, or decompilation of this software, unless required by law for interoperability, is prohibited.

The information contained herein is subject to change without notice and is not warranted to be error-free. If you find any errors, please report them to us in writing.

If this is software, software documentation, data (as defined in the Federal Acquisition Regulation), or related documentation that is delivered to the U.S. Government or anyone licensing it on behalf of the U.S. Government, then the following notice is applicable:

U.S. GOVERNMENT END USERS: Oracle programs (including any operating system, integrated software, any programs embedded, installed, or activated on delivered hardware, and modifications of such programs) and Oracle computer documentation or other Oracle data delivered to or accessed by U.S. Government end users are "commercial computer software," "commercial computer software documentation," or "limited rights data" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, the use, reproduction, duplication, release, display, disclosure, modification, preparation of derivative works, and/or adaptation of i) Oracle programs (including any operating system, integrated software, any programs embedded, installed, or activated on delivered hardware, and modifications of such programs), ii) Oracle computer documentation and/or iii) other Oracle data, is subject to the rights and limitations specified in the license contained in the applicable contract. The terms governing the U.S. Government's use of Oracle cloud services are defined by the applicable contract for such services. No other rights are granted to the U.S. Government.

This software or hardware is developed for general use in a variety of information management applications. It is not developed or intended for use in any inherently dangerous applications, including applications that may create a risk of personal injury. If you use this software or hardware in dangerous applications, then you shall be responsible to take all appropriate fail-safe, backup, redundancy, and other measures to ensure its safe use. Oracle Corporation and its affiliates disclaim any liability for any damages caused by use of this software or hardware in dangerous applications.

Oracle®, Java, MySQL, and NetSuite are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

Intel and Intel Inside are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. AMD, Epyc, and the AMD logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark of The Open Group.

This software or hardware and documentation may provide access to or information about content, products, and services from third parties. Oracle Corporation and its affiliates are not responsible for and expressly disclaim all warranties of any kind with respect to third-party content, products, and services unless otherwise set forth in an applicable agreement between you and Oracle. Oracle Corporation and its affiliates will not be responsible for any loss, costs, or damages incurred due to your access to or use of third-party content, products, or services, except as set forth in an applicable agreement between you and Oracle.

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

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

Contents

1	Introduction	
1.1	Overview	1
1.2	References	1
2	Logs	
2.1	Log Levels	1
2.2	Understanding Logs	2
3	Using Debug Tool	
3.1	Configurable Parameters for Debug Tool	14
4	Troubleshooting BSF	
4.1	Deployment Issues	1
4.1.1	Helm Install Failure	1
4.1.2	Loss of Session Binding data	2
4.2	Startup Probes	3
4.3	Upgrade or Rollback Failure	4
4.4	Bulk Import and Export Issues	5
5	BSF Alerts	
5.1	List of Alerts	1
5.1.1	AAA_RX_FAIL_COUNT_EXCEEDS_CRITICAL_THRESHOLD	2
5.1.2	AAA_RX_FAIL_COUNT_EXCEEDS_MAJOR_THRESHOLD	2
5.1.3	AAA_RX_FAIL_COUNT_EXCEEDS_MINOR_THRESHOLD	3
5.1.4	SCP_PEER_UNAVAILABLE	3
5.1.5	SCP_PEER_SET_UNAVAILABLE	4
5.1.6	BSF_SERVICES_DOWN	4
5.1.7	BSF_TRAFFIC_RATE_ABOVE_MINOR_THRESHOLD	5
5.1.8	BSF_TRAFFIC_RATE_ABOVE_MAJOR_THRESHOLD	6
5.1.9	BSF_TRAFFIC_RATE_ABOVE_CRITICAL_THRESHOLD	7

5.1.10	BINDING_QUERY_RESPONSE_ERROR_MINOR	7
5.1.11	BINDING_QUERY_RESPONSE_ERROR_MAJOR	8
5.1.12	BINDING_QUERY_RESPONSE_ERROR_CRITICAL	8
5.1.13	DIAM_RESPONSE_NETWORK_ERROR_MINOR	8
5.1.14	DIAM_RESPONSE_NETWORK_ERROR_MAJOR	9
5.1.15	DIAM_RESPONSE_NETWORK_ERROR_CRITICAL	10
5.1.16	DUPLICATE_BINDING_REQUEST_ERROR_MINOR	11
5.1.17	DUPLICATE_BINDING_REQUEST_ERROR_MAJOR	11
5.1.18	DUPLICATE_BINDING_REQUEST_ERROR_CRITICAL	12
5.1.19	INGRESS_TOTAL_ERROR_RATE_ABOVE_MINOR_THRESHOLD	12
5.1.20	INGRESS_TOTAL_ERROR_RATE_ABOVE_MAJOR_THRESHOLD	13
5.1.21	INGRESS_TOTAL_ERROR_RATE_ABOVE_CRITICAL_THRESHOLD	13
5.1.22	PCF_BINDING_ERROR_RATE_ABOVE_MINOR_THRESHOLD	14
5.1.23	PCF_BINDING_ERROR_RATE_ABOVE_MAJOR_THRESHOLD	14
5.1.25	INGRESS_CREATE_ERROR_RATE_ABOVE_MINOR_THRESHOLD	15
5.1.26	INGRESS_CREATE_ERROR_RATE_ABOVE_CRITICAL_THRESHOLD	15
5.1.27	INGRESS_CREATE_ERROR_RATE_ABOVE_MAJOR_THRESHOLD	16
5.1.28	INGRESS_DELETE_ERROR_RATE_ABOVE_MINOR_THRESHOLD	17
5.1.29	INGRESS_DELETE_ERROR_RATE_ABOVE_MAJOR_THRESHOLD	17
5.1.30	INGRESS_DELETE_ERROR_RATE_ABOVE_CRITICAL_THRESHOLD	18
5.1.31	DB_TIER_DOWN_ALERT	19
5.1.32	CPU_USAGE_PER_SERVICE_ABOVE_MINOR_THRESHOLD	19
5.1.33	CPU_USAGE_PER_SERVICE_ABOVE_MAJOR_THRESHOLD	20
5.1.34	CPU_USAGE_PER_SERVICE_ABOVE_CRITICAL_THRESHOLD	21
5.1.35	MEMORY_USAGE_PER_SERVICE_ABOVE_MINOR_THRESHOLD	22
5.1.36	MEMORY_USAGE_PER_SERVICE_ABOVE_MAJOR_THRESHOLD	23
5.1.37	MEMORY_USAGE_PER_SERVICE_ABOVE_CRITICAL_THRESHOLD	24
5.1.38	NRF_COMMUNICATION_FAILURE	24
5.1.39	NRF_SERVICE_REQUEST_FAILURE	25
5.1.40	PERF_INFO_ACTIVE_OVERLOAD_THRESHOLD_FETCH_FAILED	25
5.1.41	POD_DOC	26
5.1.42	POD_CONGESTED	27
5.1.43	POD_CONGESTION_L1	28
5.1.44	POD_CPU_CONGESTION_L1	29
5.1.45	POD_CONGESTION_L2	29
5.1.46	POD_CPU_CONGESTION_L2	30
5.1.47	POD_PENDING_REQUEST_DOC	30
5.1.48	POD_PENDING_REQUEST_CONGESTED	31
5.1.49	POD_CPU_DOC	32
5.1.50	POD_CPU_CONGESTED	33
5.1.51	POD_MEMORY_DOC	33
5.1.52	POD_MEMORY_CONGESTED	34

5.1.53	SERVICE_OVERLOADED	35
5.1.54	SERVICE_RESOURCE_OVERLOADED	38
5.1.55	SYSTEM_IMPAIRMENT_MAJOR	46
5.1.56	SYSTEM_IMPAIRMENT_CRITICAL	48
5.1.57	SYSTEM_OPERATIONAL_STATE_PARTIAL_SHUTDOWN	50
5.1.58	SYSTEM_OPERATIONAL_STATE_COMPLETE_SHUTDOWN	51
5.1.59	DIAM_CONN_PEER_DOWN	52
5.1.60	DIAM_CONN_NETWORK_DOWN	53
5.1.61	DIAM_RESPONSE_REALM_VALIDATION_ERROR_CRITICAL	54
5.1.62	DIAM_RESPONSE_REALM_VALIDATION_ERROR_MAJOR	55
5.1.63	DIAM_RESPONSE_REALM_VALIDATION_ERROR_MINOR	56
5.1.64	AUDIT_STALE_NOTIFY_ERROR_RESPONSE_MINOR	56
5.1.65	AUDIT_STALE_NOTIFY_ERROR_RESPONSE_MAJOR	57
5.1.66	AUDIT_STALE_NOTIFY_ERROR_RESPONSE_CRITICAL	58
5.1.67	BSF_CONNECTION_FAILURE	58
5.1.68	INGRESS_GATEWAY_DD_UNREACHABLE_MAJOR	59
5.1.69	EGRESS_GATEWAY_DD_UNREACHABLE_MAJOR	59
5.1.70	DIAM_GATEWAY_CERTIFICATE_EXPIRY_MAJOR	60
5.1.71	DIAM_GATEWAY_CERTIFICATE_EXPIRY_CRITICAL	61
5.1.72	DGW_TLS_CONNECTION_FAILURE	61
5.1.73	BINDING_REVALIDATION_PCF_BINDING_MISSING_MINOR	62
5.1.74	BINDING_REVALIDATION_PCF_BINDING_MISSING_MAJOR	63
5.1.75	BINDING_REVALIDATION_PCF_BINDING_MISSING_CRITICAL	63
5.1.76	BSF_STATE_NON_FUNCTIONAL_CRITICAL	64
5.1.77	POD_PENDING_REQUEST_CONGESTION_L1	64
5.1.78	POD_PENDING_REQUEST_CONGESTION_L2	65
5.1.79	AUDIT_NOT_RUNNING	66
5.1.80	BSF_PCF_BINDING_TABLE_MIGRATED_PERCENTAGE	67
5.1.81	BSF_PCF_BINDING_TABLE_MIGRATION_INVALID_CONFIGURATION	67
5.1.82	CERTIFICATE_EXPIRY_MINOR	68
5.1.83	CERTIFICATE_EXPIRY_MAJOR	69
5.1.84	CERTIFICATE_EXPIRY_CRITICAL	69
5.1.85	AUDIT_STALE_NOTIFY_404_RESPONSE_CRITICAL	70
5.1.86	AUDIT_STALE_NOTIFY_404_RESPONSE_MAJOR	70
5.1.87	AUDIT_STALE_NOTIFY_404_RESPONSE_MINOR	71
5.1.88	RAA_RX_FAIL_COUNT_EXCEEDS_CRITICAL_THRESHOLD	72
5.1.89	RAA_RX_FAIL_COUNT_EXCEEDS_MAJOR_THRESHOLD	72
5.1.90	RAA_RX_FAIL_COUNT_EXCEEDS_MINOR_THRESHOLD	73
5.1.91	STA_RX_FAIL_COUNT_EXCEEDS_CRITICAL_THRESHOLD	73
5.1.92	STA_RX_FAIL_COUNT_EXCEEDS_MAJOR_THRESHOLD	74
5.1.93	STA_RX_FAIL_COUNT_EXCEEDS_MINOR_THRESHOLD	75
5.1.94	ASA_RX_FAIL_COUNT_EXCEEDS_CRITICAL_THRESHOLD	76

5.1.95	ASA_RX_FAIL_COUNT_EXCEEDS_MAJOR_THRESHOLD	77
5.1.96	ASA_RX_FAIL_COUNT_EXCEEDS_MINOR_THRESHOLD	77
5.1.97	EGRESS_GATEWAY_DNS_RESOLUTION_FAILURE	78
5.1.98	DIAMETER_POD_ERROR_RESPONSE_CRITICAL	79
5.1.99	DIAMETER_POD_ERROR_RESPONSE_MAJOR	79
5.1.100	DIAMETER_POD_ERROR_RESPONSE_MINOR	80
5.1.101	DIAM_GATEWAY_CERTIFICATE_EXPIRY_MINOR	80

Preface

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

Documentation Accessibility

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

Access to Oracle Support

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

Diversity and Inclusion

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

Conventions

The following text conventions are used in this document:

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

My Oracle Support

My Oracle Support (<https://support.oracle.com>) is your initial point of contact for all product support and training needs.

Use My Oracle Support to access support resources, create and manage Service Requests (SRs), and get assistance with product issues.

For help with My Oracle Support registration, Service Request creation, or other support access questions, use the available support options in My Oracle Support, including MOSFS and MOS Chat.

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	Definition
3GPP	3rd Generation Partnership Project
AF	Application Function
AMF	Access and Mobility Management Function
ASM	Aspen Service Mesh
BSF	Oracle Communications Cloud Native Core, Binding Support Function
CA	Certificate Authority
CCA	Client Credentials Assertion
CDS	Cache data Service. Caching Microservice which is responsible for syncing and caching the data
CHF	Charging Function
CM	Configuration Management
CMPv2	Certificate Management Protocol Version 2
CNC	Cloud Native Core
CNC Console	Oracle Communications Cloud Native Configuration Console
CNE	Oracle Communication Cloud Native Core, Cloud Native Environment
CNPCRF	Oracle Communications Cloud Native Core, Policy and Charging Rules Function
CUSTOMER_REPO	Docker registry address including the port number, if the docker registry has an associated port.
cnDBTier	Oracle Communications Cloud Native Core, cnDBTier
GUAMI	Globally Unique AMF Identifier
HTTPS	Hypertext Transfer Protocol Secure
IMAGE_TAG	Image tag from release tar file. You can use any tag number. However, make sure that you use that specific tag number while pushing docker image to the docker registry.
IMS	IP Multimedia Subsystem
JOSE	Javascript Object Signing and Encryption
JSON	JavaScript Object Notation
JWT	JSON Web Token
LCM	Life Cycle Management
MCC	Mobile Country Code
MCPTT	Mission-critical push-to-talk
METALLB_ADDRESS_POOL	Address pool configured on metallb to provide external IPs
MNC	Mobile Network Code
NEF	Oracle Communications Cloud Native Core, Network Exposure Function
NPLI	Network Provided Location Information
NRF	Oracle Communications Cloud Native Core, Network Repository Function
NSI	Network Slice Instances
NWDAF	Oracle Communications Networks Data Analytics Function

Table (Cont.) Acronyms and Terminologies

Acronym	Definition
OCI	Oracle Cloud Infrastructure
OCCM	Oracle Communications Cloud Native Core, Certificate Management
OCNADD	Oracle Communications Networks Analytics Data Director
PDB	Pod Disruption Budget
PCF	Oracle Communications Cloud Native Core, Policy Control Function
PCRF	Oracle Communications Cloud Native Core, Policy and Charging Rules Function
PCEF	Policy and Charging Enforcement Function
PDS	Policy Data Service
PRA	Presence Reporting Area
PRE	Policy Runtime Engine
PDU	Protocol Data Unit
PVC	Persistent Volume Chain
Policy	Oracle Communications Cloud Native Core, Converged Policy
QoS	Quality of Service
RAN	Radio Access Network
SAN	Subject Alternate Name
SNSSAI	Single Network Slice Selection Assistance Information
SMF	Session Management Function
UDR	Oracle Communications Cloud Native Core, Unified Data Repository
SRA	Successful Resource Allocation
STR	Session Termination Request (Rx)
UE	User Equipment
UPF	User Plane Function
UPSI	UE Policy Section Identifier
URSP	UE Route Selection Policies
Local NRF set state data	NF instances and NF subscription data that is stored in the local cnDBTier of an NRF instance. It is the state data from the local cnDBTier and replicated state data from the georedundant NRF.
Remote NRF set state data	NF registrations and NF subscriptions data that is stored in remote NRF set instances within the same segment. It is the state data from peer NRF set in the same segment
Periodic syncing	It refreshes the local cache periodically.
State Data	The state data is the dynamic data that contains the NFInstances and NFSubscriptions of the NFs.

What's New In This Guide

This section introduces the documentation updates for release 26.1.2xx.

Release 26.1.201 - G56171-01, August 2026

- Added the following alerts in the [List of Alerts](#) section:
 - [AUDIT_STALE_NOTIFY_404_RESPONSE_CRITICAL](#)
 - [AUDIT_STALE_NOTIFY_404_RESPONSE_MAJOR](#)
 - [AUDIT_STALE_NOTIFY_404_RESPONSE_MINOR](#)
 - [RAA_RX_FAIL_COUNT_EXCEEDS_CRITICAL_THRESHOLD](#)
 - [RAA_RX_FAIL_COUNT_EXCEEDS_MAJOR_THRESHOLD](#)
 - [RAA_RX_FAIL_COUNT_EXCEEDS_MINOR_THRESHOLD](#)
 - [STA_RX_FAIL_COUNT_EXCEEDS_CRITICAL_THRESHOLD](#)
 - [STA_RX_FAIL_COUNT_EXCEEDS_MAJOR_THRESHOLD](#)
 - [STA_RX_FAIL_COUNT_EXCEEDS_MINOR_THRESHOLD](#)
 - [ASA_RX_FAIL_COUNT_EXCEEDS_CRITICAL_THRESHOLD](#)
 - [ASA_RX_FAIL_COUNT_EXCEEDS_MAJOR_THRESHOLD](#)
 - [ASA_RX_FAIL_COUNT_EXCEEDS_MINOR_THRESHOLD](#)
 - [EGRESS_GATEWAY_DNS_RESOLUTION_FAILURE](#)
- Updated the following alerts:
 - [BSF_PCF_BINDING_TABLE_MIGRATED_PERCENTAGE](#)
 - [BSF_PCF_BINDING_TABLE_MIGRATION_INVALID_CONFIGURATION](#)
 - [AUDIT_STALE_NOTIFY_ERROR_RESPONSE_MINOR](#)
 - [AUDIT_STALE_NOTIFY_ERROR_RESPONSE_MAJOR](#)
 - [AUDIT_STALE_NOTIFY_ERROR_RESPONSE_CRITICAL](#)

1

Introduction

This document provides information about troubleshooting Oracle Communications Cloud Native Core, Binding Support Function (BSF) services and managed objects.

1.1 Overview

The Binding Support Function Troubleshooting Guide provides extensive information about resolving problems you might experience while installing and configuring the network function. This document also provides information about tools available to help you collect and analyze diagnostic data.

This Troubleshooting Guide describes in detail common problems that may arise while installing, configuring, and using BSF. After a user has identified the issue, perform the provided steps to resolve the issue.

1.2 References

Refer to the following documents while deploying BSF:

- *Oracle Communications Cloud Native Binding Support Function Installation, Upgrade, and Fault Recovery Guide*
- *Oracle Communications Cloud Native Core, Binding Support Function REST Specification Guide*
- *Oracle Communications Cloud Native Core, Binding Support Function Network Impact Report Guide*
- *Oracle Communications Cloud Native Core, Binding Support Function Troubleshooting Guide*
- *Oracle Communications Cloud Native Core, Cloud Native Environment Installation and Upgrade Guide*
- *Oracle Communications Cloud Native Core, cnDBTier User Guide*
- *Oracle Communications Cloud Native Core, Data Collector User Guide*
- *Oracle Communications Cloud Native Core Automated Test Suite Guide*
- *Oracle Communications Cloud Native Core Release Notes*
- *Oracle Communications Cloud Native Core Solution Upgrade Guide*

2

Logs

Log files are used to register system events, together with their date and time of occurrence. They can be valuable tools for troubleshooting. Not only do logs indicate that specific events occurred, they also provide important clues about a chain of events that led to an error or problem.

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 kubectl commands, see Kubernetes [website](#).

2.1 Log Levels

This section provides information on log levels supported by Oracle Communications Cloud Native Core, Binding Support Function (BSF).

A log level helps in defining the severity level of a log message. 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.

As shown in the following image, only log messages with level defined as WARN are shown, after adding filter:



The screenshot shows the Kibana log viewer interface. At the top, there is a filter bar with 'level: WARN' selected and a '+ Add filter' button. Below the filter bar, a table displays the filtered log entries. The table has three columns: 'Time', 'level', and 'kubernetes.container_name'. Two log entries are visible, both with a 'WARN' level and from the 'diam-gateway' container.

Time	level	kubernetes.container_name
> Aug 15, 2021 @ 12:13:56.354	WARN	diam-gateway
> Aug 15, 2021 @ 12:13:56.353	WARN	diam-gateway

Supported Log Levels

For BSF, the log level for a micro-service can be set to any of the following valid values:

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

Configuring Log Levels

To view logging configurations and update logging levels, use the Logging Level page under **Logging Configurations** on the CNC Console. For more information, see the section "Log Level" in *Oracle Communications Cloud Native Core, Binding Support Function User Guide*.

2.2 Understanding Logs

This section provides information on how to read logs for various services of BSF in Kibana.

Understanding Logs

The following is a sample log for BSF services:

```
{
  "instant": {
    "epochSecond": 1627016656,
    "nanoOfSecond": 137175036
  },
  "thread": "Thread-2",
  "level": "INFO",
  "loggerName": "ocbsf.framework.domain.orchestration.AbstractProcess",
  "marker": {
    "name": "ALWAYS"
  },
  "message": "Received RECONFIGURE request",
  "endOfBatch": false,
  "loggerFqn": "org.apache.logging.slf4j.Log4jLogger",
  "threadId": 34,
  "threadPriority": 5,
  "messageTimestamp": "2021-07-23T05:04:16.137+0000"
}
```

The log message format is same for all the BSF services.

The following table describes key attributes of a log message:

Table 2-1 Log Attributes

Attribute	Description
level	Log level of the log printed
loggerName	Class/Module which printed the log
message	Message related to the log providing brief details
loggerFqcn	Log4j2 Internal, Fully Qualified class name of logger module
thread	Thread name
threadId	Thread ID generated internally by Log4j2
threadPriority	Thread priority generated internally by Log4j2
messageTimeStamp	Timestamp of log from application container
kubernetes.labels.application	NF Application Name
kubernetes.labels.engineVersion	Engineering version of software
kubernetes.labels.mktgVersion	Marketing version of software
kubernetes.labels.microservice	Name of the microservice
kubernetes.namespace_name	Namespace of BSF deployment
kubernetes.host	worker node name on which container is running
kubernetes.pod_name	Pod Name
kubernetes.container_name	Container Name
Docker.container_id	Process ID internally assigned
kubernetes.labels.vendor	Vendor of product

3

Using Debug Tool

Overview

The Debug Tool provides third-party troubleshooting tools for debugging the runtime issues in the lab environment.

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

Prerequisites

This section describes the prerequisites for using 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

The following configurations must be performed in the Bastion Host.

- a. When BSF is installed on CNE version 23.2.0 or above:

Note

- In CNE version 23.2.0 or above, the default CNE 23.2.0 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

- 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: {}
```

- 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":["ocbsf"]} }]'
```

Sample output:

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

- 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": "ocbsf" }]'
```

Example:

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

Sample output:

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

- 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 BSF is installed on CNE version prior to 23.2.0
PodSecurityPolicy (PSP) Creation

The following configurations must be performed in the Bastion Host.

- i. Log in to the Bastion Host.
- ii. Create a new 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. **Default values** are recommended.

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

apiVersion: bsf/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: ocbsf
rules:
  - apiGroups:
```

```

- policy
resources:
- podsecuritypolicies
verbs:
- use
resourceNames:
- debug-tool-psp
EOF

```

RoleBinding Creation

Run the following command to attach the service account for your NF namespace with the role created for the tool PSP:

```

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

```

Refer to *Debug Tool Configuration Parameters* for parameter details.

2. Configuration in NF specific Helm

Following updates must be performed in custom_values.yaml file.

- a. Log in to the NF server.
- b. Open the custom_values file:

```
$ vim <custom_values file>
```

- c. Under global configuration, add the following:

```

global:
  extraContainers: ENABLED

```

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 "-" }}
```

It ensures that necessary values are added as prefix and suffix to the container name.

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

```
bsf-management-service:
  #extraContainers: DISABLED
  envMySQLDatabase: ocpm_bsf_1.7.0
  resources:
    limits:
      cpu: 1
      memory: 1Gi
    requests:
      cpu: 0.5
      memory: 1Gi
  minReplicas: 1
```

Note

- At the global level, extraContainers flag can be used to enable or 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.

Execution of Debug Tool

Following is the procedure to run Debug Tool.

Run the following command to enter Debug Tool Container:

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

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

Example:

```
$ kubectl get pods -n ocbsf
```

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

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

Tools Tested in Debug Container

Following is the list of debug tools that are tested.

tcpdump

Table 3-1 tcpdump

Options Tested	Description	Output	Capabilities
-D	Print the list of the network interfaces available on the system and on which <i>tcpdump</i> can capture packets.	<pre>tcpdump -D</pre> 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

Table 3-1 (Cont.) tcpdump

Options Tested	Description	Output	Capabilities
-i	Listen on <i>interface</i> .	<pre>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)</pre>	NET_ADMIN, NET_RAW
-w	Write the raw packets to file rather than parsing and printing them out.	<pre>tcpdump -w capture.pcap -i eth0</pre>	NET_ADMIN, NET_RAW
-r	Read packets from <i>file</i> (which was created with the -w option).	<pre>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</pre>	NET_ADMIN, NET_RAW

ip

Table 3-2 ip

Options Tested	Description	Output	Capabilities
addr show	Look at protocol addresses.	<pre>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/loopback 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</pre>	--

Table 3-2 (Cont.) ip

Options Tested	Description	Output	Capabilities
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	--

netstat

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:tpoxy 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.default: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	--
-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:tpoxy 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	--

Table 3-3 (Cont.) netstat

Options Tested	Description	Output	Capabilities
-i	Display a table of all network interfaces.	netstat -i Kernel Interface table face MTU RX-OK RX-ERR RX-DRP RX-OVR TX-OK TX-ERR TX-DRP TX-OVR Flg eth0 1440 4131 0 0 0 4355 0 0 0 BMRUlo 65536 0 0 0 0 0 0 0 LRU	--

jq

Table 3-4 jq

Options Tested	Description	Output	Capabilities
<jq filter> [file...]	Use it to slice and filter and map and transform structured data. Sample JSON file: { "fruit": { "name": "apple", "color": "green", "price": 1.2 } }	jq '.fruit' sample.json { "name": "apple", "color": "green", "price": 1.2 }	--
Sample Json	Sample JSON file: { "fruit": { "name": "apple", "color": "green", "price": 1.2 } }	jq '.fruit.color,.fruit.price' sample.json "green" 1.2	--

curl

Table 3-5 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	--

ping

Table 3-6 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

*nmap***Table 3-7 nmap**

Options Tested	Description	Output	Capabilities
<ip>	Scan for Live hosts, Operating systems, packet filters and open ports running on remote hosts.	<pre>nmap 10.178.254.194 Starting Nmap 6.40 (http://nmap.org) at 2020-09-29 05:54 UTC Nmap 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 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: 1 IP address (1 host up) scanned in 0.04 seconds</pre>	--

Table 3-7 (Cont.) nmap

Options Tested	Description	Output	Capabilities
-v	Increase verbosity level	<pre> 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 </pre>	--

Table 3-7 (Cont.) nmap

Options Tested	Description	Output	Capabilities
-iL	Scan all the listed IP addresses in a file. Sample file	<pre>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</pre>	--

dig

Table 3-8 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.	<pre>dig 10.178.254.194</pre> <i>Note: The IP should be reachable from inside the container.</i>	--
-x	Query DNS Reverse Look-up.	<pre>dig -x 10.178.254.194</pre>	--

3.1 Configurable Parameters for Debug Tool

This section describes the configurable parameters that users can customize to configure debug tool.

CNE Parameters**Table 3-9 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 (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 3-10 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 white list of names that the rule applies to.

Table 3-11 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 3-12 Debug Tool Configuration Parameters**

Parameter	Description
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
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.readOnlyRootFilesystem	Whether this container has a read-only root filesystem. Default is false.

Table 3-12 (Cont.) Debug Tool Configuration Parameters

Parameter	Description
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
securityContext.capabilities.add	Added capabilities
securityContext.runAsUser	The UID to run the entrypoint of the container process.
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.
volumeMounts.mountPath	Indicates the path for volume mount.
volumeMounts.name	Indicates the name of the directory for debug tool logs storage.

4

Troubleshooting BSF

This section provides information on how to troubleshoot the common errors that may occur during the installation and upgrade of Oracle Communications Cloud Native Core, Binding Support Function.

4.1 Deployment Issues

This section provides information on how to troubleshoot deployment issues for BSF.

4.1.1 Helm Install Failure

If `helm install` command Fails

This section describes various scenarios and troubleshooting procedures if the `helm install` command fails.

Reasons for `helm install` failure:

- **Chart syntax issue**
Please resolve the chart specific things and rerun the `helm install` command, because in this case, no hooks should have begun.
- **Most Possible Reason [Timeout]**
If any job is stuck in a pending/error state and not able to execute, it will result in the timeout after 5 minutes as default timeout for helm command is "5 minutes". In this case, follow the below steps to troubleshoot.
- **`helm install` command failed in case of duplicated chart**

```
helm install /home/cloud-user/bsf_1.5.0/sprint3.1/ocbsf-1.5.0-  
sprint.3.1.tgz --name ocbsf2 --namespace ocbsf2 -f cust-ashish.yaml  
Error: release ocbsf2 failed: configmaps "perfinfo-config-ocbsf2" already  
exists
```

Here, configmap `perfinfo-config-ocbsf2` exists multiple times while creating the Kubernetes objects after the pre-upgrade hooks. This is a failure scenario. In this case, perform the following troubleshooting steps.

Troubleshooting steps:

1. Run the following command to cleanup the databases created by the `helm install` command :

```
DROP DATABASE IF EXISTS ocbsf_commonconfig;  
DROP DATABASE IF EXISTS ocbsf_config_server;  
DROP DATABASE IF EXISTS ocpm_bsf;  
DROP DATABASE IF EXISTS ocbsf_release;  
DROP DATABASE IF EXISTS ocbsf_audit_service;
```

2. Run the following command to get kubernetes objects:

```
kubectl get all -n <release_namespace>
```

This gives a detailed overview of which objects are stuck or in a failed state.

3. Run the following command to delete all kubernetes objects:

```
kubectl delete all --all -n <release_namespace>  
kubectl delete cm --all -n <release_namespace>
```

4. Run the following command to list all the kubernetes objects:

```
helm ls --all
```

If this is in a failed state, please purge the namespace using the command

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

Note

If the execution of this command is taking more time, run the below command parallelly in another session to clear all the delete jobs.

```
while true; do kubectl delete jobs --all -n <release_namespace>;  
sleep 5;done
```

Monitor the below command:

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

Once that is succeeded, press "ctrl+c" to stop the above script.

5. After the database cleanup, run the `helm install` command.

4.1.2 Loss of Session Binding data

Symptom

When the NDB Cluster goes down (due to node group being down), it results in loss of session binding data from `pcf_binding` table in BSF.

Problem

NOLOGGING is enabled for the `pcf_binding` table in the `ocpm_bsf` database which means there is non-persistent storage of session table in PV. Hence, when the NDB Cluster goes down there is loss of session binding data.

Resolution Steps

To resolve this issue, perform the following steps:

1. Monitor the alerts and the KPI to determine the status of the BSF cnDBTier and the NF.
2. If there is an issue due to cnDBTier cluster down, perform a controlled shutdown of the BSF in the affected site.
3. Perform a GRR procedure in the cnDBTier to recover the database and replication.
4. Start the BSF again using the controlled shutdown procedure to reinitiate the BSF to accept traffic.

4.2 Startup Probes

To increase the application's reliability and availability, startup probes are introduced in BSF. Consider a scenario where the configuration is not loaded or partially loaded but the service goes into a ready state. This may result in different pods showing different behaviour for the same service. With the introduction of startup probe, the readiness and liveness checks for a pod are not initiated until the configuration is loaded completely and startup probe is successful. However, if the startup probe fails, the container restarts.

To check the status of startup probe or investigate the reason of failing, perform the following steps:

1. Log in to a container by running the following command:

```
kubectl exec -it podname -n namespace -- bash
curl -kv http://localhost:<monitoring-port>/<startup-probe-url>
```

Example:

```
kubectl exec -it test-bsf-797cf5997-2zlgf -- curl -kv http://
localhost:9000/actuator/health/startup
```

The sample output can be as follow:

```
[cloud-user@bastion-1 ~]$
* Trying ::1...
* TCP_NODELAY set
* connect to ::1 port 9000 failed: Connection refused
* Trying 127.0.0.1...
* TCP_NODELAY set
* connect to 127.0.0.1 port 9000 failed: Connection refused
* Failed to connect to localhost port 9000: Connection refused
* Closing connection 0
curl: (7) Failed to connect to localhost port 9000: Connection refused
command terminated with exit code 7
[cloud-user@bastion-1 ~]$ k exec -it test-bsf-797cf5997-2zlgf -- curl -kv
http://localhost:9000/actuator/health/startup
* Trying ::1...
* TCP_NODELAY set
* Connected to localhost (::1) port 9000 (#0)
> GET /actuator/health/startup HTTP/1.1
> Host: localhost:9000
> User-Agent: curl/7.61.1
> Accept: */*
>
< HTTP/1.1 503 Service Unavailable
```

```

< Date: Thu, 21 Apr 2022 11:18:03 GMT
< Content-Type: application/json; charset=utf-8
< Transfer-Encoding: chunked
< Server: Jetty(9.4.43.v20210629)
<
* Connection #0 to host localhost left intact
{"status":"DOWN"}[cloud-user@bastion-1 ~]$ k exec -it test-
bsf-797cf5997-2zlgf -- curl -kv http://localhost:9000/actuator/health/
startup
* Trying ::1...
* TCP_NODELAY set
* Connected to localhost (::1) port 9000 (#0)
> GET /actuator/health/startup HTTP/1.1
> Host: localhost:9000
> User-Agent: curl/7.61.1
> Accept: */*
>
< HTTP/1.1 200 OK
< Date: Thu, 21 Apr 2022 11:18:04 GMT
< Content-Type: application/json; charset=utf-8
< Transfer-Encoding: chunked
< Server: Jetty(9.4.43.v20210629)
<
* Connection #0 to host localhost left intact
{"status":"UP"}[cloud-user@bastion-1 ~]$

```

2. To check why the startup probe failed, describe the output:

Describe output:

```

Warning Unhealthy <invalid> (x10 over 2m45s) kubelet
Startup probe failed: Get "http://10.233.81.231:9000/actuator/health/
startup": dial tcp 10.233.81.231:9000: connect: connection refused

```

The following could be the possible reasons for startup probe failure:

- Network connectivity issue
 - Database connection issue due to which server is not coming up
 - Due to any other exception
3. If the reason for startup probe failure is not clear, check the logs to determine if it is due to an issue with config-server connection or any issue with fetching configurations from the config-server.

4.3 Upgrade or Rollback Failure

When Oracle Communications Cloud Native Core, Binding Support Function (BSF) upgrade or rollback fails, perform the following procedure.

1. Check the pre or post upgrade or rollback hook logs in Kibana as applicable. Users can filter upgrade or rollback logs using the following filters:
 - For upgrade: lifeCycleEvent=9001 or 9011
 - For rollback: lifeCycleEvent=9002

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 is not related to database or network connectivity issue and is observed during the preupgrade phase, do not perform rollback because BSF deployment remains in the source or older release.
 - 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.
4. If the issue persists, contact [My Oracle Support](#).

4.4 Bulk Import and Export Issues

This section describes the most common bulk import and export 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 Oracle Support.

Bulk Import and Export – Partial Success Scenarios

Partial success messages encountered during bulk import/export operations can be ignored in the following scenarios, as they are aligned with recent updates.

Congestion Load Shedding Profile

When importing legacy congestion configuration data into version 25.1.100 or later, partial success can occur if the data contains a legacy congestion load shedding profile.

Reason: This behavior is expected due to migration changes in congestion configuration screens introduced in these versions. Partial success in this case can be ignored.

5

BSF Alerts

This section provides information on Oracle Communications Cloud Native Core, Binding Support Function (BSF) alerts and their configuration.

Note

The performance and capacity of the BSF system may vary based on the call model, Feature/Interface configuration, and underlying CNE and hardware environment.

You can configure alerts in Prometheus and `Alertrules.yaml` file.

The following table describes the various severity types of alerts generated by BSF:

Table 5-1 Alert Levels or Severity Types

Alert Levels / Severity Types	Definition
Critical	Indicates a severe issue that poses a significant risk to safety, security, or operational integrity. It requires immediate response to address the situation and prevent serious consequences. Raised for conditions can affect the service of BSF.
Major	Indicates a more significant issue that has an impact on operations or poses a moderate risk. It requires prompt attention and action to mitigate potential escalation. Raised for conditions can affect the service of BSF.
Minor	Indicates a situation that is low in severity and does not pose an immediate risk to safety, security, or operations. It requires attention but does not demand urgent action. Raised for conditions can affect the service of BSF.
Info or Warn (Informational)	Provides general information or updates that are not related to immediate risks or actions. These alerts are for awareness and do not typically require any specific response. WARN and INFO alerts may not impact the service of BSF.

For more details on how to configure alerts, see *Configuring BSF Alerts* section in *Oracle Communications Cloud Native Core, Binding Support Function Installation, Upgrade, and Fault Recovery Guide*.

5.1 List of Alerts

This section lists the alerts available for Oracle Communications Cloud Native Core, Binding Support Function (BSF).

5.1.1 AAA_RX_FAIL_COUNT_EXCEEDS_CRITICAL_THRESHOLD

Table 5-2 AAA_RX_FAIL_COUNT_EXCEEDS_CRITICAL_THRESHOLD

Field	Details
Description	AAA Rx fail count exceeds the critical threshold limit.
Summary	AAA Rx fail count exceeds the critical threshold limit.
Severity	CRITICAL
Expression	(sum by (namespace) (rate(ocbsf_diam_response_network_total{appld="16777236", msgType="AAA", responseCode!~"2.*",direction="in"}[5m])) / sum by (namespace) (rate(ocbsf_diam_response_network_total{appld="16777236", msgType="AAA", direction="in"}[5m]))) * 100 > 90
OID	1.3.6.1.4.1.323.5.3.37.1.2.40
Metric Used	ocbsf_diam_response_network_total
Recommended Actions	For any additional guidance, contact My Oracle Support.

5.1.2 AAA_RX_FAIL_COUNT_EXCEEDS_MAJOR_THRESHOLD

Table 5-3 AAA_RX_FAIL_COUNT_EXCEEDS_MAJOR_THRESHOLD

Field	Details
Description	AAA Rx fail count exceeds the major threshold limit
Summary	AAA Rx fail count exceeds the major threshold limit.
Severity	MAJOR
Expression	(sum by (namespace) (rate(ocbsf_diam_response_network_total{appld="16777236", msgType="AAA", responseCode!~"2.*",direction="in"}[5m])) / sum by (namespace) (rate(ocbsf_diam_response_network_total{appld="16777236", msgType="AAA", direction="in"}[5m]))) * 100 > 80 <=90
OID	1.3.6.1.4.1.323.5.3.37.1.2.40
Metric Used	ocbsf_diam_response_network_total
Recommended Actions	For any additional guidance, contact My Oracle Support.

5.1.3 AAA_RX_FAIL_COUNT_EXCEEDS_MINOR_THRESHOLD

Table 5-4 AAA_RX_FAIL_COUNT_EXCEEDS_MINOR_THRESHOLD

Field	Details
Description	AAA Rx fail count exceeds the minor threshold limit.
Summary	AAA Rx fail count exceeds the minor threshold limit.
Severity	MINOR
Expression	(sum by (namespace) (rate(ocbsf_diam_response_network_total{appld="16777236", msgType="AAA", responseCode! ~"2.*", direction="in"}[5m])) / sum by (namespace) (rate(ocbsf_diam_response_network_total{appld="16777236", msgType="AAA", direction="in"}[5m]))) * 100 > 60 <=80
OID	1.3.6.1.4.1.323.5.3.37.1.2.40
Metric Used	ocbsf_diam_response_network_total
Recommended Actions	For any additional guidance, contact My Oracle Support.

5.1.4 SCP_PEER_UNAVAILABLE

Table 5-5 SCP_PEER_UNAVAILABLE

Field	Details
Description	Configured SCP peer is unavailable.
Summary	SCP peer [{{\$labels.peer}}] is unavailable.
Severity	Major
Expression	ocbsf_oc_egressgateway_peer_health_status == 1
OID	1.3.6.1.4.1.323.5.3.37.1.2.38
Metric Used	ocbsf_oc_egressgateway_peer_health_status
Recommended Actions	Cause: One or more SCP peer nodes are unavailable or unhealthy. This can affect service routing, request forwarding, and signaling reliability. Diagnostic Information: Identify unhealthy peers and verify network connectivity, health-check endpoints, heartbeat mechanisms, connection failures, timeouts, DNS resolution, IP addresses, ports, certificates, routing rules, load balancer or gateway behavior, TLS handshake failures, and whether multiple peers are affected. Recovery: Restore network connectivity, restart or recover unavailable peer services, correct endpoint, port, certificate, or routing configuration, resolve DNS or service-discovery issues, confirm health checks function correctly, fail over traffic to healthy peers where supported, and scale or redistribute load if peers are overloaded. Monitor until peer health returns to zero.

5.1.5 SCP_PEER_SET_UNAVAILABLE

Table 5-6 SCP_PEER_SET_UNAVAILABLE

Field	Details
Description	None of the SCP peer available for configured peerset.
Summary	{{ \$value }} SCP peers under peer set {{ \$labels.peerset }} are currently unavailable.
Severity	Critical
Expression	(ocbsf_oc_egressgateway_peer_count > 0 and (ocbsf_oc_egressgateway_peer_available_count) == 0)
OID	1.3.6.1.4.1.323.5.3.37.1.2.39
Metric Used	ocbsf_oc_egressgateway_peer_*
Recommended Actions	Cause: One or more SCP peers in a configured peer set are unavailable. This reduces redundancy and can affect load balancing, failover capability, and service reliability. Diagnostic Information: Confirm the total and available peer counts, identify unavailable peers, and verify connectivity, latency, packet loss, health checks, heartbeat status, logs, TLS configuration, IP addresses, ports, routing policies, maintenance activity, scaling events, and crashes. Monitor the load on remaining peers. Recovery: Restore connectivity or restart unavailable peers. Correct endpoint, certificate, or routing configuration; resolve network or DNS issues; re-enable peers that were intentionally disabled; ensure suitable scaling and load distribution; validate health-check mechanisms; and monitor until all peers are healthy and stable.

5.1.6 BSF_SERVICES_DOWN

Table 5-7 BSF_SERVICES_DOWN

Field	Details
Description	{{ \$labels.service }} service is not running!
Summary	{{ \$labels.service }} is not running!
Severity	Critical
Expression	appinfo_service_running{application="ocbsf"} != 1
OID	1.3.6.1.4.1.323.5.3.37.1.2.1
Metric Used	appinfo_service_running

Table 5-7 (Cont.) BSF_SERVICES_DOWN

Field	Details
Recommended Actions	Perform the following steps: - Check for service specific alerts that may be causing the issues with service exposure. - Verify if the POD is in a <i>Running</i> state by using the following command: <pre>kubectl -n <namespace> get pod</pre> If the output shows any pod that is not running, copy the pod name and run the following command: <pre>kubectl describe pod <podname> -n <namespace></pre> - Check the application logs on Kibana and look for database related failures such as connectivity, invalid secrets, and so on. The logs can be easily filtered for different services. - Check for Helm status to ensure no errors are present by using the following command: <pre>helm status <release-name> -n <namespace></pre> If it is not in STATUS: DEPLOYED, capture the logs and events again. In case the issue persists, capture the outputs for the preceding steps and contact My Oracle Support.

5.1.7 BSF_TRAFFIC_RATE_ABOVE_MINOR_THRESHOLD

Table 5-8 BSF_TRAFFIC_RATE_ABOVE_MINOR_THRESHOLD

Field	Details
Description	BSF service Ingress traffic Rate is above threshold of Max MPS(1000) (current value is: {{ \$value }}) The total Binding Management service Ingress traffic rate has crossed the configured threshold of 700 TPS. The default value of this alert trigger point in the <code>BSF_Alertrules.yaml</code> file is when the Binding management service Ingress Rate crosses 70% of maximum ingress requests per second.
Summary	Traffic Rate is above 70 Percent of Max requests per second(1000)
Severity	Minor

Table 5-8 (Cont.) BSF_TRAFFIC_RATE_ABOVE_MINOR_THRESHOLD

Field	Details
Expression	sum(rate(ocbsf_ingress_request_total[2m])) >= 700
OID	1.3.6.1.4.1.323.5.3.37.1.2.2
Metric Used	ocbsf_ingress_request_total
Recommended Actions	The alert gets cleared when the Ingress traffic rate falls below the threshold. Note: Threshold levels can be configured using the <code>BSF_Alertrules.yaml</code> file. It is recommended to assess the reason for additional traffic. Perform the following steps to analyze the cause of increased traffic: 1. Refer Ingress Gateway section in Grafana to determine an increase in 4xx and 5xx error response codes. 2. Check Ingress Gateway logs on Kibana to determine the reason for the errors. For any assistance, contact My Oracle Support.

5.1.8 BSF_TRAFFIC_RATE_ABOVE_MAJOR_THRESHOLD

Table 5-9 BSF_TRAFFIC_RATE_ABOVE_MAJOR_THRESHOLD

Field	Details
Description	BSF service Ingress traffic Rate is above threshold of Max MPS(1000) (current value is: {{ \$value }})
Summary	Traffic Rate is above 80 Percent of Max requests per second(1000)
Severity	Major
Expression	sum(rate(ocbsf_ingress_request_total[2m])) >= 800
OID	1.3.6.1.4.1.323.5.3.37.1.2.2
Metric Used	ocbsf_ingress_request_total
Recommended Actions	The alert gets cleared when the Ingress traffic rate falls below the threshold. Note: Threshold levels can be configured using the <code>BSF_Alertrules.yaml</code> file. It is recommended to assess the reason for additional traffic. Perform the following steps to analyze the cause of increased traffic: 1. Refer Ingress Gateway section in Grafana to determine an increase in 4xx and 5xx error response codes. 2. Check Ingress Gateway logs on Kibana to determine the reason for the errors. For any assistance, contact My Oracle Support.

5.1.9 BSF_TRAFFIC_RATE_ABOVE_CRITICAL_THRESHOLD

Table 5-10 BSF_TRAFFIC_RATE_ABOVE_CRITICAL_THRESHOLD

Field	Details
Description	BSF service Ingress traffic Rate is above threshold of Max MPS(1000) (current value is: {{ \$value }})
Summary	Traffic Rate is above 90 Percent of Max requests per second(1000)
Severity	Critical
Expression	sum(rate(ocbsf_ingress_request_total[2m])) >= 900
OID	1.3.6.1.4.1.323.5.3.37.1.2.2
Metric Used	ocbsf_ingress_request_total
Recommended Actions	The alert gets cleared when the Ingress traffic rate falls below the threshold. Note: Threshold levels can be configured using the <code>BSF_Alertrules.yaml</code> file. It is recommended to assess the reason for additional traffic. Perform the following steps to analyze the cause of increased traffic: 1. Refer Ingress Gateway section in Grafana to determine an increase in 4xx and 5xx error response codes. 2. Check Ingress Gateway logs on Kibana to determine the reason for the errors. For any assistance, contact My Oracle Support.

5.1.10 BINDING_QUERY_RESPONSE_ERROR_MINOR

Table 5-11 BINDING_QUERY_RESPONSE_ERROR_MINOR

Field	Details
Description	At least 30% of the Binding Query connection requests failed.
Summary	At least 30% of the Binding Query requests failed.
Severity	Minor
Expression	(sum(rate(ocbsf_binding_query_response_total{response_code!~"2.*"}[10m]) or (appinfo_service_running * 0)) / sum(rate(ocbsf_binding_query_response_total[10m]))) * 100 >= 30
OID	1.3.6.1.4.1.323.5.3.37.1.2.36
Metric Used	ocbsf_binding_query_response_total.
Recommended Actions	For any assistance, contact My Oracle Support .

5.1.11 BINDING_QUERY_RESPONSE_ERROR_MAJOR

Table 5-12 BINDING_QUERY_RESPONSE_ERROR_MAJOR

Field	Details
Description	At least 50% of the Binding Query connection requests failed.
Summary	At least 50% of the Binding Query requests failed.
Severity	Major
Expression	$(\text{sum}(\text{rate}(\text{ocbsf_binding_query_response_total}\{\text{response_code!}\sim\text{"2.*"}\}[10\text{m}]) \text{ or } (\text{appinfo_service_running} * 0)) / \text{sum}(\text{rate}(\text{ocbsf_binding_query_response_total}[10\text{m}])) * 100 \geq 50$
OID	1.3.6.1.4.1.323.5.3.37.1.2.36
Metric Used	ocbsf_binding_query_response_total.
Recommended Actions	For any assistance, contact My Oracle Support .

5.1.12 BINDING_QUERY_RESPONSE_ERROR_CRITICAL

Table 5-13 BINDING_QUERY_RESPONSE_ERROR_CRITICAL

Field	Details
Description	At least 70% of the Binding Query connection requests failed.
Summary	At least 70% of the Binding Query requests failed.
Severity	Critical
OID	1.3.6.1.4.1.323.5.3.37.1.2.36
Expression	$(\text{sum}(\text{rate}(\text{ocbsf_binding_query_response_total}\{\text{response_code!}\sim\text{"2.*"}\}[10\text{m}]) \text{ or } (\text{appinfo_service_running} * 0)) / \text{sum}(\text{rate}(\text{ocbsf_binding_query_response_total}[10\text{m}])) * 100 \geq 70$
Metric Used	ocbsf_binding_query_response_total.
Recommended Actions	For any assistance, contact My Oracle Support .

5.1.13 DIAM_RESPONSE_NETWORK_ERROR_MINOR

Table 5-14 DIAM_RESPONSE_NETWORK_ERROR_MINOR

Field	Details
Description	At least 20% of the Diam Response connection requests failed with error 'DIAMETER_UNABLE_TO_DELIVER'.
Summary	At least 20% of the Diam Response requests failed with error 'DIAMETER_UNABLE_TO_DELIVER'.
Severity	Minor

Table 5-14 (Cont.) DIAM_RESPONSE_NETWORK_ERROR_MINOR

Field	Details
Expression	$\frac{(\text{sum}(\text{rate}(\text{ocbsf_diam_response_network_total}\{\text{responseCode}="3002"\}[10\text{m}]) \text{ or } (\text{appinfo_service_running} * 0))}{\text{sum}(\text{rate}(\text{ocbsf_diam_response_network_total}[10\text{m}]))} * 100 \geq 20$
OID	1.3.6.1.4.1.323.5.3.37.1.2.35
Metric Used	ocbsf_diam_response_network_total
Recommended Actions	Cause: At least one pod is reporting an elevated rate of Diameter result code 3002, DIAMETER_UNABLE_TO_DELIVER. This typically indicates routing or peer-selection issues, an unreachable destination realm or host, a peer connection outage, an invalid or missing route, or transient network issues. It can also occur during rolling restarts when peer tables or routes are unstable. Diagnostic Information: Identify the worst affected pod and list all affected pods. Compare the 3002 response rate with total Diameter response volume to distinguish a true volume issue from a ratio artifact. Break down failures by available routing labels, such as destination host, destination realm, peer, application ID, or message type. Correlate the condition with peer connection state, disconnects, and recent deployment or restart events for the affected pod. Recovery: If the issue is isolated to one pod, drain or replace it after confirming that there is no broader peer outage, then validate routing tables and peer connections. If it is isolated to a destination realm or host, correct routing configuration, realm definitions, and peer reachability. If peer downtime is responsible, fail over or reroute traffic to alternate peers where supported. Confirm that the 3002 response rate returns to baseline and successful 2xx responses dominate.

5.1.14 DIAM_RESPONSE_NETWORK_ERROR_MAJOR

Table 5-15 DIAM_RESPONSE_NETWORK_ERROR_MAJOR

Field	Details
Description	At least 50% of the Diam Response connection requests failed with error 'DIAMETER_UNABLE_TO_DELIVER'.
Summary	At least 50% of the Diam Response requests failed with error 'DIAMETER_UNABLE_TO_DELIVER'.
Severity	Major

Table 5-15 (Cont.) DIAM_RESPONSE_NETWORK_ERROR_MAJOR

Field	Details
Expression	(sum(rate(ocbsf_diam_response_network_total{responseCode="3002"}[10m]) or (appinfo_service_running * 0)) / sum(rate(ocbsf_diam_response_network_total[10m]))) * 100 >= 50
OID	1.3.6.1.4.1.323.5.3.37.1.2.35
Metric Used	ocbsf_diam_response_network_total
Recommended Actions	Cause: Significant DIAMETER_UNABLE_TO_DELIVER failures are occurring and are likely to affect sessions or transactions. Diagnostic Information: Perform the same checks as for the minor alert and determine whether multiple pods are affected. Confirm the number of pods whose 3002 response ratio is at or above 5%. Recovery: Prioritize restoration of reachability and routing. Fail over peers, correct realm or route configuration, and restart or replace pods if the issue is localized. Treat the condition as an incident if it persists.

5.1.15 DIAM_RESPONSE_NETWORK_ERROR_CRITICAL

Table 5-16 DIAM_RESPONSE_NETWORK_ERROR_CRITICAL

Field	Details
Description	At least 75% of the Diam Response connection requests failed with error 'DIAMETER_UNABLE_TO_DELIVER'.
Summary	At least 75% of the Diam Response requests failed with error 'DIAMETER_UNABLE_TO_DELIVER'.
Severity	Critical
Expression	(sum(rate(ocbsf_diam_response_network_total{responseCode="3002"}[10m]) or (appinfo_service_running * 0)) / sum(rate(ocbsf_diam_response_network_total[10m]))) * 100 >= 75
OID	1.3.6.1.4.1.323.5.3.37.1.2.35
Metric Used	ocbsf_diam_response_network_total

Table 5-16 (Cont.) DIAM_RESPONSE_NETWORK_ERROR_CRITICAL

Field	Details
Recommended Actions	Cause: A severe DIAMETER_UNABLE_TO_DELIVER condition is occurring. This typically indicates a major peer or routing outage, widespread disconnects, or an incorrect configuration rollout. Diagnostic Information: Urgently identify the dominant destination or destinations and determine whether the condition correlates with connection-failure or peer-down alerts. Recovery: Immediately restore routing or peer connectivity, or roll back recent routing or configuration changes. Fail traffic over to healthy peers or realms and replace unhealthy pods if the issue is pod-local. Validate recovery when the 3002 error rate falls well below 10% and end-to-end transaction success recovers.

5.1.16 DUPLICATE_BINDING_REQUEST_ERROR_MINOR

Table 5-17 DUPLICATE_BINDING_REQUEST_ERROR_MINOR

Field	Details
Description	At least 30% of the Binding Registration requests failed were duplicate failures.
Summary	At least 30% of the Binding Registration requests failed were duplicate failures.
Severity	Minor
Expression	$(\text{sum}(\text{rate}(\{_\text{name_}=\sim\text{ocbsf_collision_detection}\}[\text{10m}]) \text{ or } (\text{appinfo_service_running} * 0)) / \text{sum}(\text{rate}(\text{ocbsf_ingress_request_total}\{\text{operation_type}=\text{"register"}\}[\text{10m}])) * 100 \geq 30$
OID	1.3.6.1.4.1.323.5.3.37.1.2.37
Metric Used	ocbsf_ingress_request_total ocbsf_collision_detection
Recommended Actions	For any assistance, contact My Oracle Support .

5.1.17 DUPLICATE_BINDING_REQUEST_ERROR_MAJOR

Table 5-18 DUPLICATE_BINDING_REQUEST_ERROR_MAJOR

Field	Details
Description	At least 50% of the Binding Registration requests failed were duplicate failures.
Summary	At least 50% of the Binding Registration requests failed were duplicate failures.
Severity	Major

Table 5-18 (Cont.) DUPLICATE_BINDING_REQUEST_ERROR_MAJOR

Field	Details
Expression	$(\text{sum}(\text{rate}(\{_\text{name_} \sim \text{"ocbsf_collision_detection."}\}[10\text{m}]) \text{ or } (\text{appinfo_service_running} * 0)) / \text{sum}(\text{rate}(\text{ocbsf_ingress_request_total}\{\text{operation_type} = \text{"register"}\}[10\text{m}])) * 100 \geq 50$
OID	1.3.6.1.4.1.323.5.3.37.1.2.37
Metric Used	ocbsf_ingress_request_total ocbsf_collision_detection
Recommended Actions	For any assistance, contact My Oracle Support .

5.1.18 DUPLICATE_BINDING_REQUEST_ERROR_CRITICAL

Table 5-19 DUPLICATE_BINDING_REQUEST_ERROR_CRITICAL

Field	Details
Description	At least 70% of the Binding Registration requests failed were duplicate failures.
Summary	At least 70% of the Binding Registration requests failed were duplicate failures.
Severity	Critical
Expression	$(\text{sum}(\text{rate}(\{_\text{name_} \sim \text{"ocbsf_collision_detection."}\}[10\text{m}]) \text{ or } (\text{appinfo_service_running} * 0)) / \text{sum}(\text{rate}(\text{ocbsf_ingress_request_total}\{\text{operation_type} = \text{"register"}\}[10\text{m}])) * 100 \geq 70$
OID	1.3.6.1.4.1.323.5.3.37.1.2.37
Metric Used	ocbsf_ingress_request_total ocbsf_collision_detection
Recommended Actions	For any assistance, contact My Oracle Support .

5.1.19 INGRESS_TOTAL_ERROR_RATE_ABOVE_MINOR_THRESHOLD

Table 5-20 INGRESS_TOTAL_ERROR_RATE_ABOVE_MINOR_THRESHOLD

Field	Details
Description	Transaction Error Rate detected above 1 Percent of Total on BSF service (current value is: {{ \$value }})
Summary	Transaction Error Rate detected above 1 Percent of Total Transactions
Severity	Minor
Expression	$(\text{sum}(\text{rate}(\text{ocbsf_ingress_response_total}\{\text{response_code!} \sim \text{"2."}\}[24\text{h}])) / \text{sum}(\text{rate}(\text{ocbsf_ingress_response_total}[24\text{h}])) * 100 \geq 1$
OID	1.3.6.1.4.1.323.5.3.37.1.2.3
Metric Used	ocbsf_ingress_response_total

Table 5-20 (Cont.) INGRESS_TOTAL_ERROR_RATE_ABOVE_MINOR_THRESHOLD

Field	Details
Recommended Actions	The alert gets cleared when the number of failed transactions is below 1% of the total transactions. For any assistance, contact My Oracle Support .

5.1.20 INGRESS_TOTAL_ERROR_RATE_ABOVE_MAJOR_THRESHOLD

Table 5-21 INGRESS_TOTAL_ERROR_RATE_ABOVE_MAJOR_THRESHOLD

Field	Details
Description	Transaction Error Rate detected above 5 Percent of Total on BSF service (current value is: {{ \$value }})
Summary	Transaction Error Rate detected above 5 Percent of Total Transactions
Severity	Major
Expression	$\frac{\text{sum}(\text{rate}(\text{ocbsf_ingress_response_total}\{\text{response_code!~"2.*"}\}[24\text{h}]))}{\text{sum}(\text{rate}(\text{ocbsf_ingress_response_total}[24\text{h}]))} * 100 \geq 5$
OID	1.3.6.1.4.1.323.5.3.37.1.2.3
Metric Used	ocbsf_ingress_response_total
Recommended Actions	The alert gets cleared when the number of failed transactions is below 5% of the total transactions. For any assistance, contact My Oracle Support .

5.1.21

INGRESS_TOTAL_ERROR_RATE_ABOVE_CRITICAL_THRESHOLD

Table 5-22 INGRESS_TOTAL_ERROR_RATE_ABOVE_CRITICAL_THRESHOLD

Field	Details
Description	Transaction Error Rate detected above 10 Percent of Total on BSF service (current value is: {{ \$value }})
Summary	Transaction Error Rate detected above 10 Percent of Total Transactions
Severity	Critical
Expression	$\frac{\text{sum}(\text{rate}(\text{ocbsf_ingress_response_total}\{\text{response_code!~"2.*"}\}[24\text{h}]))}{\text{sum}(\text{rate}(\text{ocbsf_ingress_response_total}[24\text{h}]))} * 100 \geq 10$
OID	1.3.6.1.4.1.323.5.3.37.1.2.3
Metric Used	ocbsf_ingress_response_total
Recommended Actions	The alert gets cleared when the number of failed transactions is below 10% of the total transactions. For any assistance, contact My Oracle Support .

5.1.22 PCF_BINDING_ERROR_RATE_ABOVE_MINOR_THRESHOLD

Table 5-23 PCF_BINDING_ERROR_RATE_ABOVE_MINOR_THRESHOLD

Field	Details
Description	PCF Binding Error Rate above 1 Percent in {{labels.microservice}} in {{labels.namespace}}
Summary	PCF Binding Error Rate in {{labels.kubernetes_node}} (current value is: {{ \$value }})
Severity	Minor
Expression	(sum by (microservice,namespace) (rate(http_server_requests_seconds_count{microservice="bsf-management-service", status!~"200 204",method="GET"}[24h])) / sum by (microservice,namespace) (rate(http_server_requests_seconds_count{microservice="bsf-management-service",method="GET"}[24h]))) * 100 >= 1
OID	1.3.6.1.4.1.323.5.3.37.1.2.5
Metric Used	http_server_requests_seconds_count
Recommended Actions	The alert gets cleared when the number of failed transactions is below 1% of the total transactions. To assess the reason for failed transactions, check the service specific metrics for the GET method. For any assistance, contact My Oracle Support .

5.1.23 PCF_BINDING_ERROR_RATE_ABOVE_MAJOR_THRESHOLD

Table 5-24 PCF_BINDING_ERROR_RATE_ABOVE_MAJOR_THRESHOLD

Field	Details
Description	PCF Binding Error Rate above 5 Percent in {{labels.microservice}} in {{labels.namespace}}
Summary	PCF Binding Error Rate in {{labels.kubernetes_node}} (current value is: {{ \$value }})
Severity	Major
Expression	(sum by (microservice,namespace) (rate(http_server_requests_seconds_count{microservice="bsf-management-service", status!~"200 204",method="GET"}[24h])) / sum by (microservice,namespace) (rate(http_server_requests_seconds_count{microservice="bsf-management-service",method="GET"}[24h]))) * 100 >= 5
OID	1.3.6.1.4.1.323.5.3.37.1.2.5
Metric Used	http_server_requests_seconds_count

Table 5-24 (Cont.) PCF_BINDING_ERROR_RATE_ABOVE_MAJOR_THRESHOLD

Field	Details
Recommended Actions	The alert gets cleared when the number of failed transactions is below 5% of the total transactions. To assess the reason for failed transactions, check the service specific metrics for the GET method. For any assistance, contact My Oracle Support .

5.1.25

INGRESS_CREATE_ERROR_RATE_ABOVE_MINOR_THRESHOLD

Table 5-25 INGRESS_CREATE_ERROR_RATE_ABOVE_MINOR_THRESHOLD

Field	Details
Description	BSF Ingress Create Error Rate above 1 Percent in <code>{{labels.microservice}}</code> in <code>{{labels.namespace}}</code>
Summary	Transaction Create Error Rate in <code>{{labels.kubernetes_node}}</code> (current value is: <code>{{ \$value }}</code>)
Severity	Minor
Expression	<pre>sum by(namespace) (rate(http_server_requests_seconds_count{microservice="bsf-management-service", status!="200 201",method="POST"}[24h])) / sum by(namespace) (rate(http_server_requests_seconds_count{microservice="bsf-management-service",method="POST"}[24h])) * 100 >= 1</pre>
OID	1.3.6.1.4.1.323.5.3.37.1.2.4
Metric Used	http_server_requests_seconds_count
Recommended Actions	The alert gets cleared when the number of failed transactions is below 1% of the total transactions. To assess the reason for failed transactions, check the service specific metrics for the POST method. For any assistance, contact My Oracle Support .

5.1.26

INGRESS_CREATE_ERROR_RATE_ABOVE_CRITICAL_THRESHOLD

Table 5-26 INGRESS_CREATE_ERROR_RATE_ABOVE_CRITICAL_THRESHOLD

Field	Details
Description	BSF Ingress Create Error Rate above 10 Percent in <code>{{labels.microservice}}</code> in <code>{{labels.namespace}}</code>
Summary	Transaction Create Error Rate in <code>{{labels.kubernetes_node}}</code> (current value is: <code>{{ \$value }}</code>)

Table 5-26 (Cont.) INGRESS_CREATE_ERROR_RATE_ABOVE_CRITICAL_THRESHOLD

Field	Details
Severity	Critical
Expression	sum by(namespace) (rate(http_server_requests_seconds_count{microservice="bsf-management-service", status!~"200 201",method="POST"}[24h])) / sum by(namespace) (rate(http_server_requests_seconds_count{microservice="bsf-management-service",method="POST"}[24h])) * 100 >= 10
OID	1.3.6.1.4.1.323.5.3.37.1.2.4
Metric Used	http_server_requests_seconds_count
Recommended Actions	The alert gets cleared when the number of failed transactions is below 10% of the total transactions. To assess the reason for failed transactions, check the service specific metrics for the POST method. For any assistance, contact My Oracle Support .

5.1.27

INGRESS_CREATE_ERROR_RATE_ABOVE_MAJOR_THRESHOLD

Table 5-27 INGRESS_CREATE_ERROR_RATE_ABOVE_MAJOR_THRESHOLD

Field	Details
Description	BSF Ingress Create Error Rate above 5 Percent in {{ \$labels.microservice }} in {{ \$labels.namespace }}
Summary	Transaction Create Error Rate in {{ \$labels.kubernetes_node }} (current value is: {{ \$value }})
Severity	Major
Expression	sum by(namespace) (rate(http_server_requests_seconds_count{microservice="bsf-management-service", status!~"200 201",method="POST"}[24h])) / sum by(namespace) (rate(http_server_requests_seconds_count{microservice="bsf-management-service",method="POST"}[24h])) * 100 >= 5
OID	1.3.6.1.4.1.323.5.3.37.1.2.4
Metric Used	http_server_requests_seconds_count
Recommended Actions	The alert gets cleared when the number of failed transactions is below 5% of the total transactions. To assess the reason for failed transactions, check the service specific metrics for the POST method. For any assistance, contact My Oracle Support .

5.1.28

INGRESS_DELETE_ERROR_RATE_ABOVE_MINOR_THRESHOLD

Table 5-28 INGRESS_DELETE_ERROR_RATE_ABOVE_MINOR_THRESHOLD

Field	Details
Description	Ingress Delete Error Rate above 1 Percent in <code>{{labels.microservice}}</code> in <code>{{labels.namespace}}</code>
Summary	Ingress Delete Error Rate in <code>{{labels.kubernetes_node}}</code> (current value is: <code>{{ \$value }}</code>)
Severity	Minor
Expression	$\frac{\text{sum by(namespace) (rate(http_server_requests_seconds_count\{microservice=\"bsf-management-service\", status!=\"204\", method=\"DELETE\", uri=\"/nbsf-management/v1/pcfBindings/{bindingId}\"}[24h]))}}{\text{sum by(namespace) (rate(http_server_requests_seconds_count\{microservice=\"bsf-management-service\", method=\"DELETE\", uri=\"/nbsf-management/v1/pcfBindings/{bindingId}\"}[24h]))}} * 100 \geq 1$
OID	1.3.6.1.4.1.323.5.3.37.1.2.6
Metric Used	http_server_requests_seconds_count
Recommended Actions	The alert gets cleared when the number of failed transactions is below 1% of the total transactions. To assess the reason for failed transactions, check the service specific metrics for the DELETE method. For any assistance, contact My Oracle Support.

5.1.29

INGRESS_DELETE_ERROR_RATE_ABOVE_MAJOR_THRESHOLD

Table 5-29 INGRESS_DELETE_ERROR_RATE_ABOVE_MAJOR_THRESHOLD

Field	Details
Description	Ingress Delete Error Rate above 5 Percent in <code>{{labels.microservice}}</code> in <code>{{labels.namespace}}</code>
Summary	Ingress Delete Error Rate in <code>{{labels.kubernetes_node}}</code> (current value is: <code>{{ \$value }}</code>)
Severity	Major

Table 5-29 (Cont.) INGRESS_DELETE_ERROR_RATE_ABOVE_MAJOR_THRESHOLD

Field	Details
Expression	$\frac{\text{sum by(namespace) (rate(http_server_requests_seconds_count\{microservice="bsf-management-service", status!="204", method="DELETE", uri="/nbsf-management/v1/pcfBindings/{bindingId}"\}[24h]))}}{\text{sum by(namespace) (rate(http_server_requests_seconds_count\{microservice="bsf-management-service", method="DELETE", uri="/nbsf-management/v1/pcfBindings/{bindingId}"\}[24h]))}} * 100 \geq 5$
OID	1.3.6.1.4.1.323.5.3.37.1.2.6
Metric Used	http_server_requests_seconds_count
Recommended Actions	The alert gets cleared when the number of failed transactions is below 5% of the total transactions. To assess the reason for failed transactions, check the service specific metrics for the DELETE method. For any assistance, contact My Oracle Support.

5.1.30

INGRESS_DELETE_ERROR_RATE_ABOVE_CRITICAL_THRESHOLD

Table 5-30 INGRESS_DELETE_ERROR_RATE_ABOVE_CRITICAL_THRESHOLD

Field	Details
Description	Ingress Delete Error Rate above 10 Percent in <code>{{labels.microservice}}</code> in <code>{{labels.namespace}}</code>
Summary	Ingress Delete Error Rate in <code>{{labels.kubernetes_node}}</code> (current value is: <code>{{ \$value }}</code>)
Severity	Critical
Expression	$\frac{\text{sum by(namespace) (rate(http_server_requests_seconds_count\{microservice="bsf-management-service", status!="204", method="DELETE", uri="/nbsf-management/v1/pcfBindings/{bindingId}"\}[24h]))}}{\text{sum by(namespace) (rate(http_server_requests_seconds_count\{microservice="bsf-management-service", method="DELETE", uri="/nbsf-management/v1/pcfBindings/{bindingId}"\}[24h]))}} * 100 \geq 10$
OID	1.3.6.1.4.1.323.5.3.37.1.2.6
Metric Used	http_server_requests_seconds_count

Table 5-30 (Cont.) INGRESS_DELETE_ERROR_RATE_ABOVE_CRITICAL_THRESHOLD

Field	Details
Recommended Actions	The alert gets cleared when the number of failed transactions is below 10% of the total transactions. To assess the reason for failed transactions, check the service specific metrics for the DELETE method. For any assistance, contact My Oracle Support.

5.1.31 DB_TIER_DOWN_ALERT

Table 5-31 DB_TIER_DOWN_ALERT

Field	Details
Description	DB cannot be reachable!
Summary	DB cannot be reachable!
Severity	Critical
Expression	appinfo_category_running{category="database", application="ocbsf"} != 1
OID	1.3.6.1.4.1.323.5.3.37.1.2.7
Metric Used	appinfo_category_running
Recommended Actions	Cause: The database tier is unavailable or unreachable, as indicated by appinfo_category_running{category="database"} != 1. Diagnostic Information: Verify database-pod status, readiness, events, network connectivity, database logs, storage and PVC health, and the appinfo_category_running metric. Recovery: Restore the failed database component or connectivity, resolve storage or configuration issues, and confirm that the database category reports as running again.

5.1.32 CPU_USAGE_PER_SERVICE_ABOVE_MINOR_THRESHOLD

Table 5-32 CPU_USAGE_PER_SERVICE_ABOVE_MINOR_THRESHOLD

Field	Details
Description	CPU usage for {{labels.microservice}} service is above 60
Summary	CPU usage for {{labels.microservice}} service is above 60
Severity	Minor
Expression	avg by (namespace,service) (avg_over_time(cgroup_cpu_usage{application="ocbsf"}[2m])) >= 60 < 80
OID	1.3.6.1.4.1.323.5.3.37.1.2.8

Table 5-32 (Cont.) CPU_USAGE_PER_SERVICE_ABOVE_MINOR_THRESHOLD

Field	Details
Metric Used	cgroup_cpu_usage
Recommended Actions	Cause: CPU utilization for a service exceeds the configured minor threshold over the defined time window. This typically indicates increased processing load due to higher traffic, inefficient code paths, resource-intensive operations, traffic spikes, background jobs, tight loops, thread contention, or misconfigured CPU limits or requests. Diagnostic Information: Check CPU-usage graphs to determine whether utilization is consistently high or intermittent. Correlate CPU usage with traffic patterns, batch jobs, and recent deployments. Check pod metrics for CPU throttling, restarts, or abnormal usage. Review latency metrics, upstream-service logs for increased egress traffic or retries, service logs for high-latency operations or repeated loops, CPU limits and requests, and node-level CPU contention. Recovery: If CPU usage remains high, optimize application logic, reduce unnecessary processing, or tune configuration. If traffic spikes are responsible, scale out the service. Cautiously adjust CPU limits or requests if required, ensuring cluster stability, and roll back recent changes if the issue started after a deployment. The alert clears when CPU usage falls below the configured threshold.

5.1.33 CPU_USAGE_PER_SERVICE_ABOVE_MAJOR_THRESHOLD

Table 5-33 CPU_USAGE_PER_SERVICE_ABOVE_MAJOR_THRESHOLD

Field	Details
Description	CPU usage for {{\$labels.microservice}} service is above 80
Summary	CPU usage for {{\$labels.microservice}} service is above 80
Severity	Major
Expression	avg by (namespace,service) (avg_over_time(cgroup_cpu_usage{application="o cbsf"}[2m])) >= 80 < 90
OID	1.3.6.1.4.1.323.5.3.37.1.2.9
Metric Used	cgroup_cpu_usage

Table 5-33 (Cont.) CPU_USAGE_PER_SERVICE_ABOVE_MAJOR_THRESHOLD

Field	Details
Recommended Actions	Cause: CPU utilization for a service exceeds the configured major threshold of 80% over the defined time window. This can indicate sustained high traffic, inefficient processing, tight loops, thread contention, retry storms, workload spikes, downstream latency, or misconfigured CPU limits or requests. Diagnostic Information: Check CPU-usage graphs to confirm sustained utilization above 80%. Correlate CPU usage with traffic patterns, batch jobs, and recent deployments. Check pod metrics for CPU throttling, restarts, and saturation indicators. Review latency metrics, upstream-service logs for increased egress traffic or retries, service logs for CPU-intensive operations or loops, CPU limits and requests, and node-level CPU contention. Recovery: Scale out the service to reduce CPU load per pod. Optimize application logic or disable non-critical processing, mitigate retry storms or abnormal traffic, cautiously adjust CPU limits or requests, and roll back recent deployments if applicable. The alert clears when CPU usage falls below the configured major threshold.

5.1.34 CPU_USAGE_PER_SERVICE_ABOVE_CRITICAL_THRESHOLD

Table 5-34 CPU_USAGE_PER_SERVICE_ABOVE_CRITICAL_THRESHOLD

Field	Details
Description	CPU usage for {{labels.microservice}} service is above 90
Summary	CPU usage for {{labels.microservice}} service is above 90
Severity	Critical
Expression	avg by (namespace,service) (avg_over_time(cgroup_cpu_usage{application="o cbsf"}[2m])) >= 90
OID	1.3.6.1.4.1.323.5.3.37.1.2.10
Metric Used	cgroup_cpu_usage

Table 5-34 (Cont.) CPU_USAGE_PER_SERVICE_ABOVE_CRITICAL_THRESHOLD

Field	Details
Recommended Actions	Cause: CPU utilization for a service exceeds the configured critical threshold over the defined time window. This indicates severe resource saturation and a risk of service degradation or unresponsiveness. Typical causes include sustained high traffic, inefficient processing, tight loops, thread contention, misconfigured CPU limits or requests, retry storms, workload spikes, or downstream latency. Diagnostic Information: Check CPU-usage graphs to confirm sustained critical utilization. Correlate CPU usage with traffic patterns, batch jobs, and recent deployments. Check pod metrics for CPU throttling, restarts, and saturation indicators. Review latency metrics, upstream-service logs for increased egress traffic or retries, service logs for CPU-intensive operations or loops, CPU limits and requests, and node-level CPU contention. Recovery: Scale out the service immediately to reduce CPU load per pod. If required, perform a rolling restart to stabilize pods under extreme load. Optimize application logic or disable non-critical processing, mitigate retry storms or abnormal traffic, cautiously adjust CPU limits or requests, and roll back recent deployments if applicable. The alert clears when CPU usage falls below the configured critical threshold.

5.1.35 MEMORY_USAGE_PER_SERVICE_ABOVE_MINOR_THRESHOLD

Table 5-35 MEMORY_USAGE_PER_SERVICE_ABOVE_MINOR_THRESHOLD

Field	Details
Description	Memory usage for {{\$labels.microservice}} service is above 60
Summary	Memory usage for {{\$labels.microservice}} service is above 60
Severity	Minor
Expression	avg by (namespace,service) (avg_over_time(cgroup_memory_usage{application=~"ocbsf"}[2m])) >= 60 < 80
OID	1.3.6.1.4.1.323.5.3.37.1.2.11
Metric Used	cgroup_memory_usage

Table 5-35 (Cont.) MEMORY_USAGE_PER_SERVICE_ABOVE_MINOR_THRESHOLD

Field	Details
Recommended Actions	Cause: Service memory usage remains at or above 60% for the configured window. This can indicate sustained memory growth, repeated traffic or payload spikes, cache expansion, a memory leak in the PRE service, recent deployment changes, higher concurrency, or a misconfigured memory limit. Diagnostic Information: Check memory graphs for the PRE service to determine whether usage is consistently high or only spiking. Review whether the alert aligns with traffic bursts, deployment changes, or a particular job or request pattern. Recovery: Scale out the service or increase memory limits only after confirming the trend and ruling out a memory leak or abnormal workload.

5.1.36

MEMORY_USAGE_PER_SERVICE_ABOVE_MAJOR_THRESHOLD

Table 5-36 MEMORY_USAGE_PER_SERVICE_ABOVE_MAJOR_THRESHOLD

Field	Details
Description	Memory usage for {{\$labels.microservice}} service is above 80
Summary	Memory usage for {{\$labels.microservice}} service is above 80
Severity	Major
Expression	avg by (namespace,service) (avg_over_time(cgroup_memory_usage{application=~"ocbsf"}[2m])) >= 80 < 90
OID	1.3.6.1.4.1.323.5.3.37.1.2.12
Metric Used	cgroup_memory_usage
Recommended Actions	Cause: Service memory usage remains at or above 80% for the configured window. This can indicate sustained memory growth, repeated traffic or payload spikes, cache expansion, a memory leak in the PRE service, recent deployment changes, higher concurrency, or a misconfigured memory limit. Diagnostic Information: Check memory graphs for the PRE service to determine whether usage is consistently high or only spiking. Review whether the alert aligns with traffic bursts, deployment changes, or a particular job or request pattern. Recovery: Scale out the service or increase memory limits only after confirming the trend and ruling out a memory leak or abnormal workload.

5.1.37

MEMORY_USAGE_PER_SERVICE_ABOVE_CRITICAL_THRESHOLD

Table 5-37 MEMORY_USAGE_PER_SERVICE_ABOVE_CRITICAL_THRESHOLD

Field	Details
Description	Memory usage for {{labels.microservice}} service is above 90
Summary	Memory usage for {{labels.microservice}} service is above 90
Severity	Critical
Expression	avg by (namespace,service) (avg_over_time(cgroup_memory_usage{application="ocbsf"}[2m])) >= 90
OID	1.3.6.1.4.1.323.5.3.37.1.2.13
Metric Used	cgroup_memory_usage
Recommended Actions	Cause: Service memory usage remains at or above 90% for the configured window. This can indicate sustained memory growth, repeated traffic or payload spikes, cache expansion, a memory leak in the PRE service, recent deployment changes, higher concurrency, or a misconfigured memory limit. Diagnostic Information: Check memory graphs for the PRE service to determine whether usage is consistently high or only spiking. Review whether the alert aligns with traffic bursts, deployment changes, or a particular job or request pattern. Recovery: If memory usage reaches approximately 90% of the defined threshold, perform a rolling restart of PRE pods to relieve memory pressure. Investigate and correct underlying causes, such as leaks, cache growth, or abnormal traffic patterns. If the issue persists or the root cause is unclear, contact Oracle Support for further analysis.

5.1.38 NRF_COMMUNICATION_FAILURE

Table 5-38 NRF_COMMUNICATION_FAILURE

Field	Details
Description	There has been an external failure communication error with NRF.
Summary	There has been an external failure communication error with NRF.
Severity	Critical
Expression	ocbsf_nrfclient_nrf_operative_status == 0
OID	1.3.6.1.4.1.323.5.3.37.1.2.33
Metric Used	ocbsf_nrfclient_nrf_operative_status

Table 5-38 (Cont.) NRF_COMMUNICATION_FAILURE

Field	Details
Recommended Actions	For any assistance, contact My Oracle Support .

5.1.39 NRF_SERVICE_REQUEST_FAILURE

Table 5-39 NRF_SERVICE_REQUEST_FAILURE

Field	Details
Description	There has been a Service Request Failure with NRF, either a Registration failure, Heartbeat failure, or Profile Update Failure.
Summary	There has been a Service Request Failure with NRF, either a Registration failure, Heartbeat failure, or Profile Update Failure.
Severity	Critical
Expression	ocbsf_nrfclient_nfUpdate_status == 0
OID	1.3.6.1.4.1.323.5.3.37.1.2.34
Metric Used	ocbsf_nrfclient_nfUpdate_status
Recommended Actions	For any assistance, contact My Oracle Support .

5.1.40 PERF_INFO_ACTIVE_OVERLOAD_THRESHOLD_FETCH_FAILED

Table 5-40 PERF_INFO_ACTIVE_OVERLOAD_THRESHOLD_FETCH_FAILED

Field	Details
Description	The application fails to get the current active overload level threshold data.
Summary	The application raises PERF_INFO_ACTIVE_OVERLOAD_THRESHOLD_FETCH_FAILED alert when it fails to fetch the current active overload level threshold data and active_overload_threshold_fetch_failed == 1.
Severity	Major
Expression	active_overload_threshold_fetch_failed == 1
OID	1.3.6.1.4.1.323.5.3.37.1.2.20
Metric Used	active_overload_threshold_fetch_failed

Table 5-40 (Cont.) PERF_INFO_ACTIVE_OVERLOAD_THRESHOLD_FETCH_FAILED

Field	Details
Recommended Actions	Cause: The system cannot fetch active overload-threshold data, as indicated by <code>active_overload_threshold_fetch_failed == 1</code>. Overload-protection parameters might be unavailable or incorrectly loaded, affecting the system's ability to handle traffic surges safely. Diagnostic Information: Verify connectivity to the configuration source, such as the database, configuration service, or NMS. Confirm that overload-threshold data is configured and available. Review logs for fetch failures or timeouts, validate credentials and access permissions, inspect recent configuration or deployment changes, confirm that dependent services are available, and check for initialization issues during service startup. Recovery: Restore connectivity to the configuration or data source. Correct missing or invalid overload-threshold configuration, restart or reinitialize the component that retrieves the data, resolve authentication or authorization failures, and roll back recent changes that affect configuration retrieval. Confirm high availability of configuration services and monitor until the metric returns to zero. If the issue persists, contact Oracle Support.

5.1.41 POD_DOC

Table 5-41 POD_DOC

Field	Details
Description	Pod Congestion status of {{\$labels.microservice}} service is DoC
Summary	Pod Congestion status of {{\$labels.microservice}} service is DoC
Severity	Major
Expression	<code>ocbsf_pod_congestion_state == 1</code>
OID	1.3.6.1.4.1.323.5.3.37.1.2.25
Metric Used	<code>ocbsf_pod_congestion_state</code>

Table 5-41 (Cont.) POD_DOC

Field	Details
Recommended Actions	Cause: The pod entered DANGER_OF_CONGESTION (DOC) due to rising CPU and/or queue close to configured limits. Diagnostic Information: Check pod_congestion_state == 1; review pod_resource_stress (cpu and queue) and pod_cong_state_report_total to see recent transitions. Recovery: Confirm the DOC thresholds and active Load Shedding rule. If DOC is triggered by brief spikes, increase stateChangeSampleCount or the calculation interval. If sustained, consider slightly increasing discard aggressiveness for low-value calls per policy.

5.1.42 POD_CONGESTED

Table 5-42 POD_CONGESTED

Field	Details
Description	Pod Congestion status of {{\$labels.microservice}} service is congested
Summary	Pod Congestion status of {{\$labels.microservice}} service is congested
Severity	Critical
Expression	ocbsf_pod_congestion_state==4
OID	1.3.6.1.4.1.323.5.3.37.1.2.26
Metric Used	ocbsf_pod_congestion_state

Table 5-42 (Cont.) POD_CONGESTED

Field	Details
Recommended Actions	Cause: The pod has entered the CONGESTED state because CPU consumption or the pending-request queue exceeds the active threshold profile. Diagnostic Information: Check pod_congestion_state and confirm a value of 4. Review pod_resource_congestion_state for CPU and queue conditions, pod_resource_stress, and ocbsf_http_congestion_message_reject_to tal, filtered by congestion state, request URI, request method, and priority. Recovery: In the CM GUI, navigate to BSF, Overload and Congestion Control, and Congestion Control. Confirm that the feature is enabled and that the intended Thresholds and Load Shedding profiles are active. If rejections are excessive, raise the discard priority for the current state or relax thresholds according to performance baselines. Consider increasing stateChangeSampleCount or the calculation interval to reduce flapping caused by short spikes.

5.1.43 POD_CONGESTION_L1

Table 5-43 POD_CONGESTION_L1

Field	Details
Description	Pod Congestion status of {{\$labels.microservice}} service is Congestion_L1.
Summary	Pod Congestion status of {{\$labels.microservice}} service is Congestion_L1.
Severity	Critical
Expression	ocbsf_pod_congestion_state == 2
OID	1.3.6.1.4.1.323.5.3.37.1.2.52
Metric Used	ocbsf_pod_congestion_state
Recommended Actions	Cause: The pod reached CONGESTION_L1 based on CPU or queue thresholds in the active profile. Diagnostic Information: Confirm that pod_congestion_state == 2. Identify the driver by reviewing pod_resource_congestion_state for CPU or queue conditions, and review ocbsf_http_congestion_message_reject_to tal with congestionState=CONGESTION_L1. Recovery: Confirm the L1 discard priority, which defaults to 24, and the configured thresholds. If important calls are being dropped, adjust the discard priority or tune thresholds to match the expected load profile.

5.1.44 POD_CPU_CONGESTION_L1

Table 5-44 POD_CPU_CONGESTION_L1

Field	Details
Description	Pod resource is in Congestion_L1 for CPU type.
Summary	Pod Resource Congestion status of {{labels.microservice}} service is Congestion_L1 for CPU type.
Severity	Critical
Expression	ocbsf_pod_resource_congestion_state{type="cpu"} == 2
OID	1.3.6.1.4.1.323.5.3.37.1.2.54
Metric Used	ocbsf_pod_resource_congestion_state
Recommended Actions	Cause: CPU utilization has reached CONGESTION_L1. Diagnostic Information: Confirm that pod_resource_congestion_state{resourceType="cpu"} == 2. Check CPU stress and state transitions. Recovery: Validate the L1 CPU thresholds and discard priority. If short spikes cause churn, increase stateChangeSampleCount; otherwise, increase shedding at L1.

5.1.45 POD_CONGESTION_L2

Table 5-45 POD_CONGESTION_L2

Field	Details
Description	Pod Congestion status of {{labels.microservice}} service is Congestion_L2
Summary	Pod Resource Congestion status of {{labels.microservice}} service is Congestion_L2.
Severity	Critical
Expression	ocbsf_pod_congestion_state == 3
OID	1.3.6.1.4.1.323.5.3.37.1.2.53
Metric Used	ocbsf_pod_congestion_state

Table 5-45 (Cont.) POD_CONGESTION_L2

Field	Details
Recommended Actions	Cause: The pod reached CONGESTION_L2, indicating higher stress than L1. Diagnostic Information: Confirm that <code>pod_congestion_state == 3</code>. Validate resource-specific congestion states and stress metrics, and inspect rejection counters at L2. Recovery: Use the L2 discard priority, which defaults to 18, to shed more low-priority traffic. Consider tuning thresholds and sample counts to balance protection and availability.

5.1.46 POD_CPU_CONGESTION_L2

Table 5-46 POD_CPU_CONGESTION_L2

Field	Details
Description	Pod Resource Congestion status of <code>{{labels.microservice}}</code> service is Congestion_L2 for CPU type.
Summary	Pod Resource Congestion status of <code>{{labels.microservice}}</code> service is Congestion_L2 for CPU type.
Severity	Critical
Expression	<code>ocbsf_pod_resource_congestion_state{type="cpu"} == 3</code>
OID	1.3.6.1.4.1.323.5.3.37.1.2.55
Metric Used	<code>ocbsf_pod_resource_congestion_state</code>
Recommended Actions	Cause: CPU utilization has reached CONGESTION_L2. Diagnostic Information: Confirm that <code>pod_resource_congestion_state{resourceType="cpu"} == 3</code>. Check CPU stress, EMA interval and ratio, and L2 rejection counters. Recovery: Raise the L2 discard priority to protect the pod. Tune CPU thresholds or EMA cadence only after comparison with test baselines.

5.1.47 POD_PENDING_REQUEST_DOC

Table 5-47 POD_PENDING_REQUEST_DOC

Field	Details
Description	Pod Resource Congestion status of <code>{{labels.microservice}}</code> service is DoC for PendingRequest type

Table 5-47 (Cont.) POD_PENDING_REQUEST_DOC

Field	Details
Summary	Pod Resource Congestion status of {{labels.microservice}} service is DoC for PendingRequest type
Severity	Major
Expression	ocbsf_pod_resource_congestion_state{type="queue"} == 1
OID	1.3.6.1.4.1.323.5.3.37.1.2.27
Metric Used	ocbsf_pod_resource_congestion_state{type="queue"}
Recommended Actions	Cause: The pending request queue is in the DANGER_OF_CONGESTION state. Diagnostic Information: Confirm that pod_resource_congestion_state{resourceType="queue"} == 1 and review queue-related pod_resource_stress. Recovery: Review queue DoC thresholds. If early protection is required, allow gentle shedding of the lowest-priority traffic at DoC; otherwise, tune thresholds to match observed load.

5.1.48 POD_PENDING_REQUEST_CONGESTED

Table 5-48 POD_PENDING_REQUEST_CONGESTED

Field	Details
Description	Pod Resource Congestion status of {{labels.microservice}} service is congested for PendingRequest type
Summary	Pod Resource Congestion status of {{labels.microservice}} service is congested for PendingRequest type
Severity	Critical
Expression	ocbsf_pod_resource_congestion_state{type="queue"} == 4
OID	1.3.6.1.4.1.323.5.3.37.1.2.28
Metric Used	ocbsf_pod_resource_congestion_state{type="queue"}

Table 5-48 (Cont.) POD_PENDING_REQUEST_CONGESTED

Field	Details
Recommended Actions	Cause: The pending HTTP request queue is in the CONGESTED state. Diagnostic Information: Confirm that <code>pod_resource_congestion_state{resourceType="queue"} == 4</code> and review <code>pod_resource_stress{resourceType="queue"}</code>. Check <code>ocbsf_http_congestion_message_reject_total</code> for low-priority discards at this level. Recovery: Validate the queue thresholds in the active profile and the CONGESTED discard priority. If backlog persists, increase shedding by raising the discard priority so that lower-priority requests are rejected earlier.

5.1.49 POD_CPU_DOC

Table 5-49 POD_CPU_DOC

Field	Details
Description	Pod Resource Congestion status of <code>{{labels.microservice}}</code> service is DoC for CPU type
Summary	Pod Resource Congestion status of <code>{{labels.microservice}}</code> service is DoC for CPU type
Severity	Major
Expression	<code>ocbsf_pod_resource_congestion_state{type="cpu"} == 1</code>
OID	1.3.6.1.4.1.323.5.3.37.1.2.29
Metric Used	<code>ocbsf_pod_resource_congestion_state{type="cpu"}</code>
Recommended Actions	Cause: CPU utilization is in the DANGER_OF_CONGESTION state according to the current thresholds and EMA settings. Diagnostic Information: Confirm that <code>pod_resource_congestion_state{resourceType="cpu"} == 1</code> and check <code>pod_resource_stress{resourceType="cpu"}</code>. Confirm the EMA parameters, including the interval and 70:30 ratios. Recovery: If the condition is transient, increase <code>stateChangeSampleCount</code> to avoid oscillation. If it is sustained, adjust the CPU DoC threshold or enable mild shedding for non-critical, low-priority requests.

5.1.50 POD_CPU_CONGESTED

Table 5-50 POD_CPU_CONGESTED

Field	Details
Description	Pod Resource Congestion status of <code>{{labels.microservice}}</code> service is congested for CPU type
Summary	Pod Resource Congestion status of <code>{{labels.microservice}}</code> service is congested for CPU type
Severity	Critical
Expression	<code>ocbsf_pod_resource_congestion_state{type="cpu"} == 4</code>
OID	1.3.6.1.4.1.323.5.3.37.1.2.30
Metric Used	<code>ocbsf_pod_resource_congestion_state</code>
Recommended Actions	Cause: CPU utilization has reached the CONGESTED state. Diagnostic Information: Validate <code>pod_resource_congestion_state{resourceType="cpu"}</code> for values 2 or 4, according to the alert rule, and check <code>pod_resource_stress{resourceType="cpu"}</code>. Review message rejections at this congestion level. Recovery: Tighten protection by raising the discard priority for this state so that more low-priority requests are dropped. Reassess CPU thresholds and EMA intervals only after reviewing performance benchmarks.

5.1.51 POD_MEMORY_DOC

Table 5-51 POD_MEMORY_DOC

Field	Details
Description	Pod Resource Congestion status of <code>{{labels.microservice}}</code> service is DoC for Memory type
Summary	Pod Resource Congestion status of <code>{{labels.microservice}}</code> service is DoC for Memory type
Severity	Major
Expression	<code>ocbsf_pod_resource_congestion_state{type="memory"} == 1</code>
OID	1.3.6.1.4.1.323.5.3.37.1.2.31
Metric Used	<code>ocbsf_pod_resource_congestion_state{type="memory"}</code>

Table 5-51 (Cont.) POD_MEMORY_DOC

Field	Details
Recommended Actions	Cause: Pod memory utilization has entered the DANGER_OF_CONGESTION state according to the configured memory thresholds and EMA. This indicates increasing memory pressure that can lead to congestion and request shedding. Diagnostic Information: Confirm that <code>pod_resource_congestion_state{type="memory"} == 1</code>. Review the <code>pod_resource_stress{resourceType="memory"}</code> trend to determine whether pressure is sustained or spiky. Correlate the condition with working-set or RSS metrics, garbage-collection activity, traffic spikes, configuration changes, pod restarts, OOM indicators, and application logs. Recovery: For transient spikes, increase <code>stateChangeSampleCount</code> or the congestion-calculation interval to reduce flapping. For sustained pressure, reduce load by adjusting shedding for low-priority traffic or increase pod memory limits and requests according to the capacity plan. Investigate memory-growth or leak patterns and large in-flight requests; remediate at the application level if memory does not stabilize.

5.1.52 POD_MEMORY_CONGESTED

Table 5-52 POD_MEMORY_CONGESTED

Field	Details
Description	Pod Resource Congestion status of <code>{{labels.microservice}}</code> service is congested for Memory type
Summary	Pod Resource Congestion status of <code>{{labels.microservice}}</code> service is congested for Memory type
Severity	Critical
Expression	<code>ocbsf_pod_resource_congestion_state{type="memory"} == 2</code>
OID	1.3.6.1.4.1.323.5.3.37.1.2.32
Metric Used	<code>ocbsf_pod_resource_congestion_state{type="memory"}</code>

Table 5-52 (Cont.) POD_MEMORY_CONGESTED

Field	Details
Recommended Actions	Cause: Memory utilization has reached the CONGESTED state according to the active thresholds and EMA configuration. The pod is under high memory stress and might reject or drop low-priority requests to protect stability. Diagnostic Information: Confirm that <code>pod_resource_congestion_state{type="memory"} == 2</code>. Inspect <code>pod_resource_stress{resourceType="memory"}</code> and container memory consumption against limits. Check for OOM kills, restarts, Kubernetes events, request-size increases, traffic surges, cache growth, and increasing request rejections. Recovery: Increase shedding aggressiveness by raising the discard priority so that lower-priority requests are rejected earlier. If congestion persists, scale out the workload or increase memory limits according to performance baselines, then reassess thresholds. If OOM kills or restarts occur, stabilize memory by reducing load or rolling back recent changes before tuning thresholds.

5.1.53 SERVICE_OVERLOADED

Table 5-53 SERVICE_OVERLOADED

Field	Details
Description	Overload Level of <code>{{labels.microservice}}</code> service is L1
Summary	Overload Level of <code>{{labels.microservice}}</code> service is L1
Severity	Minor
Expression	<code>load_level == 1</code>
OID	1.3.6.1.4.1.323.5.3.37.1.2.14
Metric Used	<code>load_level</code>

Table 5-53 (Cont.) SERVICE_OVERLOADED

Field	Details
Recommended Actions	Cause: The service can be overloaded because of high CPU utilization and a large backlog, an extreme backlog with increasing failures caused by downstream latency or an outage, or widespread failures caused by network, DNS, certificate, routing, release, or dependency issues. Diagnostic Information: Determine whether CPU saturation, a runaway pending-request queue, or failures and timeouts are the primary condition. Check cluster and platform health, including node readiness, network, DNS, CNI, ingress, and load balancer status. Confirm that critical traffic is correctly classified. Recovery: Apply strict L3 discard rules to reject low-priority and, if required, medium-priority traffic to preserve high-priority flows. Increase percentage-based rejection where supported. If dependency failures are responsible, fix the dependency before scaling the service. If compute capacity is the issue and dependencies are healthy, rapidly scale replicas or nodes and temporarily reduce expensive traffic or features. Engage the relevant platform and dependency owners because L3 is typically a cross-team incident.

Table 5-54 SERVICE_OVERLOADED

Field	Details
Description	Overload Level of {{labels.microservice}} service is L2
Summary	Overload Level of {{labels.microservice}} service is L2
Severity	Major
Expression	load_level == 2
OID	1.3.6.1.4.1.323.5.3.37.1.2.14
Metric Used	load_level

Table 5-54 (Cont.) SERVICE_OVERLOADED

Field	Details
Recommended Actions	Cause: The service can be overloaded because of high CPU utilization and a large backlog, an extreme backlog with increasing failures caused by downstream latency or an outage, or widespread failures caused by network, DNS, certificate, routing, release, or dependency issues. Diagnostic Information: Determine whether CPU saturation, a runaway pending-request queue, or failures and timeouts are the primary condition. Check cluster and platform health, including node readiness, network, DNS, CNI, ingress, and load balancer status. Confirm that critical traffic is correctly classified. Recovery: Apply strict L3 discard rules to reject low-priority and, if required, medium-priority traffic to preserve high-priority flows. Increase percentage-based rejection where supported. If dependency failures are responsible, fix the dependency before scaling the service. If compute capacity is the issue and dependencies are healthy, rapidly scale replicas or nodes and temporarily reduce expensive traffic or features. Engage the relevant platform and dependency owners because L3 is typically a cross-team incident.

Table 5-55 SERVICE_OVERLOADED

Field	Details
Description	Overload Level of {{\$labels.service}} service is L3
Summary	Overload Level of {{\$labels.service}} service is L3
Severity	Critical
Expression	load_level == 3
OID	1.3.6.1.4.1.323.5.3.37.1.2.14
Metric Used	load_level

Table 5-55 (Cont.) SERVICE_OVERLOADED

Field	Details
Recommended Actions	Cause: The service can be overloaded because of high CPU utilization and a large backlog, an extreme backlog with increasing failures caused by downstream latency or an outage, or widespread failures caused by network, DNS, certificate, routing, release, or dependency issues. Diagnostic Information: Determine whether CPU saturation, a runaway pending-request queue, or failures and timeouts are the primary condition. Check cluster and platform health, including node readiness, network, DNS, CNI, ingress, and load balancer status. Confirm that critical traffic is correctly classified. Recovery: Apply strict L3 discard rules to reject low-priority and, if required, medium-priority traffic to preserve high-priority flows. Increase percentage-based rejection where supported. If dependency failures are responsible, fix the dependency before scaling the service. If compute capacity is the issue and dependencies are healthy, rapidly scale replicas or nodes and temporarily reduce expensive traffic or features. Engage the relevant platform and dependency owners because L3 is typically a cross-team incident.

5.1.54 SERVICE_RESOURCE_OVERLOADED

Alerts when service is in overload state due to memory usage

Table 5-56 SERVICE_RESOURCE_OVERLOADED

Field	Details
Description	{{labels.microservice}} service is L1 for {{labels.type}} type
Summary	{{labels.microservice}} service is L1 for {{labels.type}} type
Severity	Minor
Expression	service_resource_overload_level{type="memory"} == 1
OID	1.3.6.1.4.1.323.5.3.37.1.2.15
Metric Used	service_resource_overload_level{type="memory"}

Table 5-56 (Cont.) SERVICE_RESOURCE_OVERLOADED

Field	Details
Recommended Actions	Cause: The service is at overload level L1 for memory. Possible drivers include a memory leak, higher concurrency, cache growth, large payloads, misconfigured memory limits or heap settings, and node memory pressure. Diagnostic Information: Review the memory-usage trend for each pod, OOM-kill and restart events, garbage-collection pressure where applicable, allocation spikes, configured limits and requests, node memory pressure, traffic, and recent changes. Recovery: Reduce memory pressure by tuning caches or payloads, fixing leaks, or rolling back regressions. Scale out to reduce per-pod memory usage and cautiously adjust memory limits. The alert clears when the overload level returns to zero.

Table 5-57 SERVICE_RESOURCE_OVERLOADED

Field	Details
Description	{{labels.microservice}} service is L2 for {{labels.type}} type
Summary	{{labels.microservice}} service is L2 for {{labels.type}} type
Severity	Major
Expression	service_resource_overload_level{type="memory"} == 2
OID	1.3.6.1.4.1.323.5.3.37.1.2.15
Metric Used	service_resource_overload_level{type="memory"}
Recommended Actions	Cause: The service is at overload level L1 for memory. Possible drivers include a memory leak, higher concurrency, cache growth, large payloads, misconfigured memory limits or heap settings, and node memory pressure. Diagnostic Information: Review the memory-usage trend for each pod, OOM-kill and restart events, garbage-collection pressure where applicable, allocation spikes, configured limits and requests, node memory pressure, traffic, and recent changes. Recovery: Reduce memory pressure by tuning caches or payloads, fixing leaks, or rolling back regressions. Scale out to reduce per-pod memory usage and cautiously adjust memory limits. The alert clears when the overload level returns to zero.

Table 5-58 SERVICE_RESOURCE_OVERLOADED

Field	Details
Description	{{labels.microservice}} service is L3 for {{labels.type}} type
Summary	{{labels.microservice}} service is L3 for {{labels.type}} type
Severity	Critical
Expression	service_resource_overload_level{type="memory"} == 3
OID	1.3.6.1.4.1.323.5.3.37.1.2.15
Metric Used	service_resource_overload_level{type="memory"}
Recommended Actions	Cause: The service is at overload level L1 for memory. Possible drivers include a memory leak, higher concurrency, cache growth, large payloads, misconfigured memory limits or heap settings, and node memory pressure. Diagnostic Information: Review the memory-usage trend for each pod, OOM-kill and restart events, garbage-collection pressure where applicable, allocation spikes, configured limits and requests, node memory pressure, traffic, and recent changes. Recovery: Reduce memory pressure by tuning caches or payloads, fixing leaks, or rolling back regressions. Scale out to reduce per-pod memory usage and cautiously adjust memory limits. The alert clears when the overload level returns to zero.

Alerts when service is in overload state due to CPU usage**Table 5-59 SERVICE_RESOURCE_OVERLOADED**

Field	Details
Description	{{labels.microservice}} service is L1 for {{labels.type}} type
Summary	{{labels.microservice}} service is L1 for {{labels.type}} type
Severity	Minor
Expression	service_resource_overload_level{type="cpu"} == 1
OID	1.3.6.1.4.1.323.5.3.37.1.2.15
Metric Used	service_resource_overload_level{type="cpu"}

Table 5-59 (Cont.) SERVICE_RESOURCE_OVERLOADED

Field	Details
Recommended Actions	Cause: The service is at overload level L1 for memory. Possible drivers include a memory leak, higher concurrency, cache growth, large payloads, misconfigured memory limits or heap settings, and node memory pressure. Diagnostic Information: Review the memory-usage trend for each pod, OOM-kill and restart events, garbage-collection pressure where applicable, allocation spikes, configured limits and requests, node memory pressure, traffic, and recent changes. Recovery: Reduce memory pressure by tuning caches or payloads, fixing leaks, or rolling back regressions. Scale out to reduce per-pod memory usage and cautiously adjust memory limits. The alert clears when the overload level returns to zero.

Table 5-60 SERVICE_RESOURCE_OVERLOADED

Field	Details
Description	{{labels.microservice}} service is L2 for {{labels.type}} type
Summary	{{labels.microservice}} service is L2 for {{labels.type}} type
Severity	Major
Expression	service_resource_overload_level{type="cpu"} == 2
OID	1.3.6.1.4.1.323.5.3.37.1.2.15
Metric Used	service_resource_overload_level{type="cpu"}
Recommended Actions	Cause: The service is at overload level L1 for memory. Possible drivers include a memory leak, higher concurrency, cache growth, large payloads, misconfigured memory limits or heap settings, and node memory pressure. Diagnostic Information: Review the memory-usage trend for each pod, OOM-kill and restart events, garbage-collection pressure where applicable, allocation spikes, configured limits and requests, node memory pressure, traffic, and recent changes. Recovery: Reduce memory pressure by tuning caches or payloads, fixing leaks, or rolling back regressions. Scale out to reduce per-pod memory usage and cautiously adjust memory limits. The alert clears when the overload level returns to zero.

Table 5-61 SERVICE_RESOURCE_OVERLOADED

Field	Details
Description	{{labels.microservice}} service is L3 for {{labels.type}} type
Summary	{{labels.microservice}} service is L3 for {{labels.type}} type
Severity	Critical
Expression	service_resource_overload_level{type="cpu"} == 3
OID	1.3.6.1.4.1.323.5.3.37.1.2.15
Metric Used	service_resource_overload_level{type="cpu"}
Recommended Actions	Cause: The service is at overload level L1 for memory. Possible drivers include a memory leak, higher concurrency, cache growth, large payloads, misconfigured memory limits or heap settings, and node memory pressure. Diagnostic Information: Review the memory-usage trend for each pod, OOM-kill and restart events, garbage-collection pressure where applicable, allocation spikes, configured limits and requests, node memory pressure, traffic, and recent changes. Recovery: Reduce memory pressure by tuning caches or payloads, fixing leaks, or rolling back regressions. Scale out to reduce per-pod memory usage and cautiously adjust memory limits. The alert clears when the overload level returns to zero.

Alerts when service is in overload state due to number of pending messages

Table 5-62 SERVICE_RESOURCE_OVERLOADED

Field	Details
Description	{{labels.microservice}} service is L1 for {{labels.type}} type
Summary	{{labels.microservice}} service is L1 for {{labels.type}} type
Severity	Minor
Expression	service_resource_overload_level{type="svc_pending_count"} == 1
OID	1.3.6.1.4.1.323.5.3.37.1.2.15
Metric Used	service_resource_overload_level{type="svc_pending_count"}

Table 5-62 (Cont.) SERVICE_RESOURCE_OVERLOADED

Field	Details
Recommended Actions	Cause: The service is at overload level L1 for memory. Possible drivers include a memory leak, higher concurrency, cache growth, large payloads, misconfigured memory limits or heap settings, and node memory pressure. Diagnostic Information: Review the memory-usage trend for each pod, OOM-kill and restart events, garbage-collection pressure where applicable, allocation spikes, configured limits and requests, node memory pressure, traffic, and recent changes. Recovery: Reduce memory pressure by tuning caches or payloads, fixing leaks, or rolling back regressions. Scale out to reduce per-pod memory usage and cautiously adjust memory limits. The alert clears when the overload level returns to zero.

Table 5-63 SERVICE_RESOURCE_OVERLOADED

Field	Details
Description	{{labels.microservice}} service is L2 for {{labels.type}} type
Summary	{{labels.microservice}} service is L2 for {{labels.type}} type
Severity	Major
Expression	service_resource_overload_level{type="svc_pending_count"} == 2
OID	1.3.6.1.4.1.323.5.3.37.1.2.15
Metric Used	service_resource_overload_level{type="svc_pending_count"}
Recommended Actions	Cause: The service is at overload level L1 for memory. Possible drivers include a memory leak, higher concurrency, cache growth, large payloads, misconfigured memory limits or heap settings, and node memory pressure. Diagnostic Information: Review the memory-usage trend for each pod, OOM-kill and restart events, garbage-collection pressure where applicable, allocation spikes, configured limits and requests, node memory pressure, traffic, and recent changes. Recovery: Reduce memory pressure by tuning caches or payloads, fixing leaks, or rolling back regressions. Scale out to reduce per-pod memory usage and cautiously adjust memory limits. The alert clears when the overload level returns to zero.

Table 5-64 SERVICE_RESOURCE_OVERLOADED

Field	Details
Description	{{labels.microservice}} service is L3 for {{labels.type}} type
Summary	{{labels.microservice}} service is L3 for {{labels.type}} type
Severity	Critical
Expression	service_resource_overload_level{type="svc_pending_count"} == 3
OID	1.3.6.1.4.1.323.5.3.37.1.2.15
Metric Used	service_resource_overload_level{type="svc_pending_count"}
Recommended Actions	Cause: The service is at overload level L1 for memory. Possible drivers include a memory leak, higher concurrency, cache growth, large payloads, misconfigured memory limits or heap settings, and node memory pressure. Diagnostic Information: Review the memory-usage trend for each pod, OOM-kill and restart events, garbage-collection pressure where applicable, allocation spikes, configured limits and requests, node memory pressure, traffic, and recent changes. Recovery: Reduce memory pressure by tuning caches or payloads, fixing leaks, or rolling back regressions. Scale out to reduce per-pod memory usage and cautiously adjust memory limits. The alert clears when the overload level returns to zero.

Alerts when service is in overload state due to number of failed requests**Table 5-65 SERVICE_RESOURCE_OVERLOADED**

Field	Details
Description	{{labels.microservice}} service is L1 for {{labels.type}} type
Summary	{{labels.microservice}} service is L1 for {{labels.type}} type
Severity	Minor
Expression	service_resource_overload_level{type="svc_failure_count"} == 1
OID	1.3.6.1.4.1.323.5.3.37.1.2.15
Metric Used	service_resource_overload_level{type="svc_failure_count"}

Table 5-65 (Cont.) SERVICE_RESOURCE_OVERLOADED

Field	Details
Recommended Actions	Cause: The service is at overload level L1 for memory. Possible drivers include a memory leak, higher concurrency, cache growth, large payloads, misconfigured memory limits or heap settings, and node memory pressure. Diagnostic Information: Review the memory-usage trend for each pod, OOM-kill and restart events, garbage-collection pressure where applicable, allocation spikes, configured limits and requests, node memory pressure, traffic, and recent changes. Recovery: Reduce memory pressure by tuning caches or payloads, fixing leaks, or rolling back regressions. Scale out to reduce per-pod memory usage and cautiously adjust memory limits. The alert clears when the overload level returns to zero.

Table 5-66 SERVICE_RESOURCE_OVERLOADED

Field	Details
Description	{{labels.microservice}} service is L2 for {{labels.type}} type
Summary	{{labels.microservice}} service is L2 for {{labels.type}} type
Severity	Major
Expression	service_resource_overload_level{type="svc_failure_count"} == 2
OID	1.3.6.1.4.1.323.5.3.37.1.2.15
Metric Used	service_resource_overload_level{type="svc_failure_count"}
Recommended Actions	Cause: The service is at overload level L1 for memory. Possible drivers include a memory leak, higher concurrency, cache growth, large payloads, misconfigured memory limits or heap settings, and node memory pressure. Diagnostic Information: Review the memory-usage trend for each pod, OOM-kill and restart events, garbage-collection pressure where applicable, allocation spikes, configured limits and requests, node memory pressure, traffic, and recent changes. Recovery: Reduce memory pressure by tuning caches or payloads, fixing leaks, or rolling back regressions. Scale out to reduce per-pod memory usage and cautiously adjust memory limits. The alert clears when the overload level returns to zero.

Table 5-67 SERVICE_RESOURCE_OVERLOADED

Field	Details
Description	{{labels.microservice}} service is L3 for {{labels.type}} type
Summary	{{labels.microservice}} service is L3 for {{labels.type}} type
Severity	Critical
Expression	service_resource_overload_level{type="svc_failure_count"} == 3
OID	1.3.6.1.4.1.323.5.3.37.1.2.15
Metric Used	service_resource_overload_level{type="svc_failure_count"}
Recommended Actions	Cause: The service is at overload level L1 for memory. Possible drivers include a memory leak, higher concurrency, cache growth, large payloads, misconfigured memory limits or heap settings, and node memory pressure. Diagnostic Information: Review the memory-usage trend for each pod, OOM-kill and restart events, garbage-collection pressure where applicable, allocation spikes, configured limits and requests, node memory pressure, traffic, and recent changes. Recovery: Reduce memory pressure by tuning caches or payloads, fixing leaks, or rolling back regressions. Scale out to reduce per-pod memory usage and cautiously adjust memory limits. The alert clears when the overload level returns to zero.

5.1.55 SYSTEM_IMPAIRMENT_MAJOR

Table 5-68 SYSTEM_IMPAIRMENT_MAJOR

Field	Details
Description	Major Impairment alert raised for REPLICATION_FAILED or REPLICATION_CHANNEL_DOWN or BINLOG_STORAGE usage > 80%
Summary	Major impairment alert raised for REPLICATION_FAILED or REPLICATION_CHANNEL_DOWN or BINLOG_STORAGE usage must be more than 80% for 10 minutes.
Severity	Major
Expression	(db_tier_replication_status{role="failed"} == 0) or (db_tier_replication_status{role="active"} == 0) or (count by (site_name) (db_tier_replication_status) == count by (site_name) (db_tier_replication_status{role="standby"})) or (count by (site_name) (db_tier_replication_status) == count by (site_name) (db_tier_replication_status{role="failed"})) or (avg_over_time(db_tier_binlog_used_bytes_percentage[5m]) >= 80)
OID	1.3.6.1.4.1.323.5.3.52.1.2.43
Metric Used	db_tier_replication_status and db_tier_binlog_used_bytes_percentage

Table 5-68 (Cont.) SYSTEM_IMPAIRMENT_MAJOR

Field	Details
Recommended Actions	Cause: SYSTEM_IMPAIRMENT_MAJOR indicates a major site impairment condition related to the DB tier, typically associated with replication health or binlog storage pressure. This alert is intended to fire when one or more of the following major-risk conditions occur (often mapped to REPLICATION_FAILED, REPLICATION_CHANNEL_DOWN, or BINLOG_STORAGE high usage): Diagnostic Information: - Replication failure (replication role reported as failed, or replication threads stopped) - Replication channel down / missing standby (expected replication roles are not present for the site, e.g., standby not reporting) - Primary/active not reporting (active role missing/unhealthy from metrics perspective) - Binlog storage growth due to replication lag, blocked purge, or high write rate, resulting in binlog usage 80% (as measured by db_tier_binlog_used_bytes_percentage) - Network issues between DB nodes/sites (timeouts, resets, packet loss) - DB load/resource saturation causing lag (CPU, IO, disk, connection pool) - Misconfiguration/regression after deploy/patch (replication settings, credentials, TLS/certs) - Determine which condition triggered the alert - Check current values and recent history for: - db_tier_replication_status{role="active"}, {role="standby"}, {role="failed"} - db_tier_binlog_used_bytes_percentage - Confirm whether the trigger is replication-related vs binlog-usage-related. - Replication health deep-dive - Identify impacted site_name and which node(s)/roles are missing or in failed. - Check replication lag, SQL/IO thread state, last error, and channel status on the DB nodes. - Channel/network validation - Check connectivity between active and standby. - Look for recent network maintenance, route changes, MTU issues, or service mesh/proxy impacts (if in path). - Binlog storage analysis (if binlog usage is high) - Confirm binlog filesystem utilization and growth rate; determine if purge is blocked. - Correlate with replication lag/backlog (binlogs often accumulate when standby can't consume). - Check for unusually high write volume or long-running transactions. - Resource saturation / DB health - Check disk IOPS/latency, CPU, memory pressure, and connection pool saturation. - Review DB logs for replication errors, disk-full warnings, corruption, or restart events. - Change correlation - Check recent deploys/patches/config changes (DB parameters, replication config, storage changes). Recovery: - If replication failed / channel down: - Restore replication connectivity (network path, DNS, firewall, TLS/certs). - Restart/repair replication channel as per DB runbook; resolve the root error (permissions, missing binlogs, etc.).

Table 5-68 (Cont.) SYSTEM_IMPAIRMENT_MAJOR

Field	Details
	– If standby is unhealthy, recover/rebuild standby from a known-good snapshot/ backup and rejoin replication. – Consider temporarily reducing write load or enabling throttling to allow the standby to catch up. • If binlog usage > 80%: – Treat as urgent to avoid running out of storage (risk of DB outage). – Reduce binlog growth by addressing replication lag (fix channel, improve IO/ CPU, reduce write amplification). – Expand binlog storage if needed (short-term mitigation) following capacity change procedures. – Ensure binlog purge/retention settings are correct and purge is not blocked (e.g., by lagging replicas). • Operational guidance: – Escalate as major if sustained, if multiple sites are impacted, or if lag/space continues to worsen. – After remediation, verify: – replication roles return to expected states (active/standby healthy; no failed) – replication lag stabilizes – db_tier_binlog_used_bytes_percentage drops below threshold and continues trending down

5.1.56 SYSTEM_IMPAIRMENT_CRITICAL

Table 5-69 SYSTEM_IMPAIRMENT_CRITICAL

Field	Details
Description	Critical Impairment alert raised for REPLICATION_FAILED or REPLICATION_CHANNEL_DOWN or BINLOG_STORAGE usage > 80%
Summary	Critical impairment alert raised for REPLICATION_FAILED or REPLICATION_CHANNEL_DOWN or BINLOG_STORAGE usage must be more than 80% for 30 minutes.
Severity	Critical
Expression	(db_tier_replication_status{role="failed"} == 0) or (db_tier_replication_status{role="active"} == 0) or (count by (site_name) (db_tier_replication_status) == count by (site_name) (db_tier_replication_status{role="standby"})) or (count by (site_name) (db_tier_replication_status) == count by (site_name) (db_tier_replication_status{role="failed"})) or (avg_over_time(db_tier_binlog_used_bytes_percentage[5m]) >= 80)
OID	1.3.6.1.4.1.323.5.3.52.1.2.43
Metric Used	db_tier_replication_status and db_tier_binlog_used_bytes_percentage

Table 5-69 (Cont.) SYSTEM_IMPAIRMENT_CRITICAL

Field	Details
Recommended Actions	Cause: SYSTEM_IMPAIRMENT_CRITICAL indicates a critical site impairment related to DB-tier replication health and/or binlog storage pressure. It is the higher-severity escalation of the impairment condition, intended to represent that one of the following has become sustained/critical (per description: replication failed, replication channel down, or binlog storage usage ≥ 80% for ~30 minutes). Common drivers: • Sustained replication failure (replication role reported as failed, replication threads stopped, unrecoverable SQL error) • Replication channel down for an extended period (standby missing/unreachable, broken link, auth/TLS failure) • Loss of expected topology (active/standby role metrics missing or inconsistent for a site) • Sustained binlog storage pressure (≥ 80%), typically due to replication lag/backlog, blocked purge, or abnormal write volume • Network instability between sites/DB nodes causing persistent disconnects and lag • Resource exhaustion on DB nodes (disk IO saturation, CPU, memory) preventing replication catch-up • Misconfiguration/regression after change (replication config, retention, credentials) Diagnostic Information: • Confirm the trigger and impacted site_name • Identify whether the critical is driven by replication status vs binlog usage trend. • Review current and historical values for: • db_tier_replication_status{role="active" "standby" "failed"} • db_tier_binlog_used_bytes_percentage (trend over last 30–60 minutes) • Replication failure/channel-down triage • On DB nodes, check replication channel state, IO/SQL thread status, last error, and lag. • Determine whether the standby is unreachable vs reachable-but-failing (auth/TLS/permissions vs data/SQL errors). • Confirm whether the "active" node is healthy and exporting metrics (avoid false positives from missing scrapes). • Topology consistency check • Validate expected number of replication endpoints per site (active + standby) and confirm none are missing. • If the alert logic is based on counts, verify whether a node is down or metrics are absent. • Binlog storage pressure investigation (if ≥ 80%) • Check filesystem utilization where binlogs reside and remaining headroom; estimate time-to-full. • Confirm whether binlog purge is blocked by lagging replica(s) or retention settings. • Correlate with write throughput and long transactions. • Dependency and resource saturation • Disk IO latency/queue depth, CPU, memory pressure, network throughput/packet loss. • DB error logs for disk-full warnings, replication errors, restarts, corruption indicators. • Change/event correlation • Look for recent deployments, DB patching, parameter changes, certificate rotations, or network maintenance coinciding with onset.

Table 5-69 (Cont.) SYSTEM_IMPAIRMENT_CRITICAL

Field	Details
	Recovery: • Immediate posture (CRITICAL): • Treat as incident-level due to risk of data protection loss (replication down) and/or imminent DB outage (binlog disk exhaustion). • Engage DB/platform/network on-call as required; follow site incident process. • If replication failed / channel down: • Restore connectivity (network/DNS/firewall) and fix auth/TLS/cert issues. • Restart/repair replication channels per runbook; remediate root cause (SQL errors, missing binlogs, GTID divergence). • If standby cannot be repaired quickly, rebuild standby from snapshot/backup and re-establish replication. • Consider temporarily throttling write load to reduce lag/backlog while recovering. • If binlog usage 80% (sustained): • Prevent disk-full: • reduce binlog generation (traffic/load throttling where feasible) • restore replication so binlogs can be consumed and purged • expand binlog storage/capacity as an emergency mitigation (per change controls) • Verify purge/retention settings and ensure purge is not blocked by an unreachable or stuck replica. • Validation or exit criteria: • Replication status returns to expected healthy roles (no failed, standby present and catching up). • Replication lag stabilizes within acceptable bounds. • db_tier_binlog_used_bytes_percentage trends downward and remains below threshold.

5.1.57 SYSTEM_OPERATIONAL_STATE_PARTIAL_SHUTDOWN

Table 5-70 SYSTEM_OPERATIONAL_STATE_PARTIAL_SHUTDOWN

Field	Details
Description	System Operational State is now in partial shutdown state.
Summary	System Operational State is now in partial shutdown state.
Severity	Major
Expression	system_operational_state == 2
OID	1.3.6.1.4.1.323.5.3.52.1.2.44
Metric Used	system_operational_state == 2

Table 5-70 (Cont.) SYSTEM_OPERATIONAL_STATE_PARTIAL_SHUTDOWN

Field	Details
Recommended Actions	Cause: The site has entered the PARTIAL SHUTDOWN operational state under the Controlled Shutdown feature. This deliberate isolation mode supports maintenance, such as upgrades or configuration changes, or recovery activities, such as prolonged replication-channel failures or replication delay, while allowing existing sessions to continue. In this state, NRF status is SUSPENDED to prevent consumer discovery; the Ingress Gateway and Diameter Gateway permit only in-session messages and reject new session establishment; and the Egress Gateway has no restriction. This state is typically set through the CM GUI or REST API when enableControlledShutdown is enabled. Diagnostic Information: • Confirm the operational state is truly 2 (PARTIAL) and record transition time (use alert timestamp + metric history). • Verify whether the change was planned (maintenance/recovery activity) and who/what initiated it (CM GUI/REST). • Validate enableControlledShutdown is enabled and configuration is consistent across the site. • Verify expected gateway behavior: • Ingress rejects new session establishment but permits in-session traffic • Diameter permits in-session only • Check NRF to confirm NF status marked SUSPENDED and are not being discovered for new sessions. • Identify the underlying reason for isolation (maintenance, replication channel failures/delay, dependency instability) and correlate with replication metrics/logs. • Assess customer impact: rate of blocked new sessions, admission control counters, error codes returned at ingress/diameter. Recovery: For a planned shutdown, proceed with maintenance or recovery activities, such as replication repair or dependency stabilization, while keeping existing sessions stable. For an unplanned shutdown, confirm that no automation or configuration error triggered the change, correct configuration or access controls, and return to NORMAL when safe. After recovery, verify that NFs return to REGISTERED in NRF and that new session establishment succeeds at Ingress and Diameter. Escalate if the partial shutdown persists unexpectedly or blocks new sessions beyond agreed SLOs.

5.1.58 SYSTEM_OPERATIONAL_STATE_COMPLETE_SHUTDOWN

Table 5-71 SYSTEM_OPERATIONAL_STATE_COMPLETE_SHUTDOWN

Field	Details
Description	System Operational State is now in complete shutdown state
Summary	System Operational State is now in complete shutdown state
Severity	Critical
Expression	system_operational_state == 3
OID	1.3.6.1.4.1.323.5.3.52.1.2.44
Metric Used	system_operational_state

Table 5-71 (Cont.) SYSTEM_OPERATIONAL_COMPLETE_SHUTDOWN

Field	Details
Recommended Actions	Cause: The site has entered the COMPLETE SHUTDOWN operational state under the Controlled Shutdown feature. This fully stops load and isolates the site during recovery or severe instability. NRF status is SUSPENDED, the Ingress Gateway accepts no requests, the Egress Gateway allows only NRF-bound requests, and Diameter Gateway external-peer connections are closed. Diagnostic Information: • Confirm the operational state is 3 (COMPLETE) and capture the exact change time and duration. • Verify whether complete isolation was intended (approved maintenance/incident action) and confirm initiator. • Validate actual enforcement at gateways: • Ingress is refusing all new connections • Diameter external peer sessions are closed • Egress allows only NRF-bound traffic • Confirm NRF state: NFs remain SUSPENDED and are not discoverable by consumers (SMF/AMF/etc.). • Identify the triggering condition that required complete shutdown (e.g., severe replication failure, corruption risk, cascading dependency failures) and gather supporting logs/metrics. • Confirm traffic has drained/failover behavior is functioning (upstream routing, peer sites capacity, no looping retries). • Check for any secondary effects: backlog growth, message queue accumulation, alarms from consumers due to loss of discovery. Recovery: Treat an unscheduled shutdown as an incident. Perform the maintenance or recovery procedure while the site is isolated, such as replication rebuild or data-consistency checks. When ready, transition to PARTIAL if required for validation, or directly to NORMAL; restore gateway connectivity and Egress behavior; update NRF status to REGISTERED; and validate discovery, new-session establishment, and clean Diameter-peer reconnection.

5.1.59 DIAM_CONN_PEER_DOWN

Table 5-72 DIAM_CONN_PEER_DOWN

Field	Details
Description	Diameter connection to peer {{ \$labels.peerHost }} is down.
Summary	Diameter connection to peer down.
Severity	Major
Expression	(sum by (namespace,peerHost) (ocbsf_diam_conn_network) == 0) and (sum by (namespace,peerHost) (max_over_time(ocbsf_diam_conn_network[24h])) != 0)
OID	1.3.6.1.4.1.323.5.3.37.1.2.18
Metric Used	ocbsf_diam_conn_network

Table 5-72 (Cont.) DIAM_CONN_PEER_DOWN

Field	Details
Recommended Actions	Cause: The Diameter Gateway network-side connection to a specific external peer is down, although it was connected at least once during the preceding 24 hours. Common causes include peer downtime or maintenance, a peer-specific network-path issue, transport instability, TLS or mTLS handshake failures, peer overload, configuration mismatches, or gateway resource pressure. Diagnostic Information: Identify the affected <code>kubernetes_namespace</code> and <code>origHost</code>, and determine whether one or multiple peers are affected. Confirm that <code>max_over_time(occnp_diam_conn_network[24h])</code> is non-zero for the same peer. Review Diameter Gateway logs for watchdog failures, disconnect procedures, CER or CEA failures, TCP errors, and TLS validation failures. Check peer health, listener availability, logs, overload indicators, routing, firewall, NAT, DNS, recent changes, and TLS certificate, trust-bundle, hostname, SNI, protocol, and cipher configuration. Recovery: Restore peer availability and repair peer-specific routing, firewall, or DNS issues. If TLS is enabled, renew or rotate certificates, update trust bundles, and correct hostname, SNI, cipher, or protocol settings. Address overload or instability by scaling or reducing load on the gateway or peer. Verify that <code>occnp_diam_conn_network</code> for the affected <code>origHost</code> is greater than zero, the connection is stable, and Diameter transaction success rates return to normal.

5.1.60 DIAM_CONN_NETWORK_DOWN

Table 5-73 DIAM_CONN_NETWORK_DOWN

Field	Details
Description	All diameter network connections are down.
Summary	All diameter network connections are down.
Severity	Critical
Expression	<code>sum by (namespace)(ocbsf_diam_conn_network) == 0</code>
OID	1.3.6.1.4.1.323.5.3.37.1.2.19
Metric Used	<code>ocbsf_diam_conn_network</code>

Table 5-73 (Cont.) DIAM_CONN_NETWORK_DOWN

Field	Details
Recommended Actions	Cause: The Diameter Gateway has no active external Diameter-network connections for a <code>kubernetes_namespace</code>. Common causes include an external-peer outage or maintenance, network-path failures, TLS or mTLS issues, peer-configuration mismatches, transport instability, gateway overload, or controlled shutdown or site isolation. Diagnostic Information: Confirm the affected namespace and whether the outage is isolated or widespread. Identify affected peers and determine whether all external peers are down. Review Diameter Gateway logs for watchdog failures, disconnect procedures, CER or CEA failures, TCP errors, and TLS errors. Validate peer health, listener availability, maintenance status, routing, firewall, NAT, DNS, and recent network changes. If TLS is enabled, verify certificate validity and chain, truststore or CA alignment, hostname or SNI matching, and protocol or cipher compatibility. Correlate the alert with backend-connection alerts and controlled-shutdown status. Recovery: Restore external-peer availability and network connectivity. If TLS-related, renew or rotate certificates, update truststores or CA bundles, correct hostname, SNI, cipher, or protocol mismatches, and re-establish sessions. Correct Diameter-peer configuration and address gateway resource pressure or instability. Verify that <code>sum by (kubernetes_namespace) (ocnp_diam_conn_network)</code> is greater than zero, external-peer sessions are stable, and end-to-end Diameter transactions succeed.

5.1.61 DIAM_RESPONSE_REALM_VALIDATION_ERROR_CRITICAL

Table 5-74 DIAM_RESPONSE_REALM_VALIDATION_ERROR_CRITICAL

Field	Details
Description	At least 75% of the Diam Response failed with error 'DIAMETER_REALM_NOT_SERVED', either of BSF realm or PCF Realm doesn't match with received destination realm in diameter message.
Summary	{{ \$value }}% of the Diam Response failed with error 'DIAMETER_REALM_NOT_SERVED'.
Severity	CRITICAL
Expression	$(\text{sum by (namespace, microservice)} (\text{increase(ocbsf_diam_realm_validation_failed_total}{responseCode="3003", appId="16777236"}[10m])) / \text{sum by (namespace, microservice)} (\text{increase(ocbsf_diam_response_network_total}{appId="16777236"}[10m])) * 100 \geq 75$

Table 5-74 (Cont.) DIAM_RESPONSE_REALM_VALIDATION_ERROR_CRITICAL

Field	Details
OID	1.3.6.1.4.1.323.5.3.37.1.2.41
Metric Used	ocbsf_diam_realm_validation_failed_total ocbsf_diam_response_network_total
Recommended Actions	1. Check if the value of the following keys under Advanced settings of diameter settings page are set to true: • DIAMETER.Enable.Validate.Realm • DIAMETER.BSF.Enable.Validate.Binding.Realm 2. Check the destination-realm in diameter request.

5.1.62 DIAM_RESPONSE_REALM_VALIDATION_ERROR_MAJOR

Table 5-75 DIAM_RESPONSE_REALM_VALIDATION_ERROR_MAJOR

Field	Details
Description	At least 50% of the Diam Response failed with error 'DIAMETER_REALM_NOT_SERVED', either of BSF realm or PCF Realm doesn't match with received destination realm in diameter message.
Summary	{{ \$value }}% of the Diam Response failed with error 'DIAMETER_REALM_NOT_SERVED'.
Severity	MAJOR
Expression	(sum by (namespace, microservice) (increase(ocbsf_diam_realm_validation_failed_total{responseCode="3003", appld="16777236"}[10m])) / sum by (namespace, microservice) (increase(ocbsf_diam_response_network_total{appld="16777236"}[10m]))) * 100 >= 50
OID	1.3.6.1.4.1.323.5.3.37.1.2.41
Metric Used	ocbsf_diam_realm_validation_failed_total ocbsf_diam_response_network_total
Recommended Actions	1. Check if the value of the following keys under Advanced settings of diameter settings page are set to true: • DIAMETER.Enable.Validate.Realm • DIAMETER.BSF.Enable.Validate.Binding.Realm 2. Check the destination-realm coming in diameter request.

5.1.63 DIAM_RESPONSE_REALM_VALIDATION_ERROR_MINOR

Table 5-76 DIAM_RESPONSE_REALM_VALIDATION_ERROR_MINOR

Field	Details
Description	At least 20% of the Diam Response failed with error 'DIAMETER_REALM_NOT_SERVED', either of BSF realm or PCF Realm doesn't match with received destination realm in diameter message.
Summary	{{ \$value }}% of the Diam Response failed with error 'DIAMETER_REALM_NOT_SERVED'.
Severity	MINOR
Expression	(sum by (namespace, microservice) (increase(ocbsf_diam_realm_validation_failed_total{responseCode="3003", appId="16777236"}[10m])) / sum by (namespace, microservice) (increase(ocbsf_diam_response_network_total{appId="16777236"}[10m]))) * 100 >= 20
OID	1.3.6.1.4.1.323.5.3.37.1.2.41
Metric Used	ocbsf_diam_realm_validation_failed_total ocbsf_diam_response_network_total
Recommended Actions	1. Check if the value of the following keys under Advanced settings of diameter settings page are set to true: • DIAMETER.Enable.Validate.Realm • DIAMETER.BSF.Enable.Validate.Binding.Realm 2. Check the destination-realm coming in diameter request.

5.1.64 AUDIT_STALE_NOTIFY_ERROR_RESPONSE_MINOR

Table 5-77 AUDIT_STALE_NOTIFY_ERROR_RESPONSE_MINOR

Field	Details
Description	At least 20 % of the BSF Audit Notification Requests sent to PCF to check for Suspected Stale Sessions have responded with a 5xx or 4xx (excluding 404) Status in the last 10 minutes.
Summary	At least 20% of the BSF Notification Request for Audit have responded with a 5xx or 4xx (not 404) Status in the last 10 minutes.
Severity	MINOR
Expression	(sum by (namespace, microservice) (increase(ocbsf_query_response_count_total{response_code=~"5xx 4xx timeout"}[10m])) / sum by (namespace, microservice) (increase(ocbsf_query_response_count_total[10m])) * 100 >= 20 < 40
OID	1.3.6.1.4.1.323.5.3.37.1.2.42

Table 5-77 (Cont.) AUDIT_STALE_NOTIFY_ERROR_RESPONSE_MINOR

Field	Details
Metric Used	ocbsf_query_response_count_total
Recommended Actions	Cause: More than 20% and less than 40% of BSF audit notification requests sent to PCF to check suspected stale sessions returned 4xx, 5xx, or timeout responses in the last 10 minutes. Diagnostic Information: Verify whether PCF returns 4xx, 5xx, or timeout responses for the BSF PCF bindings used by audit query or notification requests. Recovery: The alert clears when BSF audit requests to PCF receive a successful 2xx response or another HTTP error response.

5.1.65 AUDIT_STALE_NOTIFY_ERROR_RESPONSE_MAJOR

Table 5-78 AUDIT_STALE_NOTIFY_ERROR_RESPONSE_MAJOR

Field	Details
Description	At least 40 % of the BSF Audit Notification Requests sent to PCF to check for Suspected Stale Sessions have responded with a 5xx or 4xx (excluding 404) Status in the last 10 minutes.
Summary	{{ \$value }} % of the BSF Audit Notification Requests sent to PCF to check for Suspected Stale Sessions have responded with a 5xx or 4xx (excluding 404) Status in the last 10 minutes.
Severity	MAJOR
Expression	(sum by (namespace, microservice) (increase(ocbsf_query_response_count_total{response_code=~"5xx 4xx timeout"}[10m])) / sum by (namespace, microservice) (increase(ocbsf_query_response_count_total[10m]))) * 100 >= 40 < 60
OID	1.3.6.1.4.1.323.5.3.37.1.2.42
Metric Used	ocbsf_query_response_count_total
Recommended Actions	Cause: More than 40% and less than 60% of BSF audit notification requests sent to PCF to check suspected stale sessions returned 4xx, 5xx, or timeout responses in the last 10 minutes. Diagnostic Information: Verify whether PCF returns 4xx, 5xx, or timeout responses for the BSF PCF bindings used by audit query or notification requests. Recovery: The alert clears when BSF audit requests to PCF receive a successful 2xx response or another HTTP error response.

5.1.66 AUDIT_STALE_NOTIFY_ERROR_RESPONSE_CRITICAL

Table 5-79 AUDIT_STALE_NOTIFY_ERROR_RESPONSE_CRITICAL

Field	Details
Description	At least 60 % of the BSF Audit Notification Requests sent to PCF to check for Suspected Stale Sessions have responded with a 5xx or 4xx (excluding 404) Status in the last 10 minutes.
Summary	At least 60% of the BSF Notification Request for Audit to PCF (or its respective NF) failed with a 5xx or 4xx (not 404) Status in the last 10 minutes. The threshold default value is defined at <code>BSF_Alertrules.yaml</code> .
Severity	CRITICAL
Expression	(sum by (namespace, microservice) (increase(ocbsf_query_response_count_total{response_code=~"5xx 4xx timeout"}[10m])) / sum by (namespace, microservice) (increase(ocbsf_query_response_count_total[10m]))) * 100 >= 60
OID	1.3.6.1.4.1.323.5.3.37.1.2.42
Metric Used	ocbsf_query_response_count_total
Recommended Actions	Cause: More than 60% of BSF audit notification requests sent to PCF to check suspected stale sessions returned 4xx, 5xx, or timeout responses in the last 10 minutes. Diagnostic Information: Verify whether PCF returns 4xx, 5xx, or timeout responses for the BSF PCF bindings used by audit query or notification requests. Recovery: The alert clears when BSF audit requests to PCF receive a successful 2xx response or another HTTP error response.

5.1.67 BSF_CONNECTION_FAILURE

Table 5-80 BSF_CONNECTION_FAILURE

Field	Details
Description	Connection failure on Egress and Ingress Gateways for incoming and outgoing connections.
Summary	Connection failure on Egress and Ingress Gateways for incoming and outgoing connections.
Severity	Major

Table 5-80 (Cont.) BSF_CONNECTION_FAILURE

Field	Details
Expression	sum(increase(ocbsf_oc_ingressgateway_connection_failure_total[5m]) >0 or (ocbsf_oc_ingressgateway_connection_failure_total unless ocbsf_oc_ingressgateway_connection_failure_total offset 5m)) by (namespace,app, error_reason) > 0 or sum(increase(ocbsf_oc_egressgateway_connection_failure_total[5m]) >0 or (ocbsf_oc_egressgateway_connection_failure_total unless ocbsf_oc_egressgateway_connection_failure_total offset 5m)) by (namespace,app, error_reason) > 0
OID	1.3.6.1.4.1.323.5.3.37.1.2.43
Metric Used	ocbsf_oc_ingressgateway_connection_failure_total ocbsf_oc_egressgateway_connection_failure_total
Recommended Actions	For any additional guidance, contact My Oracle Support (https://support.oracle.com).

5.1.68 INGRESS_GATEWAY_DD_UNREACHABLE_MAJOR

Table 5-81 INGRESS_GATEWAY_DD_UNREACHABLE_MAJOR

Field	Details
Description	'BSF Ingress Gateway Data Director unreachable for {{\$labels.namespace}}'
Summary	'kubernetes_namespace: {{\$labels.kubernetes_namespace}}, timestamp: {{ with query "time()" }}{{ . first value humanizeTimestamp }}{{ end }} BSF Ingress Gateway Data Director unreachable'
Severity	Major
Expression	sum(oc_ingressgateway_dd_unreachable) by(namespace,container) > 0
OID	1.3.6.1.4.1.323.5.3.37.1.2.48
Metric Used	oc_ingressgateway_dd_unreachable
Recommended Actions	

5.1.69 EGRESS_GATEWAY_DD_UNREACHABLE_MAJOR

Table 5-82 EGRESS_GATEWAY_DD_UNREACHABLE_MAJOR

Field	Details
Description	'BSF Egress Gateway Data Director unreachable for {{\$labels.namespace}}'

Table 5-82 (Cont.) EGRESS_GATEWAY_DD_UNREACHABLE_MAJOR

Field	Details
Summary	'kubernetes_namespace: {{ \$labels.kubernetes_namespace }}, timestamp: {{ with query "time()" }}{ . first value humanizeTimestamp }}{ end }} BSF Egress Gateway Data Director unreachable'
Severity	Major
Expression	sum(oc_egressgateway_dd_unreachable) by(namespace,container) > 0
OID	1.3.6.1.4.1.323.5.3.37.1.2.49
Metric Used	oc_egressgateway_dd_unreachable
Recommended Actions	Cause: No cause information is provided in the Confluence Resolutions/Workarounds column. Diagnostic Information: No diagnostic information is provided in the Confluence Resolutions/Workarounds column. Recovery: No recovery information is provided in the Confluence Resolutions/Workarounds column.

5.1.70 DIAM_GATEWAY_CERTIFICATE_EXPIRY_MAJOR

Table 5-83 DIAM_GATEWAY_CERTIFICATE_EXPIRY_MAJOR

Field	Details
Description	Diam-gw certificate expiry in less than 3 months for {{ \$labels.namespace }}.
Summary	Diam-gw certificate expiry in less than 3 months.
Severity	Major
Expression	dgw_tls_cert_expiration_seconds - time() <= 7862400
OID	1.3.6.1.4.1.323.5.3.37.1.2.47
Metric Used	dgw_tls_cert_expiration_seconds
Recommended Actions	Cause: A Diameter Gateway TLS certificate is approaching expiry with increased urgency. Diagnostic Information: Identify the expiring certificate and affected instance from available metric labels. Confirm that dependent peers trust the issuing CA and that intermediate certificates are correctly configured. Recovery: Execute the certificate-rotation plan. Deploy the new certificate and key through the approved secret-management process, then reload or restart the Diameter Gateway according to the runbook. Verify successful TLS handshakes with peers.

5.1.71 DIAM_GATEWAY_CERTIFICATE_EXPIRY_CRITICAL

Table 5-84 DIAM_GATEWAY_CERTIFICATE_EXPIRY_CRITICAL

Field	Details
Description	Diam-gw certificate expiry in less than a month for <code>{{\$labels.namespace}}</code> .
Summary	Diam-gw certificate expiry in less than a month.
Severity	Critical
Expression	<code>dgw_tls_cert_expiration_seconds - time() <= 2592000</code>
OID	1.3.6.1.4.1.323.5.3.37.1.2.47
Metric Used	<code>dgw_tls_cert_expiration_seconds</code>
Recommended Actions	Cause: A Diameter Gateway certificate is close to expiry, creating a high risk of a production outage when it expires. Diagnostic Information: Verify the exact certificate-expiry timestamp and determine whether multiple instances are affected. Recovery: Expedite certificate renewal or rotation. After deployment, verify peer-connection stability without TLS alerts or handshake failures, confirm that Diameter Gateway logs show the new certificate has loaded, and confirm that the metric resets to a longer time to expiry.

5.1.72 DGW_TLS_CONNECTION_FAILURE

Table 5-85 DGW_TLS_CONNECTION_FAILURE

Field	Details
Description	<code>{{\$labels.reason}}</code> failure when Diam gateway is an initiator.
Summary	TLS Connection failure when Diam gateway is an initiator.
Severity	Major
Expression	<code>sum by (namespace,reason) (ocbsf_diam_failed_conn_network) > 0</code>
OID	1.3.6.1.4.1.323.5.3.37.1.2.50.
Metric Used	<code>ocbsf_diam_failed_conn_network</code>

Table 5-85 (Cont.) DGW_TLS_CONNECTION_FAILURE

Field	Details
Recommended Actions	Cause: Diameter Gateway is failing to establish TLS connections as a client because of handshake exceptions. Common causes include expired or rotated certificates, a missing intermediate certificate chain, truststore mismatch or an unknown CA, SAN or CN mismatch, TLS policy or cipher mismatch, clock skew, or network middlebox interference. Diagnostic Information: Review failures by namespace and reason using <code>sum by(namespace, reason) (ocnp_diam_failed_conn_network)</code>. If peer or endpoint labels are available, isolate the affected target. Correlate the issue with certificate-expiry alerts, validate NTP synchronization on Diameter Gateway nodes, and review Diameter Gateway logs for <code>unknown_ca</code>, <code>bad_certificate</code>, <code>certificate_expired</code>, <code>handshake_failure</code>, or <code>protocol_version</code> errors. Recovery: Renew or rotate the Diameter Gateway client certificate, update trust stores or CA bundles including intermediate certificates, and reload or restart Diameter Gateway according to the runbook. Align TLS versions and ciphers on both sides. If the issue is endpoint-specific, remediate the server certificate or configuration, or fail over to an alternate server or peer. Confirm that handshake failures stop increasing and peer connections remain established.

5.1.73 BINDING_REVALIDATION_PCF_BINDING_MISSING_MINOR

Table 5-86 BINDING_REVALIDATION_PCF_BINDING_MISSING_MINOR

Field	Details
Description	At least 30% but less than 50% of the PCF BINDING missing among all binding revalidation records in the last 5 minutes.
Summary	At least 30% but less than 50% of the PCF BINDING missing among all Binding Revalidation records in the last 5 minutes.
Severity	Minor
Expression	<code>(sum by (namespace) (rate(ocbsf_binding_revalidation_pcfBinding_missing_total[5m])) / sum by (namespace) (rate(ocbsf_binding_revalidation_response_total[5m]))) * 100 >= 30 < 50</code>
OID	1.3.6.1.4.1.323.5.3.37.1.2.51
Metric Used	<code>ocbsf_binding_revalidation_pcfBinding_missing_total</code>

Table 5-86 (Cont.) BINDING_REVALIDATION_PCF_BINDING_MISSING_MINOR

Field	Details
Recommended Actions	Check BSF Management service health history. Increase binding audit frequency. For any additional guidance, contact My Oracle Support (https://support.oracle.com).

5.1.74 BINDING_REVALIDATION_PCF_BINDING_MISSING_MAJOR

Table 5-87 BINDING_REVALIDATION_PCF_BINDING_MISSING_MAJOR

Field	Details
Description	At least 50% but less than 70% of the PCF BINDING missing among all binding revalidation records in the last 5 minutes.
Summary	At least 50% but less than 70% of the PCF BINDING missing among all binding revalidation records in the last 5 minutes.
Severity	Major
Expression	(sum by (namespace) (rate(ocbsf_binding_revalidation_pcfBinding_missing_total[5m])) / sum by (namespace) (rate(ocbsf_binding_revalidation_response_total[5m]))) * 100 >= 50 < 70
OID	1.3.6.1.4.1.323.5.3.37.1.2.51
Metric Used	ocbsf_binding_revalidation_pcfBinding_missing_total
Recommended Actions	Check BSF Management service health history. Increase binding audit frequency. For any additional guidance, contact My Oracle Support (https://support.oracle.com).

5.1.75 BINDING_REVALIDATION_PCF_BINDING_MISSING_CRITICAL

Table 5-88 BINDING_REVALIDATION_PCF_BINDING_MISSING_CRITICAL

Field	Details
Description	At least 70% of the PCF BINDING missing among all binding revalidation records in the last 5 minutes.
Summary	At least 70% of the PCF BINDING missing among all binding revalidation records in the last 5 minutes.
Severity	Critical
Expression	(sum by (namespace) (rate(ocbsf_binding_revalidation_pcfBinding_missing_total[5m])) / sum by (namespace) (rate(ocbsf_binding_revalidation_response_total[5m]))) * 100 >= 70

Table 5-88 (Cont.) BINDING_REVALIDATION_PCF_BINDING_MISSING_CRITICAL

Field	Details
OID	1.3.6.1.4.1.323.5.3.37.1.2.51
Metric Used	ocbsf_binding_revalidation_pcfBinding_missing_total
Recommended Actions	Check BSF Management service health history. Increase binding audit frequency. For any additional guidance, contact My Oracle Support (https://support.oracle.com).

5.1.76 BSF_STATE_NON_FUNCTIONAL_CRITICAL

Table 5-89 BSF_STATE_NON_FUNCTIONAL_CRITICAL

Field	Details
Description	BSF is in non functional state due to DB Cluster state down
Summary	BSF is in non functional state due to DB Cluster state down
Severity	Critical
Expression	appinfo_nfDbFunctionalState_current{nfDbFunctionalState="Not_Running"} == 1
OID	1.3.6.1.4.1.323.5.3.37.1.2.56
Metric Used	appinfo_nfdbfunctionalstate_current
Recommended Actions	Cause: The alert is raised because the BSF network function is non-functional due to the database cluster being down. Diagnostic Information: System monitoring indicates that the database cluster state is "Not Running" and is unreachable, preventing the BSF network function from operating normally. Recovery: Check and restore the database cluster to a running state. After recovery, verify that the BSF network function returns to operational status. Escalate to database administration if the issue persists.

5.1.77 POD_PENDING_REQUEST_CONGESTION_L1

Table 5-90 POD_PENDING_REQUEST_CONGESTION_L1

Field	Details
Description	Pod Resource Congestion status of {{\$labels.microservice}} service is Congestion_L1 for resource type queue.

Table 5-90 (Cont.) POD_PENDING_REQUEST_CONGESTION_L1

Field	Details
Summary	Pod Resource Congestion status of <code>{{labels.microservice}}</code> service is Congestion_L1 for resource type queue.
Severity	Critical
Expression	<code>ocbsf_pod_resource_congestion_state{type="queue"} == 2</code>
OID	1.3.6.1.4.1.323.5.3.37.1.2.54
Metric Used	<code>ocbsf_pod_resource_congestion_state</code>
Recommended Actions	Cause: The pending HTTP request queue has reached CONGESTION_L1. Diagnostic Information: Confirm that <code>pod_resource_congestion_state{resourceType="queue"} == 2</code>. Verify queue-related <code>pod_resource_stress</code> values and review rejections by request URI and request method at L1. Recovery: Ensure that L1 queue thresholds are correct. If queues continue to grow, raise the L1 discard priority so that lower-priority requests are rejected earlier.

5.1.78 POD_PENDING_REQUEST_CONGESTION_L2

Table 5-91 POD_PENDING_REQUEST_CONGESTION_L2

Field	Details
Description	Pod Resource Congestion status of <code>{{labels.microservice}}</code> service is Congestion_L2 for resource type queue.
Summary	Pod Resource Congestion status of <code>{{labels.microservice}}</code> service is Congestion_L2 for resource type queue.
Severity	Critical
Expression	<code>ocbsf_pod_resource_congestion_state{type="queue"} == 3</code>
OID	1.3.6.1.4.1.323.5.3.37.1.2.55
Metric Used	<code>ocbsf_pod_resource_congestion_state</code>
Recommended Actions	Cause: The pending HTTP request queue has reached CONGESTION_L2. Diagnostic Information: Confirm that <code>pod_resource_congestion_state{resourceType="queue"} == 3</code>. Examine queue-stress trends and rejection counters for low-priority traffic at L2. Recovery: Increase shedding at L2 by raising the discard priority, and review queue thresholds to prevent saturation.

5.1.79 AUDIT_NOT_RUNNING

Table 5-92 AUDIT_NOT_RUNNING

Field	Details
Description	Audit has not been running for at least 1 hour in pod <code>{{labels.pod}}</code> .
Summary	Audit has been stuck in an unhealthy state for over 1 hour.
Severity	Critical
Expression	<code>(increase(data_repository_invocations_seconds_count{method="getQueuedTablesToAudit",state="SUCCESS"}[1h])) == 0</code>
OID	1.3.6.1.4.1.323.5.3.37.1.2.45
Metric Used	<code>data_repository_invocations_seconds_count</code>
Recommended Actions	Cause: No successful executions of <code>getQueuedTablesToAudit</code> have been recorded for at least one hour, as indicated by <code>increase(...[1h]) == 0</code>. This typically indicates that the audit job is not running, is stuck, or is failing before it records a successful execution. Diagnostic Information: Check in Prometheus or Grafana: - Graph <code>data_repository_invocations_seconds_count{method="getQueuedTablesToAudit"}</code> by state to determine whether attempts are failing, for example with <code>ERROR</code> or <code>FAILURE</code>, or whether no attempts are occurring. - Confirm that the metric series exists for the pod. If it is missing, verify Prometheus scrape health and target status for that pod. Check pod health and logs for the previous hour: - Check for recent restarts, OOM kills, and readiness or liveness failures. - Review application logs for audit scheduler or worker-thread issues, database connectivity problems, timeouts, deadlocks, or queue starvation. Recovery: - If the pod is unhealthy or the audit worker is stuck, restart or roll the pod according to the team's approved change and restart procedures. - If non-success states are present, remediate the underlying error, for example by restoring database connectivity, correcting credentials, resolving timeouts or locks, or clearing blocking conditions. Then verify that successful invocations resume. - If metrics are missing while the service is healthy, correct metrics emission or scraping configuration, including the Prometheus target, ServiceMonitor, or endpoint authentication. Then confirm that the counter is increasing. - Validate recovery by confirming that the metric increases over a short interval, for example <code>increase(...[5m]) > 0</code>, and that audit backlog and health indicators return to normal.

5.1.80 BSF_PCF_BINDING_TABLE_MIGRATED_PERCENTAGE

Table 5-93 BSF_PCF_BINDING_TABLE_MIGRATED_PERCENTAGE

Field	Details
Description	Pcf binding table migration configuration should be updated to only use the pcf binding v2 table
Summary	Pcf binding table migration configuration should be updated to only use the pcf binding v2 table
Severity	Minor
Expression	max_over_time(((100 * sum by (namespace) (oc_db_active_session_count{Service="bsf-management-service",Table="pcf_binding_v2"}) / clamp_min(sum by (namespace) (oc_db_active_session_count{Service="bsf-management-service",Table=~"pcf_binding pcf_binding_v2"}),1) == 100) and (100 * sum by (namespace) (oc_db_active_session_count{Service="bsf-management-service",Table="pcf_binding_v2"} offset 1m) / clamp_min(sum by (namespace) (oc_db_active_session_count{Service="bsf-management-service",Table=~"pcf_binding pcf_binding_v2"} offset 1m),1) < 100))[3h:]) == 100
OID	1.3.6.1.4.1.323.5.3.37.1.2.57
Metric Used	oc_db_active_session_count
Recommended Actions	Cause: The alert is raised because all BSF pcf binding records in legacy v1 on current site have been migrated to pcf binding v2 table. Diagnostic Information: Verify that pcf_binding table is empty and transition Advanced Settings PCF_BINDING_TABLE_LOOKUP to 3. Recovery: No recovery steps are needed as it indicates to move into a migration complete status. Alert is cleared after 24 hours.

5.1.81

BSF_PCF_BINDING_TABLE_MIGRATION_INVALID_CONFIGURATION

Table 5-94 BSF_PCF_BINDING_TABLE_MIGRATION_INVALID_CONFIGURATION

Field	Details
Description	Pcf binding table migration configuration should be reviewed and updated to a valid configuration, invalid configurations: {{labels.incompatibleFeatures}}. Examples of configurations that could trigger the alert: • ENABLE_PCF_BINDING_TABLE_MIGRATION :true, PCF_BINDING_TABLE_LOOKUP : 0 • ENABLE_PCF_BINDING_TABLE_MIGRATION :false, PCF_BINDING_TABLE_LOOKUP : 3
Summary	Pcf binding table migration, invalid configuration was set, latest valid values are used.
Severity	Critical
Expression	ocbsf_feature_incompatibility == 1
OID	1.3.6.1.4.1.323.5.3.37.1.2.58

Table 5-94 (Cont.)
BSF_PCF_BINDING_TABLE_MIGRATION_INVALID_CONFIGURATION

Field	Details
Metric Used	ocbsf_feature_incompatibility
Recommended Actions	Cause: The alert is raised because the current configuration for Remove Index Based Lookup feature is having an incorrect combination for ENABLE_PCF_BINDING_TABLE_MIGRATION and PCF_BINDING_TABLE_LOOKUP_VALUE. Diagnostic Information: Verify configuration is valid according to the following rules: 1. If ENABLE_PCF_BINDING_TABLE_MIGRATION is false, PCF_BINDING_TABLE_LOOKUP_VALUE can only have value 0. 2. If ENABLE_PCF_BINDING_TABLE_MIGRATION is true, PCF_BINDING_TABLE_LOOKUP_VALUE can only have values between 1-3. Recovery: Alert is cleared once the configuration is updated to a valid configuration.

5.1.82 CERTIFICATE_EXPIRY_MINOR

Table 5-95 CERTIFICATE_EXPIRY_MINOR

Field	Details
Description	Certificate expiry in less than 6 months for {{\$labels.namespace}}
Summary	Certificate expiry in less than 6 months
Severity	Minor
Expression	security_cert_x509_expiration_seconds - time() <= 15724800
OID	1.3.6.1.4.1.323.5.3.37.1.2.44
Metric Used	security_cert_x509_expiration_seconds
Recommended Actions	Cause: A certificate monitored by the security exporter is approaching expiry. Diagnostic Information: Identify the certificate with the shortest remaining validity and the affected resources. For example, review <code>min by (namespace, secret, cert_subject) (security_cert_x509_expiration_seconds - time())</code>. Recovery: Plan renewal or rotation through the approved certificate or PKI process. Update Kubernetes secrets or keystores, schedule the rollout before the major-alert window, and confirm that the expiration metric increases after rotation.

5.1.83 CERTIFICATE_EXPIRY_MAJOR

Table 5-96 CERTIFICATE_EXPIRY_MAJOR

Field	Details
Description	Certificate expiry in less than 3 months for <code>{{labels.namespace}}</code>
Summary	Certificate expiry in less than 3 months.
Severity	Major
Expression	<code>security_cert_x509_expiration_seconds - time() <= 7862400</code>
OID	1.3.6.1.4.1.323.5.3.37.1.2.44
Metric Used	<code>security_cert_x509_expiration_seconds</code>
Recommended Actions	Cause: A certificate is approaching expiry with increased urgency. Diagnostic Information: Identify the workloads that use the certificate, such as ingress, egress, internal services, and databases. Recovery: Rotate the certificate and confirm that dependent clients trust the new certificate chain. Consider a staged rollout to avoid widespread TLS handshake failures.

5.1.84 CERTIFICATE_EXPIRY_CRITICAL

Table 5-97 CERTIFICATE_EXPIRY_CRITICAL

Field	Details
Description	Certificate expiry in less than a month for <code>{{labels.namespace}}</code>
Summary	Certificate expiry in less than a month.
Severity	Critical
Expression	<code>security_cert_x509_expiration_seconds - time() <= 2592000</code>
OID	1.3.6.1.4.1.323.5.3.37.1.2.44
Metric Used	<code>security_cert_x509_expiration_seconds</code>
Recommended Actions	Cause: A certificate is close to expiry, creating a high likelihood of service disruption when it expires. Diagnostic Information: Validate the exact expiry time and determine the blast radius by identifying the services and peers that depend on the certificate. Recovery: Perform an emergency certificate rotation. Renew the certificate, deploy it through the approved secrets process, and reload or restart affected workloads according to the runbook. Confirm successful TLS handshakes, the absence of connection-failure alerts, and reset of the expiration metric.

5.1.85 AUDIT_STALE_NOTIFY_404_RESPONSE_CRITICAL

Table 5-98 AUDIT_STALE_NOTIFY_404_RESPONSE_CRITICAL

Field	Details
Description	At least 60 % of the BSF Audit Notification Requests sent to PCF to check for Suspected Stale Sessions have responded with a 5xx or 4xx (excluding 404) Status in the last 10 minutes.
Summary	{{ \$value }} % of the BSF Audit Notification Requests sent to PCF to check for Suspected Stale Sessions have responded with a 5xx or 4xx (excluding 404) Status in the last 10 minutes.
Severity	CRITICAL
Expression	(sum by (namespace) (increase(ocbsf_query_response_count_total{response_code="404"}[10m])) / sum by (namespace) (increase(ocbsf_query_response_count_total[10m])) * 100 >= 60
OID	1.3.6.1.4.1.323.5.3.37.1.2.62
Metric Used	ocbsf_query_response_count_total
Recommended Actions	Cause: More than 60% of BSF audit notification requests sent to PCF to check suspected stale sessions returned HTTP 404 in the last 10 minutes. Diagnostic Information: Verify whether PCF returns 404 NOT_FOUND for the BSF PCF bindings used by audit query or notification requests. Recovery: The alert clears when BSF audit requests to PCF receive a successful 2xx response or another HTTP error response.

5.1.86 AUDIT_STALE_NOTIFY_404_RESPONSE_MAJOR

Table 5-99 AUDIT_STALE_NOTIFY_404_RESPONSE_MAJOR

Field	Details
Description	At least 40 % of the BSF Audit Notification Requests sent to PCF to check for Suspected Stale Sessions have responded with a 5xx or 4xx (excluding 404) Status in the last 10 minutes.
Summary	{{ \$value }} % of the BSF Audit Notification Requests sent to PCF to check for Suspected Stale Sessions have responded with a 5xx or 4xx (excluding 404) Status in the last 10 minutes.
Severity	MAJOR
Expression	(sum by (namespace) (increase(ocbsf_query_response_count_total{response_code="404"}[10m])) / sum by (namespace) (increase(ocbsf_query_response_count_total[10m])) * 100 >= 40 < 60

Table 5-99 (Cont.) AUDIT_STALE_NOTIFY_404_RESPONSE_MAJOR

Field	Details
OID	1.3.6.1.4.1.323.5.3.37.1.2.62
Metric Used	ocbsf_query_response_count_total
Recommended Actions	Cause: More than 40% and less than 60% of BSF audit notification requests sent to PCF to check suspected stale sessions returned HTTP 404 in the last 10 minutes. Diagnostic Information: Verify whether PCF returns 404 NOT_FOUND for the BSF PCF bindings used by audit query or notification requests. Recovery: The alert clears when BSF audit requests to PCF receive a successful 2xx response or another HTTP error response.

5.1.87 AUDIT_STALE_NOTIFY_404_RESPONSE_MINOR

Table 5-100 AUDIT_STALE_NOTIFY_404_RESPONSE_MINOR

Field	Details
Description	At least 20 % of the BSF Audit Notification Requests sent to PCF to check for Suspected Stale Sessions have responded with a 5xx or 4xx (excluding 404) Status in the last 10 minutes.
Summary	{{ \$value }} % of the BSF Audit Notification Requests sent to PCF to check for Suspected Stale Sessions have responded with a 5xx or 4xx (excluding 404) Status in the last 10 minutes.
Severity	MINOR
Expression	(sum by (namespace) (increase(ocbsf_query_response_count_total{response_code="404"}[10m])) / sum by (namespace) (increase(ocbsf_query_response_count_total[10m])) * 100 >= 20 < 40
OID	1.3.6.1.4.1.323.5.3.37.1.2.62
Metric Used	ocbsf_query_response_count_total
Recommended Actions	Cause: More than 20% and less than 40% of BSF audit notification requests sent to PCF to check suspected stale sessions returned HTTP 404 in the last 10 minutes. Diagnostic Information: Verify whether PCF returns 404 NOT_FOUND for the BSF PCF bindings used by audit query or notification requests. Recovery: The alert clears when BSF audit requests to PCF receive a successful 2xx response or another HTTP error response.

5.1.88 RAA_RX_FAIL_COUNT_EXCEEDS_CRITICAL_THRESHOLD

Table 5-101 RAA_RX_FAIL_COUNT_EXCEEDS_CRITICAL_THRESHOLD

Field	Details
Description	RAA Rx fail count exceeds the critical threshold limit.
Summary	RAA Rx fail count exceeds the critical threshold limit.
Severity	CRITICAL
Expression	(sum by (namespace) (rate(ocbsf_diam_response_network_total{applId="16777236", msgType="RAA", responseCode!~"2.*", direction="in"}[5m])) / sum by (namespace) (rate(ocbsf_diam_response_network_total{applId="16777236", msgType="RAA", direction="in"}[5m]))) * 100 > 90
OID	1.3.6.1.4.1.323.5.3.37.1.2.59
Metric Used	ocbsf_diam_response_network_total
Recommended Actions	Cause: More than 90% of received Diameter RAA responses are non-success responses. Diagnostic Information: Check the response-code distribution, peer reachability, TLS, routing, realm configuration, and logs for protocol or AVP errors. Recovery: Restore connectivity or correct the peer configuration. If load-related, reduce traffic, scale capacity, or apply traffic shedding. Escalate with diagnostics if the issue persists.

5.1.89 RAA_RX_FAIL_COUNT_EXCEEDS_MAJOR_THRESHOLD

Table 5-102 RAA_RX_FAIL_COUNT_EXCEEDS_MAJOR_THRESHOLD

Field	Details
Description	RAA Rx fail count exceeds the major threshold limit.
Summary	RAA Rx fail count exceeds the major threshold limit.
Severity	MAJOR
Expression	(sum by (namespace) (rate(ocbsf_diam_response_network_total{applId="16777236", msgType="RAA", responseCode!~"2.*", direction="in"}[5m])) / sum by (namespace) (rate(ocbsf_diam_response_network_total{applId="16777236", msgType="RAA", direction="in"}[5m]))) * 100 > 80 <= 90
OID	1.3.6.1.4.1.323.5.3.37.1.2.59
Metric Used	ocbsf_diam_response_network_total

Table 5-102 (Cont.) RAA_RX_FAIL_COUNT_EXCEEDS_MAJOR_THRESHOLD

Field	Details
Recommended Actions	Cause: Between 80% and 90% of received Diameter RAA responses are non-success responses. Diagnostic Information: Check the response-code distribution, peer reachability, TLS, routing, and realm configuration. Recovery: Correct connectivity or routing issues. If systemic, reduce load or scale capacity. Monitor the trend and escalate if it persists or worsens.

5.1.90 RAA_RX_FAIL_COUNT_EXCEEDS_MINOR_THRESHOLD

Table 5-103 RAA_RX_FAIL_COUNT_EXCEEDS_MINOR_THRESHOLD

Field	Details
Description	RAA Rx fail count exceeds the minor threshold limit.
Summary	RAA Rx fail count exceeds the minor threshold limit.
Severity	MINOR
Expression	(sum by (namespace) (rate(ocbsf_diam_response_network_total{applId="16777236", msgType="RAA", responseCode! ~"2.*", direction="in"}[5m])) / sum by (namespace) (rate(ocbsf_diam_response_network_total{applId="16777236", msgType="RAA", direction="in"}[5m]))) * 100 > 60 <=80
OID	1.3.6.1.4.1.323.5.3.37.1.2.59
Metric Used	ocbsf_diam_response_network_total
Recommended Actions	Cause: Between 60% and 80% of received Diameter RAA responses are non-success responses. Diagnostic Information: Check the response-code distribution and correlate the increase with peer latency, recent configuration changes, and logs. Recovery: Correct any identified peer, routing, TLS, or protocol issue, and monitor the trend.

5.1.91 STA_RX_FAIL_COUNT_EXCEEDS_CRITICAL_THRESHOLD

Table 5-104 STA_RX_FAIL_COUNT_EXCEEDS_CRITICAL_THRESHOLD

Field	Details
Description	STA Rx fail count exceeds the critical threshold limit.

Table 5-104 (Cont.) STA_RX_FAIL_COUNT_EXCEEDS_CRITICAL_THRESHOLD

Field	Details
Summary	STA Rx fail count exceeds the critical threshold limit.
Severity	CRITICAL
Expression	(sum by (namespace) (rate(ocbsf_diam_response_network_total{appId="16777236", msgType="STA", responseCode! ~"2.*", direction="in"}[5m])) / sum by (namespace) (rate(ocbsf_diam_response_network_total{appId="16777236", msgType="STA", direction="in"}[5m]))) * 100 > 90
OID	1.3.6.1.4.1.323.5.3.37.1.2.60
Metric Used	ocbsf_diam_response_network_total
Recommended Actions	Cause: More than 90% of received Diameter STA responses for Rx, application ID 16777236, are non-2xx. This indicates severe failures in session termination or related Rx procedures. Diagnostic Information: Validate the response ratio using ocbsf_diam_response_network_total{msgType="STA", appId="16777236", direction="in"}, grouped by response code. Identify the dominant failing response code and whether failures originate from a specific peer. Correlate the condition with Diameter peer connectivity, request timeouts, and congestion or overload state. Recovery: Restore Diameter peer connectivity or routing if impaired. If the condition is overload-driven, reduce load, increase resources, or apply stricter shedding for low-priority traffic. Collect the failing response-code distribution and logs, then escalate to peer or dependency owners with supporting evidence if required.

5.1.92 STA_RX_FAIL_COUNT_EXCEEDS_MAJOR_THRESHOLD

Table 5-105 STA_RX_FAIL_COUNT_EXCEEDS_MAJOR_THRESHOLD

Field	Details
Description	STA Rx fail count exceeds the major threshold limit.
Summary	STA Rx fail count exceeds the major threshold limit.
Severity	MAJOR

Table 5-105 (Cont.) STA_RX_FAIL_COUNT_EXCEEDS_MAJOR_THRESHOLD

Field	Details
Expression	(sum by (namespace) (rate(ocbsf_diam_response_network_total{appId="16777236", msgType="STA", responseCode! ~"2.*", direction="in"}[5m])) / sum by (namespace) (rate(ocbsf_diam_response_network_total{appId="16777236", msgType="STA", direction="in"}[5m]))) * 100 > 80 <=90
OID	1.3.6.1.4.1.323.5.3.37.1.2.60
Metric Used	ocbsf_diam_response_network_total
Recommended Actions	Cause: Between 80% and 90% of received STA responses are non-2xx, indicating major degradation. Diagnostic Information: Perform the same checks as for the critical alert. Determine whether one response code dominates and whether the issue is isolated to specific pods or peers. Recovery: Apply early mitigation by correcting connectivity or routing, performing controlled restarts where the issue is localized, or reducing load where it is systemic. Escalate if the percentage increases or customer KPIs are affected.

5.1.93 STA_RX_FAIL_COUNT_EXCEEDS_MINOR_THRESHOLD

Table 5-106 STA_RX_FAIL_COUNT_EXCEEDS_MINOR_THRESHOLD

Field	Details
Description	STA Rx fail count exceeds the minor threshold limit.
Summary	STA Rx fail count exceeds the minor threshold limit.
Severity	MINOR
Expression	(sum by (namespace) (rate(ocbsf_diam_response_network_total{appId="16777236", msgType="STA", responseCode! ~"2.*", direction="in"}[5m])) / sum by (namespace) (rate(ocbsf_diam_response_network_total{appId="16777236", msgType="STA", direction="in"}[5m]))) * 100 > 60 <=80
OID	1.3.6.1.4.1.323.5.3.37.1.2.60
Metric Used	ocbsf_diam_response_network_total

Table 5-106 (Cont.) STA_RX_FAIL_COUNT_EXCEEDS_MINOR_THRESHOLD

Field	Details
Recommended Actions	Cause: Between 60% and 80% of received STA responses are non-2xx, providing an early warning of degradation. Diagnostic Information: Confirm the response ratio and dominant response code, and correlate the condition with peer health and latency. Recovery: Correct the most common failure cause, monitor the trend closely, and apply proactive capacity or load measures if the condition worsens.

5.1.94 ASA_RX_FAIL_COUNT_EXCEEDS_CRITICAL_THRESHOLD

Table 5-107 ASA_RX_FAIL_COUNT_EXCEEDS_CRITICAL_THRESHOLD

Field	Details
Description	ASA Rx fail count exceeds the critical threshold limit.
Summary	ASA Rx fail count exceeds the critical threshold limit.
Severity	CRITICAL
Expression	<pre>(sum by (namespace) (rate(ocbsf_diam_response_network_total{appld=" 16777236", msgType="ASA", responseCode! ~"2.*",direction="in"}[5m])) / sum by (namespace) (rate(ocbsf_diam_response_network_total{appld=" 16777236", msgType="ASA", direction="in"}[5m]))) * 100 > 90</pre>
OID	1.3.6.1.4.1.323.5.3.37.1.2.61
Metric Used	ocbsf_diam_response_network_total
Recommended Actions	Cause: More than 90% of received Diameter ASA responses are non-success responses. Diagnostic Information: Check the response-code distribution, peer reachability, TLS, routing, realm configuration, and logs for protocol or AVP errors. Recovery: Restore connectivity or correct the peer configuration. If load-related, reduce traffic, scale capacity, or apply traffic shedding. Escalate with diagnostics if the issue persists.

5.1.95 ASA_RX_FAIL_COUNT_EXCEEDS_MAJOR_THRESHOLD

Table 5-108 ASA_RX_FAIL_COUNT_EXCEEDS_MAJOR_THRESHOLD

Field	Details
Description	ASA Rx fail count exceeds the major threshold limit.
Summary	ASA Rx fail count exceeds the major threshold limit.
Severity	MAJOR
Expression	(sum by (namespace) (rate(ocbsf_diam_response_network_total{appld="16777236", msgType="ASA", responseCode! ~"2.*", direction="in"}[5m])) / sum by (namespace) (rate(ocbsf_diam_response_network_total{appld="16777236", msgType="ASA", direction="in"}[5m]))) * 100 > 80 <=90
OID	1.3.6.1.4.1.323.5.3.37.1.2.61
Metric Used	ocbsf_diam_response_network_total
Recommended Actions	Cause: Between 80% and 90% of received Diameter ASA responses are non-success responses. Diagnostic Information: Check the response-code distribution, peer reachability, TLS, routing, and realm configuration. Recovery: Correct connectivity or routing issues. If systemic, reduce load or scale capacity. Monitor the trend and escalate if it persists or worsens.

5.1.96 ASA_RX_FAIL_COUNT_EXCEEDS_MINOR_THRESHOLD

Table 5-109 ASA_RX_FAIL_COUNT_EXCEEDS_MINOR_THRESHOLD

Field	Details
Description	ASA Rx fail count exceeds the minor threshold limit.
Summary	ASA Rx fail count exceeds the minor threshold limit.
Severity	MINOR
Expression	(sum by (namespace) (rate(ocbsf_diam_response_network_total{appld="16777236", msgType="ASA", responseCode! ~"2.*"}[5m])) / sum by (namespace) (rate(ocbsf_diam_response_network_total{appld="16777236", msgType="ASA", direction="in"}[5m]))) * 100 > 60 <=80
OID	1.3.6.1.4.1.323.5.3.37.1.2.61
Metric Used	ocbsf_diam_response_network_total

Table 5-109 (Cont.) ASA_RX_FAIL_COUNT_EXCEEDS_MINOR_THRESHOLD

Field	Details
Recommended Actions	Cause: Between 60% and 80% of received Diameter ASA responses are non-success responses. Diagnostic Information: Check the response-code distribution and correlate the increase with peer latency, recent configuration changes, and logs. Recovery: Correct any identified peer, routing, TLS, or protocol issue, and monitor the trend.

5.1.97 EGRESS_GATEWAY_DNS_RESOLUTION_FAILURE

Table 5-110 EGRESS_GATEWAY_DNS_RESOLUTION_FAILURE

Field	Details
Description	Egress Gateway encountered DNS SRV resolution failure.
Summary	"\$labels.microservice is unable to route to SCP in namespace \$labels.namespace due to DNS-SRV lookup failure for the FQDN \$labels.vfqdn"
Severity	Major
Expression	$(\text{increase}(\text{ocbsf_oc_egressgateway_resolve_fqdn_from_ars_failure_total}[2m])) > 0$
OID	1.3.6.1.4.1.323.5.3.37.1.2.63
Metric Used	ocbsf_oc_egressgateway_resolve_fqdn_from_ars_failure_total
Recommended Actions	Cause: A misconfiguration prevents the Egress Gateway from routing requests to the SCP because DNS SRV lookup fails. Diagnostic Information: • Confirm that SBI routing is enabled in the Egress Gateway ConfigMap. • Check the configured FQDN or virtual FQDN value in the Egress Gateway ConfigMap. • Verify that the FQDN or virtual FQDN value in the custom values file is correct. Recovery: Configure the correct FQDN or virtual FQDN. The alert clears after the Egress Gateway can successfully resolve and route to the SCP.

5.1.98 DIAMETER_POD_ERROR_RESPONSE_CRITICAL

Table 5-111 DIAMETER_POD_ERROR_RESPONSE_CRITICAL

Field	Details
Description	At least 10% of the Diam Response connection requests failed with error DIAMETER_UNABLE_TO_DELIVER in pod {{\${labels.pod}}}
Summary	At least 10% of the Diam Response requests failed with error DIAMETER_UNABLE_TO_DELIVER.
Severity	Critical
Expression	(sum by (pod) (rate(ocbsf_diam_response_network_total{responseCode="3002"}[2m])))/(sum by (pod) (rate(ocbsf_diam_response_network_total[2m]))) * 100 >=10
OID	1.3.6.1.4.1.323.5.3.37.1.2.46
Metric Used	ocbsf_diam_response_network_total
Recommended Actions	Cause: The Diameter pod is returning a critical level of error responses. Diagnostic Information: Check the affected pod, peer connectivity, TLS status, routing, and error logs. Recovery: Restore connectivity, correct configuration or protocol errors, and restart only the affected pod if necessary. Reduce traffic or scale capacity if overload is contributing. Escalate if recovery is unsuccessful.

5.1.99 DIAMETER_POD_ERROR_RESPONSE_MAJOR

Table 5-112 DIAMETER_POD_ERROR_RESPONSE_MAJOR

Field	Details
Description	At least 5% of the Diam Response connection requests failed with error DIAMETER_UNABLE_TO_DELIVER in pod {{\${labels.pod}}}
Summary	At least 5% of the Diam Response requests failed with error DIAMETER_UNABLE_TO_DELIVER.
Severity	Major
Expression	(sum by (pod) (rate(ocbsf_diam_response_network_total{responseCode="3002"}[2m])))/(sum by (pod) (rate(ocbsf_diam_response_network_total[2m]))) * 100 >=5
OID	1.3.6.1.4.1.323.5.3.37.1.2.46
Metric Used	ocbsf_diam_response_network_total
Recommended Actions	Cause: The Diameter pod is returning a major level of error responses. Diagnostic Information: Check the affected pod, peer connectivity, TLS status, routing, and recent changes. Recovery: Correct identified issues and monitor recovery. If the condition persists or increases, reduce load, scale capacity, or escalate with diagnostics.

5.1.100 DIAMETER_POD_ERROR_RESPONSE_MINOR

Table 5-113 DIAMETER_POD_ERROR_RESPONSE_MINOR

Field	Details
Description	At least 1% of the Diam Response connection requests failed with error DIAMETER_UNABLE_TO_DELIVER in pod {{labels.pod}}
Summary	At least 1% of the Diam Response requests failed with error DIAMETER_UNABLE_TO_DELIVER.
Severity	Minor
Expression	(sum by (pod) (rate(ocbsf_diam_response_network_total{responseCode="3002"}[2m])))/(sum by (pod) (rate(ocbsf_diam_response_network_total[2m]))) * 100 >=1
OID	1.3.6.1.4.1.323.5.3.37.1.2.46
Metric Used	ocbsf_diam_response_network_total
Recommended Actions	Cause: The Diameter pod is returning a minor level of error responses. Diagnostic Information: Review pod logs, response codes, peer latency, and recent configuration changes. Recovery: Correct any identified connectivity, routing, TLS, or protocol issue, and monitor for recurrence.

5.1.101 DIAM_GATEWAY_CERTIFICATE_EXPIRY_MINOR

Table 5-114 DIAM_GATEWAY_CERTIFICATE_EXPIRY_MINOR

Field	Details
Description	Diam-gw certificate expiry in less than 6 months for {{labels.namespace}}
Summary	Diam-gw certificate expiry in less than 6 months
Severity	Minor
Expression	dgw_tls_cert_expiration_seconds - time() <= 15724800
OID	1.3.6.1.4.1.323.5.3.37.1.2.47
Metric Used	dgw_tls_cert_expiration_seconds

Table 5-114 (Cont.) DIAM_GATEWAY_CERTIFICATE_EXPIRY_MINOR

Field	Details
Recommended Actions	Cause: A Diameter Gateway TLS certificate is approaching expiry, creating a risk of future connection failures if it is not rotated. Diagnostic Information: Identify the expiring certificate and affected instance. Where available, review <code>min_by (instance,cert_subject,cert_name) (dgw_tls_cert_expiration_seconds - time())</code>, adjusting labels for the environment. Recovery: Start the certificate renewal process. Create and approve the new certificate, validate the certificate chain and SANs, and schedule rotation before the major-alert window. Confirm that Diameter Gateway can load the new certificate in a staging or non-production environment first.