

Oracle® Communications

Cloud Native Core, Binding Support Function

User Guide



Release 26.1.201
G58218-01
August 2026

ORACLE®

Copyright © 2019, 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	2
2	BSF Architecture	
2.1	BSF Architecture	1
3	BSF Features	
3.1	Support for SBI Message Priority Header	1
3.2	Application Framework Change	10
3.3	Support for Optimizing N7 Session Lookup for AAR messages in Rx Call Flows	11
3.4	BSF Management Service Pod Congestion Control	13
3.5	Congestion Control Discard Message Response Code	18
3.6	Stale Binding Detection Audit, Report and Recover	22
3.7	Traffic Segregation	25
3.8	Support for Dual Stack	27
3.9	Logging Support for Error Response	41
3.10	Support for TLS	44
3.10.1	Support for TLS Using HTTPs	48
3.10.2	Support for TLS in Diameter Gateway	48
3.10.3	TLS 1.3 Support for Kubernetes API Server	50
3.11	Enhancements to Error Response	50
3.12	Validating Destination-Realm Attribute-Value Pair (AVP) Received in AAR-I Message	51
3.13	Support for Automated Certificate Lifecycle Management	54
3.14	Support for cnDBTier APIs in CNC Console	58
3.15	Diameter Session Retry	60
3.16	Support for BSF Status on NRF on CNC Console	74
3.17	Network Policies	75
3.18	Monitoring the Availability of SCP using HTTP2 OPTIONS	79
3.19	Supports 3gpp-Sbi-Correlation-Info Header	81
3.20	Configurations for Pre and Post Upgrade/Install Validations	83
3.21	Detection and Handling of Late Arrival Requests	84

3.22	Support for Timer Configuration	87
3.23	Support for Session Retry and Alternate Route Service	90
3.24	XFCC Header Validation	90
3.25	Georedundancy Support	96
3.26	Diameter Gateway Pod Congestion Control	100
3.26.1	Support for Server Header	107
3.27	Overload Control	110
3.27.1	Overload Control - Diameter	114
3.27.2	Overload Control - SBI	115
3.28	Rate Limiting - SBI	116
3.29	Pod Protection at Ingress Gateway	118
3.30	Service Mesh for Intra-NF Communication	118
3.31	Turning off AccessToken signature Validation	120
3.32	Automated Test Suite Support	120
3.33	SBI Error Codes	120
3.34	Handling Stale Session in BSF	121
3.34.1	Enhanced Logging of BSF Stale Binding Deletion	133
3.35	Support Multiple Cluster Deployment at CNC Console	136
3.36	Support for 3GPP NF Sets and Binding Headers	138
3.37	Support for User-Agent Header	140
3.38	Support for Active Sessions Counter	143
3.39	Controlled Shutdown of an Instance	147
3.40	Graceful Termination of Kubernetes Pods	153
3.41	NF Scoring for a Site	154
3.42	NRF Client Retry and Health Check	157
3.43	BSF Message Feed for Monitoring	158
3.44	Optimizing Lookup in BSF Management Service Database	162
3.45	Diameter Gateway Connection Repaving	186
3.46	Security Context	194

4 Configuring BSF Using CNC Console

4.1	General Configurations	2
4.1.1	General Settings	2
4.1.2	SBI Error Codes Configurations	4
4.1.3	SBI Ingress Error Code Profiles Collection	9
4.2	Error Handling	10
4.2.1	Error Configurations	10
4.3	Logging Configurations	14
4.3.1	Logging Level	15
4.3.2	Subscriber Activity Logging	16
4.4	Service Configurations	19

4.4.1	Management Service	19
4.4.2	Audit Service	28
4.4.2.1	Audit	28
4.4.2.2	Audit Schedule Data	29
4.4.3	IGW	31
4.4.4	Common	32
4.4.4.1	Message Profiles	32
4.4.4.2	Retry Profiles	35
4.4.4.3	NF Communication Profiles	39
4.4.5	NRF Agent Service	42
4.5	Diameter Configurations	44
4.5.1	Settings	45
4.5.2	Peer Nodes	49
4.5.3	Diameter Routing Table	51
4.5.4	Diameter Error Codes Configurations	55
4.5.5	Diameter Repaving	58
4.6	Status and Query	59
4.6.1	Session Viewer	59
4.6.2	BSF NF Data	61
4.6.2.1	BSF Registration Profile	61
4.6.2.2	BSF NRF Status	61
4.6.3	Active Session Query	63
4.6.4	Custom Query	63
4.7	Administration	64
4.7.1	Import and Export	64
4.7.1.1	Exporting BSF Configurations	65
4.7.1.2	Importing BSF Configurations	67
4.7.1.3	Using REST API for BSF Import & Export	70
4.8	Controlled Shutdown Configurations	71
4.8.1	Operational State	71
4.8.2	Diameter Error Mapping	72
4.8.3	SBI Ingress Error Mapping	73
4.9	Overload and Congestion Control Configurations	74
4.9.1.1	Load Shedding Profiles	74
4.9.1.2	Message Priority Profiles	80
4.9.2.1.1	Rate Limiting Policy	82
4.9.2.1.2	Route Level Mapping	83
4.9.2.2.1	Discard Policy Mapping	85
4.9.2.2.2	Discard Policy	86
4.9.3	Congestion Control	87
4.9.3.1	Settings	87
4.9.3.2	Thresholds	92

4.9.3.3	Load Shedding Rules	94
4.9.4	Overload Control Threshold	96
4.10	NF Scoring Configurations	99
4.11	Viewing cnDBTier APIs in CNC Console	101
4.11.1	Backup List	101
4.11.2	cnDBTier Backup Status	101
4.11.3	Database Statistics Report	102
4.11.4	Georeplication Recovery	102
4.11.5	Georeplication Status	104
4.11.6	Georeplication Status Across All Sites	105
4.11.7	cnDBTier Health	107
4.11.8	Heartbeat Status	109
4.11.9	IP Family Preference	110
4.11.10	Local Cluster Status	110
4.11.11	On Demand Backup	111
4.11.12	Version	111
4.11.13	Gracefully Stop and Start Georeplication	112

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
5.1.24	PCF_BINDING_ERROR_RATE_ABOVE_CRITICAL_THRESHOLD	15
5.1.25	INGRESS_CREATE_ERROR_RATE_ABOVE_MINOR_THRESHOLD	15
5.1.26	INGRESS_CREATE_ERROR_RATE_ABOVE_CRITICAL_THRESHOLD	16
5.1.27	INGRESS_CREATE_ERROR_RATE_ABOVE_MAJOR_THRESHOLD	17
5.1.28	INGRESS_DELETE_ERROR_RATE_ABOVE_MINOR_THRESHOLD	17
5.1.29	INGRESS_DELETE_ERROR_RATE_ABOVE_MAJOR_THRESHOLD	18
5.1.30	INGRESS_DELETE_ERROR_RATE_ABOVE_CRITICAL_THRESHOLD	19
5.1.31	DB_TIER_DOWN_ALERT	19
5.1.32	CPU_USAGE_PER_SERVICE_ABOVE_MINOR_THRESHOLD	20
5.1.33	CPU_USAGE_PER_SERVICE_ABOVE_MAJOR_THRESHOLD	21
5.1.34	CPU_USAGE_PER_SERVICE_ABOVE_CRITICAL_THRESHOLD	22
5.1.35	MEMORY_USAGE_PER_SERVICE_ABOVE_MINOR_THRESHOLD	23
5.1.36	MEMORY_USAGE_PER_SERVICE_ABOVE_MAJOR_THRESHOLD	24
5.1.37	MEMORY_USAGE_PER_SERVICE_ABOVE_CRITICAL_THRESHOLD	25
5.1.38	NRF_COMMUNICATION_FAILURE	25
5.1.39	NRF_SERVICE_REQUEST_FAILURE	26
5.1.40	PERF_INFO_ACTIVE_OVERLOAD_THRESHOLD_FETCH_FAILED	26
5.1.41	POD_DOC	27
5.1.42	POD_CONGESTED	28
5.1.43	POD_CONGESTION_L1	29
5.1.44	POD_CPU_CONGESTION_L1	30
5.1.45	POD_CONGESTION_L2	30
5.1.46	POD_CPU_CONGESTION_L2	31
5.1.47	POD_PENDING_REQUEST_DOC	31
5.1.48	POD_PENDING_REQUEST_CONGESTED	32
5.1.49	POD_CPU_DOC	33
5.1.50	POD_CPU_CONGESTED	34
5.1.51	POD_MEMORY_DOC	34
5.1.52	POD_MEMORY_CONGESTED	35
5.1.53	SERVICE_OVERLOADED	36
5.1.54	SERVICE_RESOURCE_OVERLOADED	39
5.1.55	SYSTEM_IMPAIRMENT_MAJOR	47
5.1.56	SYSTEM_IMPAIRMENT_CRITICAL	49
5.1.57	SYSTEM_OPERATIONAL_STATE_PARTIAL_SHUTDOWN	51
5.1.58	SYSTEM_OPERATIONAL_STATE_COMPLETE_SHUTDOWN	52
5.1.59	DIAM_CONN_PEER_DOWN	53
5.1.60	DIAM_CONN_NETWORK_DOWN	54
5.1.61	DIAM_RESPONSE_REALM_VALIDATION_ERROR_CRITICAL	55
5.1.62	DIAM_RESPONSE_REALM_VALIDATION_ERROR_MAJOR	56
5.1.63	DIAM_RESPONSE_REALM_VALIDATION_ERROR_MINOR	57

5.1.64	AUDIT_STALE_NOTIFY_ERROR_RESPONSE_MINOR	57
5.1.65	AUDIT_STALE_NOTIFY_ERROR_RESPONSE_MAJOR	58
5.1.66	AUDIT_STALE_NOTIFY_ERROR_RESPONSE_CRITICAL	59
5.1.67	BSF_CONNECTION_FAILURE	59
5.1.68	INGRESS_GATEWAY_DD_UNREACHABLE_MAJOR	60
5.1.69	EGRESS_GATEWAY_DD_UNREACHABLE_MAJOR	60
5.1.70	DIAM_GATEWAY_CERTIFICATE_EXPIRY_MAJOR	61
5.1.71	DIAM_GATEWAY_CERTIFICATE_EXPIRY_CRITICAL	62
5.1.72	DGW_TLS_CONNECTION_FAILURE	62
5.1.73	BINDING_REVALIDATION_PCF_BINDING_MISSING_MINOR	63
5.1.74	BINDING_REVALIDATION_PCF_BINDING_MISSING_MAJOR	64
5.1.75	BINDING_REVALIDATION_PCF_BINDING_MISSING_CRITICAL	64
5.1.76	BSF_STATE_NON_FUNCTIONAL_CRITICAL	65
5.1.77	POD_PENDING_REQUEST_CONGESTION_L1	65
5.1.78	POD_PENDING_REQUEST_CONGESTION_L2	66
5.1.79	AUDIT_NOT_RUNNING	67
5.1.80	BSF_PCF_BINDING_TABLE_MIGRATED_PERCENTAGE	68
5.1.81	BSF_PCF_BINDING_TABLE_MIGRATION_INVALID_CONFIGURATION	68
5.1.82	CERTIFICATE_EXPIRY_MINOR	69
5.1.83	CERTIFICATE_EXPIRY_MAJOR	70
5.1.84	CERTIFICATE_EXPIRY_CRITICAL	70
5.1.85	AUDIT_STALE_NOTIFY_404_RESPONSE_CRITICAL	71
5.1.86	AUDIT_STALE_NOTIFY_404_RESPONSE_MAJOR	71
5.1.87	AUDIT_STALE_NOTIFY_404_RESPONSE_MINOR	72
5.1.88	RAA_RX_FAIL_COUNT_EXCEEDS_CRITICAL_THRESHOLD	73
5.1.89	RAA_RX_FAIL_COUNT_EXCEEDS_MAJOR_THRESHOLD	73
5.1.90	RAA_RX_FAIL_COUNT_EXCEEDS_MINOR_THRESHOLD	74
5.1.91	STA_RX_FAIL_COUNT_EXCEEDS_CRITICAL_THRESHOLD	74
5.1.92	STA_RX_FAIL_COUNT_EXCEEDS_MAJOR_THRESHOLD	75
5.1.93	STA_RX_FAIL_COUNT_EXCEEDS_MINOR_THRESHOLD	76
5.1.94	ASA_RX_FAIL_COUNT_EXCEEDS_CRITICAL_THRESHOLD	77
5.1.95	ASA_RX_FAIL_COUNT_EXCEEDS_MAJOR_THRESHOLD	78
5.1.96	ASA_RX_FAIL_COUNT_EXCEEDS_MINOR_THRESHOLD	78
5.1.97	EGRESS_GATEWAY_DNS_RESOLUTION_FAILURE	79
5.1.98	DIAMETER_POD_ERROR_RESPONSE_CRITICAL	80
5.1.99	DIAMETER_POD_ERROR_RESPONSE_MAJOR	80
5.1.100	DIAMETER_POD_ERROR_RESPONSE_MINOR	81
5.1.101	DIAM_GATEWAY_CERTIFICATE_EXPIRY_MINOR	81

6 BSF KPIs

7

BSF Metrics

7.1	Egress Gateway Metrics for SCP	5
7.2	Correlation-Info Header Metrics	8
7.3	Configuration Server Metrics	9
7.4	Active Sessions Count Metrics	9
7.5	AppInfo Metrics	10
7.6	Audit Service Metrics	12
7.7	BSF Management Service	14
7.8	Collision Detection Metrics	35
7.9	CM Service Metrics	37
7.10	Diameter Gateway Metrics	38
7.11	NRF Management Service	42
7.12	Overload Control Metrics	44
7.13	PerfInfo Metrics	45
7.14	Pod Congestion Metrics	49
7.15	User-Agent Header Metrics	51
7.16	Query Service Metrics	51
7.17	TLS Metrics	52
7.18	NRF Client Metrics	54
7.19	Metrics for Automated Certificate Lifecycle Management	58

A

Ingress Gateway Metrics

B

Egress Gateway Metrics

C

HTTP Status Codes Supported on SBI

D

Error Code Dictionary

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. A representative at Customer Access Support can assist you with My Oracle Support registration.

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

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

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

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

Acronyms

This following table lists the acronyms and the terminologies used in this document.

Table Acronyms

Acronym	Description
AF	Application Function
BSF	Binding Support Function
DNN	Domain Network Name
FQDN	Fully Qualified Domain Names
GPSI	Generic Public Subscription Identifier
HTTP	Hypertext Transfer Protocol
MCC	Mobile Country Code
MNC	Mobile Network Code
NEF	Network Exposure Function
NF	Network Function
NID	Network Identifier
NRF	Oracle Communications Cloud Native Core, Network Repository Function
PCF	Oracle Communications Cloud Native Core, Policy Control Function
OCPM	Oracle Communications Policy Management
PDU	Protocol Data Unit
RDBMS	Relational Database Management System
S-NSSAI	Single Network Slice Selection Assistance Information. An S-NSSAI is comprised of: - A Slice/Service type (SST), which refers to the expected Network Slice behaviour in terms of features and services; - A Slice Differentiator (SD), which is an optional information that complements the Slice/Service type(s) to differentiate amongst multiple Network Slices of the same Slice/Service type.
SMF	Session Management Function
SNPN	Stand-alone Non-Public Network
SUPI	Subscription Permanent Identifier
UDSF	Unstructured Data Storage network function
UE	User Equipment

What's New in This Guide

This section introduces the new and updated features for release 26.1.2xx.

Release 26.1.201 - G58218-01, August 2026

- **Feature Updates:**

- **New Features:**

- * **Optimizing Lookup in BSF Management Service Database**

- * Added [Optimizing Lookup in BSF Management Service Database](#) to support optimized BSF PCF binding storage and lookup using VARCHAR-based binding data, alternate lookup tables for DISCOVERY, and lazy migration of existing records.

- * Added the following new metrics to [BSF Management Service](#):

- * ocbsf_feature_incompatibility
 - * ocbsf_binding_record_lookup_latency
 - * ocbsf_binding_table_lookup_total
 - * ocbsf_binding_query_processing_latency
 - * ocbsf_alt_key_lookup_total
 - * ocbsf_index_based_lookup_fallback_total
 - * ocbsf_lookup_table_reinsertions_total
 - * ocbsf_lookup_table_stale_deletions_total
 - * ocbsf_db_latency_seconds_count
 - * ocbsf_db_latency_seconds
 - * ocbsf_db_latency_seconds_sum
 - * ocbsf_db_latency_seconds_max

- * Updated the following alerts:

- * [BSF_PCF_BINDING_TABLE_MIGRATED_PERCENTAGE](#)
 - * [BSF_PCF_BINDING_TABLE_MIGRATION_INVALID_CONFIGURATION](#)

- * **Diameter Gateway Connection Repaving**

- * Added [Diameter Gateway Connection Repaving](#) to describe the Diameter Gateway Connection Repaving feature that provides a controlled mechanism to re-balance Diameter connections across the cluster without restarting services or causing a complete outage.
 - * Added [Diameter Repaving](#) to describe Diameter Gateway connection repaving configuration.
 - * Added the following advanced settings keys to [Settings](#):
 - * DGW_CONNECTION_REPAVING.MAX_STAGGER_JITTER_MS
 - * DGW_CONNECTION_REPAVING.REINVOCATION_DELAY_MIN
 - * DGW_CONNECTION_REPAVING.CONNECTION_DISCONNECT_INTERVAL_MS

- * DGW_CONNECTION_REPAVING.DPR_DELAY_MS
 - * DGW_CONNECTION_REPAVING.DPR_DISCONNECT_CAUSE
 - * Added `diam_repaving_disconnections_total` metric to [Diameter Gateway Metrics](#) that counts the total number of connections terminated by the "force disconnect" mechanism.
- * **Security Context**
 - * Added [Security Context](#) which applies Kubernetes pod and container security-context settings to strengthen the security posture of BSF deployments.
- **Enhancements:**
 - **Feature Enhancements**
 - * **Diameter Session Retry**
 - * Added details about Diameter Message Retry support for message types Rx-STR, RX-RAR and Rx-ASR in the section [Diameter Session Retry](#).
 - * Updated the dimensions of the following metrics in [BSF Management Service](#):
 - * `ocbsf_query_request_count_total`
 - * `ocbsf_query_response_count_total`
 - * `ocbsf_diam_request_network_total`
 - * `ocbsf_diam_request_inter_total`
 - * Added the following alerts in the [List of Alerts](#) section:
 - * [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](#)
 - * **Stale Session Audit Retry:**
 - * Added details about session retry support for stale session audit in the section [Handling Stale Session in BSF](#).
 - * Added the following metrics in [BSF Management Service](#):
 - * `ocpm_send_internal_failure_total`
 - * `ocbsf_ar_request_total`
 - * `ocbsf_ar_response_total`
 - * `ocbsf_ar_failure_total`
 - * `ocpm_nfDiscovery_request_total`
 - * `ocpm_nfDiscovery_response_total`
 - * 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](#)
 - * Updated the following alerts:
 - * [AUDIT_STALE_NOTIFY_ERROR_RESPONSE_MINOR](#)
 - * [AUDIT_STALE_NOTIFY_ERROR_RESPONSE_MAJOR](#)
 - * [AUDIT_STALE_NOTIFY_ERROR_RESPONSE_CRITICAL](#)
 - * **Support for cnDBTier APIs in CNC Console**
 - * Updated the [Support for cnDBTier APIs in CNC Console](#) section with details of the APIs IP Family Reference and Gracefully Start and Stop Georeplication.
 - * Added the sections [IP Family Preference](#) and [Gracefully Stop and Start Georeplication](#) to [Viewing cnDBTier APIs in CNC Console](#).
- **General Updates:**
 - Added `ocbsf_mesh_map_missing_dgw_peer_connection_total` to count the number of times the `DGW_IP_ADDRESS` entry in the mesh map is detected as missing.
 - Added `ocbsf_oc_egressgateway_resolve_fqdn_from_ars_failure_total` to count Egress Gateway failures to resolve an SCP peer's FQDN through DNS SRV lookup from ARS.
 - Added the alert [EGRESS_GATEWAY_DNS_RESOLUTION_FAILURE](#).

1

Introduction

This document provides information on how to configure the Oracle Communications Cloud Native Core Binding Support Function (BSF).

1.1 Overview

Binding Support Function (BSF) provides a Policy Control Function (PCF) session binding functionality and ensures an Application Function request for a certain PCF session that reaches the relevant Binding Support Function (BSF) holding the PCF Session information.

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.

BSF supports the following functions:

- Allows BSF users to register, discover, and remove the binding information
- Allows network function consumers to retrieve the binding information

CNC BSF Availability

Oracle Communications Cloud Native Core (CNC) BSF availability is dependent on many factors. BSF applications are designed to achieve 99.999% availability, according to the applicable Telecommunications Industry Association TL9000 standards, with the following deployment requirements:

- Deploy on a Cloud Native Environment with at least 99.999% Availability.
- Deploy with $n + k$ application redundancy, where k is greater than or equal to one.
- Maintain production software within $n-2$ software releases, where n is the current general availability release.

Note

BSF 25.1.100 supports upgrade from 24.3.x, 24.2.x, and 23.4.6.
Also, 25.1.100 can be rolled back to 24.3.x, 24.2.x, 23.4.6.

- Apply bug fixes, critical patches, and configuration recommendations provided by Oracle promptly.
- Maintain disaster recovery procedures external to the applications for the reconstruction of lost or altered files, data, programs, or Cloud Native environment.
- Install, configure, operate, and maintain CNC BSF as per Oracle's applicable installation, operation, administration, and maintenance specifications.

- Maintain an active support contract and provide access to the deployed CNC BSF and your personnel to assist Oracle in addressing any outage.

CNC BSF availability is measured for each calendar year and is calculated as follows:

Table 1-1 Measuring CNC BSF Availability

Availability	Description
Planned Product Availability	(Product available time in each month) less (Excluded Time (defined below) in each month).
Actual Product Availability	(Planned Product Availability) less (any Unscheduled Outage)
Product Availability Level	(Actual Product Availability across all Production instances divided by Planned Product Availability across all Production instances) x 100

Note

Excluded Time means:

- Scheduled maintenance time.
- Lack of power or backhaul connectivity, except to the extent that such lack of backhaul connectivity was caused directly by the CNC NF.
- Hardware failure.
- Issues arising out of configuration errors or omissions.
- Failures caused by third-party equipment or software not provided by Oracle.
- Occurrence of any event under Force Majeure.
- Any time associated with failure to maintain the recommended architecture and redundancy model requirements above.

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

BSF Architecture


This section provides information about Oracle Communications Cloud Native Core Binding Support Function (BSF) Architecture.

2.1 BSF Architecture

The BSF network function is a cloud native application that consists of multiple microservices running in cloud native environment.

The following figure describes the component level architecture of BSF:

Figure 2-1 Architecture of Binding Support Function



The microservices are logically categorized as following categories:

- **Connectivity:** It includes external entities that interact with BSF and vice-versa.
- **Adaptation & Mobility:** Components under this category provide an interface that ensures seamless interaction between external and internal entities. It contains the following components:
 - **Diameter Gateway:** Acts as a Gateway for all diameter traffic to BSF.
 - **Ingress Gateway:** Acts as a gateway for all incoming Ingress HTTP traffic.
 - **Egress Gateway:** Acts as a gateway for all outgoing Egress HTTP traffic.
 - **Alternate Route Service:** Provides alternate routing destinations for BSF (NRF-management) to re-route during failures. There are two options for alternate routing:
 1. **DNS (SRV) based routing** - A DNS server is required to resolve SRV records having alternate destinations with higher priority.
 2. **Static routing** - Static configuration for alternate destinations with weight or priority.
- **Business Logic:** Components under this category can be enabled based on deployment needs. It includes the following components:
 - **Bsf-Management:** It implements the nbsf interface as defined in 3GPP Specification 29.521.
 - **Nrf-Management:** It helps in the autonomous discovery of network functions.
 - **Nrf-Client-NfDiscovery:** It discovers eligible network functions from the NRF to enable alternate routing.
- **Operations & Maintenance:** Components under this category perform specific tasks as follows:
 - **Config-Mgmt:** It provides GUI and REST OAM interfaces for service provisioning.
 - **Config Server:** It abstracts the database for storage and retrieval of configuration.
 - **Query:** It processes session viewer queries triggered from config management service.
 - **App-Info:** It monitors application (microservice) health and status.
 - **Perf-Info:** It monitors application (microservice) capacity and load status.
 - **Audit-service:** It runs the Audit engine to detect and process stale session records. Also, Audit-service counts the maximum number of active sessions for a particular service.
- **Data Management:** Components under this category are responsible for storing various types of persistent data.
- **Ingress and Egress Gateway Traffic Management**
For more information on Ingress and Egress Gateway Traffic Management, see *Oracle Communications Cloud Native Environment User Guide*.

3

BSF Features

This section describes the key features of Oracle Communications Cloud Native Core, Binding Support Function (BSF).

3.1 Support for SBI Message Priority Header

BSF supports 3gpp-Sbi-Message-Priority header in order to enable overload and congestion control through informed throttling decisions.

Using 3gpp-Sbi-Message-Priority header, BSF and the associated producers, consumers and proxies such as SCP, NRF, and PCF can distinguish between high-priority and low-priority Service-Based Interface (SBI) messages.

The 3gpp-Sbi-Message-Priority header specifies HTTP/2 message priority values ranging from 0 (highest) to 31 (lowest). Integrating this header enables priority-based processing of both ingress and egress HTTP/2 SBI messages.

BSF as a consumer: BSF sends the 3gpp-Sbi-Message-Priority header in outgoing HTTP requests and also supports receiving it in responses in all the interaction between BSF and NRF, consuming NRF Management services and BSF interaction with PCF for stale session notification confirmations.


BSF as a producer: BSF receives the 3gpp-Sbi-Message-Priority header in the incoming HTTP requests and propagates the appropriate header value in responses in BSF interaction with PCF for Nbsf services such as Register, Deregister, and Discovery.

Note

The 3gpp-Sbi-Message-Priority header should be integrated before configuring congestion control mechanism.

Call Flow

Figure 3-1 Sample call flow for session binding registration with BSF as a producer



1. PCF Sends Session Binding Register Request

- a. The Policy Control Function (PCF) sends a Session Binding Register request to the BSF (Binding-Support-Function) through the nbsf_Management_Register interface. BSF receives the request through Ingress Gateway.
- b. BSF Management service records the binding data and updates the 3gpp-Sbi-Message-Priority header based on configuration.

- c. BSF responds to PCF with 201 Created message, indicating successful registration, along with updated 3gpp-Sbi-Message-Priority headers.

2. AF Initiates App Session Establishment


- a. The AF initiates application session establishment by sending an Rx AAR-I request to the BSF Diameter Gateway.
- b. The BSF Diameter Gateway performs binding discovery by sending a request to the BSF Management service. The BSF Management service validates the binding information and returns a 200 OK response with the selected PCF binding information to the BSF Diameter Gateway.
- c. During this internal SBI/HTTP interaction, BSF Management service evaluates the 3gpp-Sbi-Message-Priority header based on the selected Message Priority Profile. Based on the profile configuration, the response may include the evaluated 3gpp-Sbi-Message-Priority header.
- d. BSF Diameter Gateway forwards the Rx AAR-I request to the selected PCF and returns the Rx AAA-I response from PCF to the AF.

3. AF/NEF Sends Retrieve Session Binding Request

- a. Application Function(AF)/Network Exposure Function (NEF) sends a Retrieve Session Binding request using the nbsf_Management_Discovery interface to BSF.
- b. The BSF performs discovery based on the current binding information, updates the priority header appropriately.
- c. BSF responds to AF/NEF with 200 OK along with the updated 3gpp-Sbi-Message-Priority headers.

4. PCF Sends De-register Session Binding

- a. The PCF sends a Deregister request using the nbsf_Management_Deregister interface to BSF.
- b. BSF Management service processes the deregistration by removing the binding information, and updates the priority header accordingly.
- c. BSF confirms de-registration by returning 204 NO CONTENT, acknowledging successful deletion of the session binding. BSF includes 3gpp-Sbi-Message-Priority header in the response to deregister request based on the configuration.

Figure 3-2 Sample call flow for session binding registration with BSF as a consumer**1. BSF Registration with NRF**

- a. BSF sends a registration request to NRF with 3gpp-Sbi-Message-Priority header in the request. The request is sent through NRF Client and Egress Gateway.
- b. NRF processes the registration request and responds with 3gpp-Sbi-Message-Priority header in its response.

2. BSF receives Heartbeat/Update Requests from NRF

- a. BSF sends heartbeat or status update requests to the NRF to confirm and update its status in the network.
The request containing 3gpp-Sbi-Message-Priority header is sent to NRF through NRF Client.
- b. NRF processes the request, updates the status of BSF, and responds to BSF with 3gpp-Sbi-Message-Priority header in the response.

3. BSF Deregistration from NRF

- a. BSF sends a deregistration request to NRF. The deregistration request containing 3gpp-Sbi-Message-Priority header is sent to NRF through NRF Client and Egress Gateway.

- b. NRF processes the deregistration request and responds to BSF with 3gpp-Sbi-Message-Priority header in its response.
4. **Audit Notify for Stale Session**
 - a. The Audit service for BSF sends an audit notification (BSF_audit_Notify) to BSF Management service.
 - b. BSF sends a stale session check request to PCF. The request containing 3gpp-Sbi-Message-Priority header is sent to PCF through Egress Gateway.
 - c. The PCF processes the request, checks for any stale session, and responds to BSF. The response from PCF contains the 3gpp-Sbi-Message-Priority header.
 - d. BSF Management service sends the response from PCF to Audit service without 3gpp-Sbi-Message-Priority header.

Upgrade Impact

While upgrading BSF from 25.2.1xx release to 25.2.2xx release or later versions, the default priority for AutonomousNfPatch and AutonomousNfDeRegistration should be 3 and 2 respectively. The default priority of all other message types must be 1.

Managing Support for SBI Message Priority Header

Enable

This feature is not enabled by default. The feature gets enabled when a message priority profile is selected for **Message Priority Profile** field on **Management Service** page on CNC Console, and 3gpp-Sbi-Message-Priority header is added to the outgoing and incoming HTTP requests and responses to BSF depending on the configuration of the Message Priority Profile selected for BSF Management Service.

Configure

Support for SBI message priority header can be configured either using CNC console or using REST API.

To configure the feature using CNC Console:

1. Configure the Message Profile as explained in [Message Profiles](#).
2. Select the required message priority profile to be used under **Message Priority Profile** section on **Management Service** page as explained in [Management Service](#).
3. Configure the traffic prioritization as explained in [NRF Agent Service](#).

Configure using REST API

The following APIs for BSF Management service are used to configure support for SBI message priority header:

- Get Message Profile API for BSF Management service is used to get the configuration of a particular message profile.
Request URL: {apiRoot}/oc-bsf-configuration/v1/services/common/messageProfiles/{messageProfileName}
- Create Message Profile API for BSF Management service is used to create the configuration of a particular message profile.
Request URL: {apiRoot}/oc-bsf-configuration/v1/services/common/messageProfiles
- Update Message Profile API for BSF Management service is used to update the configuration of a particular message profile.

Request URL: {apiRoot}/oc-bsf-configuration/v1/services/common/messageProfiles/{messageProfileName}

- Delete Message Profile API for BSF Management service is used to delete the configuration of a particular message profile.

Request URL: {apiRoot}/oc-bsf-configuration/v1/services/common/messageProfiles/{messageProfileName}

- Export Message Profile API for BSF Management service is used to export the configuration of a particular message profile.

Request URL: {apiRoot}/oc-bsf-configuration/v1/services/common/messageProfiles/export

- Import Message Profile API for BSF Management service is used to import the configuration of a particular message profile.

Request URL: {apiRoot}/oc-bsf-configuration/v1/services/common/messageProfiles/import

The following APIs are used to configure NRF Agent Service:

- Get NRF Agent Service configurations API to get the configuration of traffic prioritization feature of NRF Agent service.

Request URL: {apiRoot}/oc-bsf-configuration/v1/services/nrfAgentServiceConfig

- Update NRF Agent Service configurations API to update the configuration of traffic prioritization feature of NRF Agent service.

Request URL: {apiRoot}/oc-bsf-configuration/v1/services/nrfAgentServiceConfig

- Export NRF Agent Service configurations API to export the configuration of traffic prioritization feature of NRF Agent service.

Request URL: {apiRoot}/oc-bsf-configuration/v1/services/nrfAgentServiceConfig/export

- Import NRF Agent Service configurations API to import the configuration of traffic prioritization feature of NRF Agent service.

Request URL: {apiRoot}/oc-bsf-configuration/v1/services/nrfAgentServiceConfig/import

For more information, see *Management Service* section in *Oracle Communications Cloud Native Core, Binding Support Function REST Specification Guide*

Observability

Metrics

The following metrics are used to get the statistics on the total number of requests received and requests sent by BSF, total number of responses sent by BSF and responses received by BSF during register, deregister, discovery, and stale session scenarios:

- ocbsf_http_message_priority_request_count_total
- ocbsf_http_message_priority_response_count_total

For more information, see [BSF Management Service](#).

Logging

The following logs are generated when 3gpp-Sbi-Message-Priority header is enabled.

Server filter logging

```

{"instant":
{"epochSecond":1758784050,"nanoOfSecond":383374000},"thread":"default-
nioEventLoopGroup-1-5","level":"INFO","loggerName":"ocpm.bsf.api.management.m
essageprofile.PriorityHeaderInboundRequestFilter","message":"incoming request
priority headers: {3gpp-sbi-message-priority=[8], oc-message-priority=[9],
Content-Type=[application/json; charset=UTF-8], Content-Length=[1520],
Host=[localhost:8080], Connection=[Keep-Alive], User-Agent=[Apache-HttpClient/
4.5.13 (Java/17.0.12)], Accept-
Encoding=[gzip,deflate]]","endOfBatch":false,"loggerFqcn":"org.apache.logging.
slf4j.Log4jLogger","threadId":66,"threadPriority":5,"messageTimestamp":"2025-0
9-25T12:37:30.383+0530","ocLogId":"${ctx:ocLogId}"}
{"instant":
{"epochSecond":1758784050,"nanoOfSecond":383578000},"thread":"default-
nioEventLoopGroup-1-5","level":"DEBUG","loggerName":"ocpm.bsf.api.management.m
essageprofile.PriorityHeaderInboundRequestFilter","message":"buildParams map
generated from request with values: {interfaceType=Nbsf_Management,
messageType=Register, 3gpp-sbi-message-priority=8, oc-message-
priority=9}","endOfBatch":false,"loggerFqcn":"org.apache.logging.slf4j.Lo
gger","threadId":66,"threadPriority":5,"messageTimestamp":"2025-09-25T12:37:30
.383+0530","ocLogId":"${ctx:ocLogId}"}
{"instant":
{"epochSecond":1758784050,"nanoOfSecond":383650000},"thread":"default-
nioEventLoopGroup-1-5","level":"DEBUG","loggerName":"ocpm.pcf.framework.ruleen
gine.messageprofile.MessageProfileRuleEngine","message":"select message
profile rule with attribute map {interfaceType=Nbsf_Management,
messageType=Register, 3gpp-sbi-message-priority=8, oc-message-
priority=9}","endOfBatch":false,"loggerFqcn":"org.apache.logging.slf4j.Lo
gger","threadId":66,"threadPriority":5,"messageTimestamp":"2025-09-25T12:37:30
.383+0530","ocLogId":"${ctx:ocLogId}"}
{"instant":
{"epochSecond":1758784050,"nanoOfSecond":383840000},"thread":"default-
nioEventLoopGroup-1-5","level":"DEBUG","loggerName":"ocpm.pcf.framework.ruleen
gine.messageprofile.MessageProfileRuleEngine","message":"Matched
messageProfileRule
messageId=HIGH_PRIORITY_ANY_I","endOfBatch":false,"loggerFqcn":"org.apache.log
ging.slf4j.Log4jLogger","threadId":66,"threadPriority":5,"messageTimestamp":"2
025-09-25T12:37:30.383+0530","ocLogId":"${ctx:ocLogId}"}
{"instant":
{"epochSecond":1758784050,"nanoOfSecond":383901000},"thread":"default-
nioEventLoopGroup-1-5","level":"DEBUG","loggerName":"ocpm.bsf.api.management.m
essageprofile.PriorityHeaderInboundRequestFilter","message":"incoming server
request contains oc-message-priority value as :
9","endOfBatch":false,"loggerFqcn":"org.apache.logging.slf4j.Log4jLogger","thr
eadId":66,"threadPriority":5,"messageTimestamp":"2025-09-25T12:37:30.383+0530"
,"ocLogId":"${ctx:ocLogId}"}
{"instant":
{"epochSecond":1758784050,"nanoOfSecond":384550000},"thread":"default-
nioEventLoopGroup-1-5","level":"DEBUG","loggerName":"ocpm.bsf.api.management.m
essageprofile.PriorityHeaderInboundRequestFilter","message":"incoming server
request evaluated 3gpp-sbi-message-priority value is : 3 for Message
Profile : MessageProfileRule(messageId=HIGH_PRIORITY_ANY_I, rulePriority=5,
interfaceType=Nbsf_Management, messageType=Any, requestPriorityValue=3,
requestPriorityOverride=Always, responsePriorityAssign=Never,
responsePriorityValue=0,

```

```

messageProfileFilterConditions=null)", "endOfBatch": false, "loggerFqcn": "org.apache.logging.slf4j.Log4jLogger", "threadId": 66, "threadPriority": 5, "messageTimestamp": "2025-09-25T12:37:30.384+0530", "ocLogId": "${ctx:ocLogId}"}
{"instant":
{"epochSecond": 1758784050, "nanoOfSecond": 384685000}, "thread": "default-nioEventLoopGroup-1-5", "level": "INFO", "loggerName": "ocpm.bsf.api.management.messageprofile.PriorityHeaderInboundRequestFilter", "message": "final incoming request priority headers: {oc-message-priority=[8], Content-Type=[application/json; charset=UTF-8], Content-Length=[1520], Host=[localhost:8080], Connection=[Keep-Alive], User-Agent=[Apache-HttpClient/4.5.13 (Java/17.0.12)], Accept-Encoding=[gzip, deflate], oc-message-priority=[3]}", "endOfBatch": false, "loggerFqcn": "org.apache.logging.slf4j.Log4jLogger", "threadId": 66, "threadPriority": 5, "messageTimestamp": "2025-09-25T12:37:30.384+0530", "ocLogId": "${ctx:ocLogId}"}
{"instant":
{"epochSecond": 1758784050, "nanoOfSecond": 384916000}, "thread": "default-nioEventLoopGroup-1-5", "level": "DEBUG", "loggerName": "ocpm.bsf.api.management.metrics.BsfMetrics", "message": "Pegging
ocbsf_http_message_priority_request_count_total with direction : in,
rcvd_priority: 9, eval_priority: 3, interfaceType: Nbsf_Management, msgType: Register", "endOfBatch": false, "loggerFqcn": "org.apache.logging.slf4j.Log4jLogger", "threadId": 66, "threadPriority": 5, "messageTimestamp": "2025-09-25T12:37:30.384+0530", "ocLogId": "${ctx:ocLogId}"}
{"instant":
{"epochSecond": 1758784050, "nanoOfSecond": 863867000}, "thread": "boundedElastic-1", "level": "DEBUG", "loggerName": "ocpm.bsf.api.management.metrics.BsfMetrics", "message": "Pegging the Ingress Response Metric for Operation Type : register, Response Code CREATED", "endOfBatch": false, "loggerFqcn": "org.apache.logging.slf4j.Log4jLogger", "threadId": 68, "threadPriority": 5, "messageTimestamp": "2025-09-25T12:37:30.863+0530", "ocLogId": "${ctx:ocLogId}"}
{"instant":
{"epochSecond": 1758784050, "nanoOfSecond": 864543000}, "thread": "boundedElastic-1", "level": "INFO", "loggerName": "ocpm.bsf.api.management.messageprofile.PriorityHeaderInboundRequestFilter", "message": "outgoing server response priority headers: {Location=[/nbsf-management/v1/pcfBindings/91795ebd-74f3-4c90-9bfd-76f22a72a5b3], Content-Type=[application/json], date=[Thu, 25 Sep 2025 07:07:30 GMT], Access-Control-Allow-Origin=[*, *], Access-Control-Allow-Methods=[GET, POST, DELETE, PUT, GET, POST, DELETE, PUT], Access-Control-Allow-Headers=[Content-Type, Content-Type]}, for request headers: {3gpp-sbi-message-priority=[8], Content-Type=[application/json; charset=UTF-8], Content-Length=[1520], Host=[localhost:8080], Connection=[Keep-Alive], User-Agent=[Apache-HttpClient/4.5.13 (Java/17.0.12)], Accept-Encoding=[gzip, deflate], oc-message-priority=[3]}", "endOfBatch": false, "loggerFqcn": "org.apache.logging.slf4j.Log4jLogger", "threadId": 68, "threadPriority": 5, "messageTimestamp": "2025-09-25T12:37:30.864+0530", "ocLogId": "${ctx:ocLogId}"}
{"instant":
{"epochSecond": 1758784050, "nanoOfSecond": 864665000}, "thread": "boundedElastic-1", "level": "DEBUG", "loggerName": "ocpm.bsf.api.management.messageprofile.PriorityHeaderInboundRequestFilter", "message": "No matching message profile is found, ignoring outgoing server response header processing for 3gpp-sbi-message-priority priority header", "endOfBatch": false, "loggerFqcn": "org.apache.logging.slf4j.Log4jLogger", "threadId": 68, "threadPriority": 5, "messageTimestamp": "2025-09-25T12:37:30.864+0530", "ocLogId": "${ctx:ocLogId}"}

```



```
{
  "instant":
  {
    "epochSecond":1758784050,"nanoOfSecond":864712000,"thread":"boundedElastic-1",
    "level":"DEBUG","loggerName":"ocpm.bsf.api.management.metrics.BsfMetrics","message":
    "Pegging ocbsf_http_message_priority_response_count_total with direction : out, rcvd_priority: NA, eval_priority: NA, interfaceType: Nbsf_Management, msgType: Register, responseCode: 201",
    "endOfBatch":false,"loggerFqcn":"org.apache.logging.slf4j.Log4jLogger","threadId":68,
    "threadPriority":5,"messageTimestamp":"2025-09-25T12:37:30.864+0530",
    "ocLogId":"${ctx:ocLogId}"
  }
}
```

Client filter logging

```
{
  "instant":
  {
    "epochSecond":1758783713,"nanoOfSecond":823759000,"thread":"Thread-0","level":
    "INFO","loggerName":"ocpm.bsf.api.management.service.BsfConfigServiceImpl","message":
    "configurationChanged attempt, topic: common.bsf.messageProfiles, item: class ConfigurationItem { key: Audit_Contract_1, value: {
    \"messageProfileName\": \"bsf-mpp-audit1\", \"serviceType\": \"BSF_Management\", \"messageProfileAttributes\": [
    { \"messageID\": \"DEFAULT_AUDIT_ALWAYS\", \"rulePriority\": 1, \"interfaceType\": \"Npcf_Binding_Audit\", \"messageType\": \"Audit_Notify\", \"messageProfileFilterConditions\": [
    [ ], \"requestPriorityOverride\": \"Always\", \"requestPriorityValue\": 10, \"responsePriorityAssign\": \"Never\" },
    { \"messageID\": \"DEFAULT_AUDIT_IF_NOT_PRESENT\", \"rulePriority\": 2, \"interfaceType\": \"Npcf_Binding_Audit\", \"messageType\": \"Audit_Notify\", \"messageProfileFilterConditions\": [
    [ ], \"requestPriorityOverride\": \"IfNotPresent\", \"requestPriorityValue\": 11, \"responsePriorityAssign\": \"Never\" },
    { \"messageID\": \"DEFAULT_AUDIT_NEVER\", \"rulePriority\": 3, \"interfaceType\": \"Npcf_Binding_Audit\", \"messageType\": \"Audit_Notify\", \"messageProfileFilterConditions\": [
    [ ], \"requestPriorityOverride\": \"Never\", \"responsePriorityAssign\": \"Never\" } ] ] },
    version: 0, topicInfo: class TopicInfo { id: 77, name: common.bsf.messageProfiles, description: Created by config server., modifydate: null, version: 89 }, labels: null }",
    "endOfBatch":false,"loggerFqcn":"org.apache.logging.slf4j.Log4jLogger","threadId":28,
    "threadPriority":5,"messageTimestamp":"2025-09-25T12:31:53.823+0530",
    "ocLogId":"${ctx:ocLogId}"
  }
  {
    "instant":
    {
      "epochSecond":1758783713,"nanoOfSecond":824031000,"thread":"Thread-0","level":
      "INFO","loggerName":"ocpm.bsf.api.management.service.BsfConfigServiceImpl","message":
      "MessageProfiles list updated for : bsf-mpp-audit1",
      "endOfBatch":false,"loggerFqcn":"org.apache.logging.slf4j.Log4jLogger","threadId":28,
      "threadPriority":5,"messageTimestamp":"2025-09-25T12:31:53.824+0530",
      "ocLogId":"${ctx:ocLogId}"
    }
    {
      "instant":
      {
        "epochSecond":1758783727,"nanoOfSecond":848469000,"thread":"Thread-0","level":
        "INFO","loggerName":"ocpm.cne.common.configclient.ConfigurationAgent","message":
        "Configuration changed. Topic:[bsf.managementservice], Key:[bsf.managementservice], Value:[{
        \"audit\": {
        \"enable\": true, \"notificationRate\": \"100\", \"ttl\": \"55\", \"maxTtl\": \"155\",
        \"frequency\": \"10\", \"queryToPcf\": true, \"vendorId\": \"vendor1\" },
        \"messagePriority\": { \"messagePriorityProfile\": \"bsf-mpp-audit1\" } } ]",
        version: 1, "endOfBatch":false,"loggerFqcn":"org.apache.logging.slf4j.Log4jLog
      }
    }
  }
}
```

```

ger", "threadId": 28, "threadPriority": 5, "messageTimestamp": "2025-09-25T12:32:07.848+0530", "ocLogId": "${ctx:ocLogId}" }
{"instant":
{"epochSecond": 1758783724, "nanoOfSecond": 148161000}, "thread": "BSFManagementService_ThreadPool_1", "level": "INFO", "loggerName": "ocpm.bsf.api.management.messageprofile.PriorityHeaderClientFilter", "message": "outgoing request priority headers: {Accept=[application/problem+json, application/json], Content-Type=[application/json], 3gpp-sbi-target-apiroot=[http://dummyURL]}", "endOfBatch": false, "loggerFqcn": "org.apache.logging.slf4j.Log4jLogger", "threadId": 70, "threadPriority": 5, "messageTimestamp": "2025-09-25T12:32:04.148+0530", "ocLogId": "${ctx:ocLogId}" }
{"instant":
{"epochSecond": 1758783724, "nanoOfSecond": 149078000}, "thread": "BSFManagementService_ThreadPool_1", "level": "INFO", "loggerName": "ocpm.bsf.api.management.messageprofile.PriorityHeaderClientFilter", "message": "final outgoing request priority headers: {Accept=[application/problem+json, application/json], Content-Type=[application/json], 3gpp-sbi-target-apiroot=[http://dummyURL], 3gpp-sbi-message-priority=[10]}", "endOfBatch": false, "loggerFqcn": "org.apache.logging.slf4j.Log4jLogger", "threadId": 70, "threadPriority": 5, "messageTimestamp": "2025-09-25T12:32:04.149+0530", "ocLogId": "${ctx:ocLogId}" }
{"instant":
{"epochSecond": 1758783724, "nanoOfSecond": 225723000}, "thread": "HttpLoggingJettyHttpClient@6ef0bb37-75", "level": "INFO", "loggerName": "ocpm.bsf.api.management.service.AuditServiceHelperImpl", "message": "Response code OK and header {3gpp-sbi-message-priority=[8], Connection=[keep-alive], Content-Length=[0]} received from PCF for notification request for id 7aea9548-4a38-4381-8934-49abf08bda38 from URI: http://dummyURL/staleSessionVerificationPath", "endOfBatch": false, "loggerFqcn": "org.apache.logging.slf4j.Log4jLogger", "threadId": 75, "threadPriority": 5, "messageTimestamp": "2025-09-25T12:32:04.225+0530", "ocLogId": "${ctx:ocLogId}" }
{"instant":
{"epochSecond": 1758783724, "nanoOfSecond": 226155000}, "thread": "HttpLoggingJettyHttpClient@6ef0bb37-75", "level": "INFO", "loggerName": "ocpm.bsf.api.management.messageprofile.PriorityHeaderClientFilter", "message": "incoming client response priority headers: {3gpp-sbi-message-priority=[8], Connection=[keep-alive], Content-Length=[0]}", "endOfBatch": false, "loggerFqcn": "org.apache.logging.slf4j.Log4jLogger", "threadId": 75, "threadPriority": 5, "messageTimestamp": "2025-09-25T12:32:04.226+0530", "ocLogId": "${ctx:ocLogId}" }

```

3.2 Application Framework Change

BSF supports Micronaut from 25.2.200 onwards.

With this feature, Spring Boot has been replaced with Micronaut as the framework for microservices. Following are the impacts as part of this feature:

- Updated name of the following Diameter Gateway Metrics:

Table 3-1 Updates in Diameter Gateway Metrics

Old Metric Name	New Metric Name
ocbsf_diam_pending_requests_total	ocbsf_diam_pending_requests

Table 3-1 (Cont.) Updates in Diameter Gateway Metrics

Old Metric Name	New Metric Name
ocbsf_diam_service_overall_processing_time_seconds	ocbsf_diam_service_overall_processing_time_seconds_bucket

- Updated name of the following BSF Management Metrics:

Table 3-2 Updates in BSF Management Metrics

Old Metric Name	New Metric Name
spring_data_repository_invocations_seconds_sum	data_repository_invocations_seconds_sum
spring_data_repository_invocations_seconds_count	data_repository_invocations_seconds_count
spring_data_repository_invocations_seconds_max	data_repository_invocations_seconds_max

- Removed the following metrics:
 - ocbsf_diam_response_latency_seconds
 - ocbsf_diam_service_overall_processing_time_seconds

Note

This metric, `ocbsf_diam_service_overall_processing_time_seconds`, is removed. In order to get quantile data, use Prometheus Histogram (which contains bucket). You can use `histogram_percentile` from PQL to get much more accurate data from the bucket as per the use cases. Use this metric, `ocbsf_diam_service_overall_processing_time_seconds_bucket` to fetch the quantile data.

- Removed the `global.isIpvSixSetup` parameter.

3.3 Support for Optimizing N7 Session Lookup for AAR messages in Rx Call Flows

BSF supports optimize the N7 session lookup for AAR-I messages in Rx call flows. BSF stores the cookie received in the binding registration request from Policy.


Whenever BSF receives a AAR-I from an AF, it searches for the PCF identity of the Rx session and the cookie stored in the database. If the details are present, BSF Management Service sends the details as a header in the request to Diameter Gateway. The Diameter Gateway uses the details in the header to prepare `ocnf-service-cookie` Custom AVP. Diameter Gateway sends the `ocnf-service-cookie` Custom AVP in the AAR-I request to Policy. These details enable Policy to identify the session owner and the session identifier to further process the AAR-I request.

Data propagation from Policy and storage in BSF Management Service

During binding registration request, BSF receives the N7 session identifier details as a `ocnf-service-cookie` in the header of the HTTP request from Policy. The `ocnf-service-cookie` header includes session owner and session identifier such as `SMPolicyAssociationID` details.

If ENABLE_PCF_COOKIE Advanced Settings key in BSF Management Service is enabled, BSF stores the cookie data in `ocpm_bsf.pcf_binding` table in BSF Management Service database.

Figure 3-3 ocnf-service-cookie received by BSF from Policy



Sample `ocnf-service-cookie` header that BSF receives from Policy is as follows:


```

ocnf-service-cookie:{
  "cookieOwner": "PCF-SM",
  "cookieData": {
    "SmPolicyAssociationId": "xxx"
  }
}
  
```

Data lookup in BSF Management Service

Whenever BSF receives a AAR-I request from an AF, if ENABLE_PCF_COOKIE Advanced Settings key in BSF Management Service is enabled, the Diameter Gateway performs initial lookup of the PCF identity by querying BSF-Management Service database.

If the Rx session details are already stored in the `ocpm_bsf.pcf_binding` table in Management Service database, Diameter Gateway retrieves the details, encodes the details, and adds them to a custom AVP called `OCNF-SERVICE-COOKIE`. BSF includes this `OCNF-SERVICE-COOKIE` AVP in the AAR-I message and sends the message to Policy.

Figure 3-4 OCNF-SERVICE-COOKIE sent by BSF to Policy

Policy extracts the details in OCNF-SERVICE-COOKIE AVP and uses it to identify the session owner and session identifier in order to further process the AAR-I message.

For more details on how Policy processes the cookie details, see *Optimizing N7 Session Lookup for AAR messages in Rx Callflows* section in *Oracle Communications Cloud Native Core, Converged Policy User Guide*.

Upgrade Impact

For this feature to work, make sure that both BSF and Policy are upgraded to 25.2.100 or later versions.

Managing Support for Optimizing N7 Session Lookup for AAR messages in Rx Call Flows

Enable and Configure

ENABLE_PCF_COOKIE Advanced Settings key in Management Service page on CNC Console is used to enable support for optimizing N7 session lookup for AAR messages in Rx call flows. For more information, see [Management Service](#).

Observability

Metrics

The following metrics are used to support optimizing N7 session lookup for AAR messages in Rx call flows:

- ocbsf_pcf_cookie_total
- ocbsf_ocnf_service_cookie_total


For more information, see [BSF Management Service](#).

3.4 BSF Management Service Pod Congestion Control

BSF Management service interacts with other services such as PCF-Binding, Audit, Diameter Gateway services and also with the database. At times, excessive traffic from consumer services toward BSF Management service can be observed in the network, which can result in

high CPU utilization and high memory utilization. The service gets congested and can cause performance degradation in its service response time and eventually reach a state of service unavailability. To handle this, the service implements congestion control mechanism and ensures that its service is protected from congestion and, also its service is consistently available to the consumer services.

Figure 3-5 BSF Management Service interactions with other Services



The Pod Congestion Control feature in BSF application is a crucial mechanism to ensure consistent service availability and performance, especially in scenarios with high traffic and resource utilization. It is designed to identify congestion conditions and apply load shedding rules to maintain optimal service levels.

The BSF Management service handles all incoming requests. The congestion control mechanism evaluates the congestion state of the pods and the defined request priority to determine whether a request should be accepted or rejected. When congestion persists in the BSF Management pods, the service triggers the load-shedding rules. The pod congestion control mechanism involves:

1. Determining Pod Congestion State
2. Triggering Pod Congestion Control

Determining Pod Congestion State

In congestion control mechanism different congestion levels are defined, each level indicates how congested the pod is, and which requests can be served when pod is on certain congestion level. The congestion control works at following five congestion levels or states:

- NORMAL
- DANGER_OF_CONGESTION (DOC)
- CONGESTION_L1
- CONGESTION_L2
- CONGESTED

Each BSF Management service pod monitors its congestion state. The pod's congestion state is decided based on CPU consumption and Queue.

1. **CPU:** The CPU usage for congestion state is calculated by comparing the CPU usage of the container (monitored using cgroup parameter, `cpuacct.usage`, which provides current cpu usage in nanoseconds) with the configured threshold.
2. **Queue:** For the DOC, CONGESTION_L1, CONGESTION_L2, and CONGESTED pod states, compare the number of pending messages in the queue with the configured pending messages threshold.

To avoid toggling between these states due to traffic pattern, it is required for the pod to be in a particular state for a given period before transitioning to another state. The below configurations are used to define the period that the pod has to be in a particular state for:

- `stateChangeSampleCount`: This REST API parameter can be configured to specify after how many continuous intervals, the pod state can be changed. This value can range from 1 to 2147483647.
- `stateCalculationInterval`: This REST API parameter can be configured to specify the time duration or interval, after which the pod congestion state will be re-verified. This interval is configured in milliseconds and can range from 50 to 2147483647.

A pod's state changes when its resources, namely the CPU and Queue buffer usage exceeds the predefined congestion threshold, checked at regular intervals. An active CPU and Queue default threshold profile can be viewed using CNC Console at *BSF → Overload and Congestion Control → Congestion Control → Thresholds* section. The user can create new threshold profile and make it active. The BSF Management service pod's different congestion states and their default congestion parameters, CPU, and Queue pending counts are provided in the following table:

Table 3-3 BSF Management Service Congestion States

Congestion States	CPU (%)	Queue Count (Pending Requests)
DANGER_OF_CONGESTION (DOC)	75	175
CONGESTION_L1	80	200
CONGESTION_L2	85	225
CONGESTED	90	250

For more information about congestion control threshold profiles, see section in CNC Console.

Triggering Pod Congestion Control

With Congestion Control feature enabled for BSF Management service, every time a message is fetched for processing, the system checks the current congestion state of the pod. If the current state is either in DOC, CONGESTION_L1, CONGESTION_L2, or Congested state, then the congestion control mechanism is triggered. After verifying that the message type is a request, a request priority is assigned to it. If the request priority is greater than or equal to the discard priority, then the message is rejected, otherwise it is accepted.

Congestion Control Load Shedding using Discard and Request Priority

The BSF Management pods performs load shedding by considering the load rule configured for the current congestion state. The load shedding rule is based on message priority configuration. A default load shedding rule is active and the messages are discarded based on it. You can configure these load rules using the Overload and Congestion Control menu in CNC Console. The priority value ranges between 0 to 31, with 0 being the highest, and 31 being the lowest priority. The default discard priority for load shedding rules used in BSF Management service:

Table 3-4 BSF Management Service Default Discard Priorities:

Congestion State	Discard Priority (Default Value)
DANGER_OF_CONGESTION	30
CONGESTION_L1	24
CONGESTION_L2	18
CONGESTED	14

You can add the following keys to configure the message priority for different rest endpoints for BSF Management Service using CNC Console at *BSF* → *Service Configurations* → *Management Service* → *Advanced Settings* section. The message priority value ranges between 0 to 31, with 0 being the highest, and 31 being the lowest priority. If an invalid value is provided, then the default priority shall be considered.

Table 3-5 BSF Management Service Request Default Priorities

Message Type	Advanced Settings Key	Default Priority
Get Pcf Bindings	BSF_MGMT_GET_PCF_BINDIN G_REQUEST_PRIORITY	15
Post Pcf Bindings	BSF_MGMT_POST_PCF_BINDI NG_REQUEST_PRIORITY	20
Delete Pcf Bindings	BSF_MGMT_DELETE_PCF_BIN DING_REQUEST_PRIORITY	16
Post Audit Notification	BSF_MGMT_POST_AUDIT_NOT IFY_REQUEST_PRIORITY	25
Delete Audit Notification	BSF_MGMT_DELETE_AUDIT_N OTIFY_REQUEST_PRIORITY	25
Get Session Viewer Pcf Bindings	BSF_MGMT_GET_SESSION_VI EWER_PCF_BINDING_REQUE ST_PRIORITY	28
Post Session Viewer Pcf Bindings Cleanup	BSF_MGMT_POST_SESSION_ VIEWER_DELETE_PCF_BINDIN G_REQUEST_PRIORITY	18
Delete Session Viewer Pcf Bindings Cleanup All	BSF_MGMT_DELETE_SESSION _VIEWER_DELETEALL_PCF_BI NDING_REQUEST_PRIORITY	18

The following points needs to be known while assigning request or message priority to an request API:

- For requests coming from Diameter Gateway, the requests priority number is retrieved from *oc-message-priority* parameter in the http request header. But if this request header parameter is not present or its value is not set, then the message priority shall be based on the configured values from Advanced Settings or Default priority values.
- For requests coming from Ingress Gateway, the priority is based on *3gpp-sbi-message-priority* parameter in http request header. But if this request header parameter is not present or its value is not set, then the priority shall be based on configured values from Advanced Settings or on the default priority values.
- The request priority value ranges between 0 to 31, with 0 being the highest, and 31 being the lowest priority. If an invalid value is provided, then the default priority shall be considered.

Thus the precedence order for assigning message priority shall be from oc-message-priority → 3gpp-sbi-message-priority → Advanced Settings Priority → Default Priority.

You can also configure the result codes for the rejected requests when configuring the load rules. By default, BSF Management service sends 503 as response code with error cause as 'NF_CONGESTION'. You can configure the response code and cause using CNC Console **SBI Error Codes** configuration in General Settings. The response code configured should be **4xx** or **5xx** error status only.

In the response body, an error header is sent to the consumer service, providing the congestion error details as understood by the error-handler framework. A sample JSON error response:

```
{
  "errorDetails":{
    "status":"503",
    "cause":"NF_CONGESTION",
    "message":"Pod <pod-name> is in CONGESTION_L1",
    "targetService":"BSF_MGMT"
  }
}
```

BSF Management service Pod Congestion Control

Enable

By default, the Pod Congestion Control is disabled for BSF Management service. You can enable this feature using CNC Console or REST API for Policy.

Configure Using CNC Console

To enable the feature using CNC Console, set the **Enable** parameter in **Settings** page under **Congestion Control** for **Overload and Congestion Control Configurations**.

To configure UE service pod congestion control feature in CNC Console, see section.

Configure Using REST API

Perform the feature configurations as described in "Congestion Control" section in *Oracle Communications Cloud Native Core, Binding Support Function REST Specification Guide*.

Observability

Metrics:

The following common service metrics from [Pod Congestion Metrics](#) are used to provide information about this feature:

- pod_congestion_state
- pod_resource_stress
- pod_resource_congestion_state
- pod_cong_state_report_total
- pod_resource_congestion_state_report_total
- congestion_http_pending_request_active_counter

The `ocbsf_http_congestion_message_reject_total` metric was added to [BSF Management Service](#) section.

Alerts

The following alerts are generated for this feature:

- [POD_CPU_DOC](#)
- [POD_DOC](#)
- [POD_CONGESTED](#)
- [POD_CPU_CONGESTED](#)
- [POD_CONGESTION_L1](#)
- [POD_CPU_CONGESTION_L1](#)
- [POD_CONGESTION_L2](#)
- [POD_CPU_CONGESTION_L2](#)

3.5 Congestion Control Discard Message Response Code

In Congestion Control mechanism, on Pod being in any of the congested states the messages are discarded based on the message priority. The discarded messages responses can have customized error response. In BSF application, the services supporting the Congestion Control mechanism can configure the discarded messages error response in their service specific Advanced Settings page in CNC Console.

The following table lists the supported HTTP response code:

Table 3-6 Supported HTTP Response Codes

Series	Code	Cause	Type
1xx	101	Continue	INFORMATIONAL
1xx	102	Switching Protocols	INFORMATIONAL
1xx	103	Processing	INFORMATIONAL
1xx	104	Early Hints	INFORMATIONAL
3xx	300	Multiple Choices	REDIRECTION
3xx	301	Moved Permanently	REDIRECTION
3xx	302	Found	REDIRECTION
3xx	303	See Other	REDIRECTION
3xx	304	Not Modified	REDIRECTION
3xx	305	Use Proxy	REDIRECTION
3xx	307	Temporary Redirect	REDIRECTION
3xx	308	Permanent Redirect	REDIRECTION
4xx	400	Bad Request	CLIENT_ERROR
4xx	401	Unauthorized	CLIENT_ERROR
4xx	402	Payment Required	CLIENT_ERROR
4xx	403	Forbidden	CLIENT_ERROR
4xx	404	Not Found	CLIENT_ERROR
4xx	405	Method Not Allowed	CLIENT_ERROR
4xx	406	Not Acceptable	CLIENT_ERROR

Table 3-6 (Cont.) Supported HTTP Response Codes

Series	Code	Cause	Type
4xx	407	Proxy Authentication Required	CLIENT_ERROR
4xx	408	Request Timeout	CLIENT_ERROR
4xx	409	Conflict	CLIENT_ERROR
4xx	410	Gone	CLIENT_ERROR
4xx	411	Length Required	CLIENT_ERROR
4xx	412	Precondition Failed	CLIENT_ERROR
4xx	413	Payload Too Large	CLIENT_ERROR
4xx	414	URI Too Long	CLIENT_ERROR
4xx	415	Unsupported Media Type	CLIENT_ERROR
4xx	416	Requested range not satisfiable	CLIENT_ERROR
4xx	417	Expectation Failed	CLIENT_ERROR
4xx	418	I'm a teapot	CLIENT_ERROR
4xx	419	Insufficient Space On Resource	CLIENT_ERROR
4xx	420	Method Failure	CLIENT_ERROR
4xx	421	Destination Locked	CLIENT_ERROR
4xx	422	Unprocessable Entity	CLIENT_ERROR
4xx	423	Locked	CLIENT_ERROR
4xx	424	Failed Dependency	CLIENT_ERROR
4xx	425	Too Early	CLIENT_ERROR
4xx	426	Upgrade Required	CLIENT_ERROR
4xx	428	Precondition Required	CLIENT_ERROR
4xx	429	Too Many Requests	CLIENT_ERROR
4xx	431	Request Header Fields Too Large	CLIENT_ERROR
4xx	451	Unavailable For Legal Reasons	CLIENT_ERROR
5xx	500	Internal Server Error	SERVER_ERROR
5xx	501	Not Implemented	SERVER_ERROR
5xx	502	Bad Gateway	SERVER_ERROR
5xx	503	Service Unavailable	SERVER_ERROR
5xx	504	Gateway Timeout	SERVER_ERROR
5xx	505	HTTP Version not supported	SERVER_ERROR
5xx	506	Variant Also Negotiates	SERVER_ERROR
5xx	507	Insufficient Storage	SERVER_ERROR
5xx	508	Loop Detected	SERVER_ERROR
5xx	509	Bandwidth Limit Exceeded	SERVER_ERROR
5xx	510	Not Extended	SERVER_ERROR
5xx	511	Network Authentication Required	SERVER_ERROR

Note

If the congestion response code other than above supported HTTP response codes is configured, then the default response code 503 shall be assigned and below log message will be logged:

Log line => Exception occurred: No matching constant for '512' Rejecting request with default response code: 503.

The following table lists the supported Diameter result codes:

Table 3-7 Supported Diameter Result Codes

Series	Code	Name
3xxx	3001	DIAMETER_COMMAND_UNSUPPORTED
3xxx	3002	DIAMETER_UNABLE_TO_DELIVER
3xxx	3003	DIAMETER_REALM_NOT_SERVED
3xxx	3004	DIAMETER_TOO_BUSY
3xxx	3005	DIAMETER_LOOP_DETECTED
3xxx	3006	DIAMETER_REDIRECT_INDICATION
3xxx	3007	DIAMETER_APPLICATION_UNSUPPORTED
3xxx	3008	DIAMETER_INVALID_HDR_BITS
3xxx	3009	DIAMETER_INVALID_AVP_BITS
3xxx	3010	DIAMETER_UNKNOWN_PEER
4xxx	4001	DIAMETER_AUTHENTICATION_REJECTED
4xxx	4002	DIAMETER_OUT_OF_SPACE
4xxx	4003	ELECTION_LOST
4xxx	4010	DIAMETER_END_USER_SERVICE_DENIED
4xxx	4011	DIAMETER_CREDIT_CONTROL_NOT_APPLICABLE
4xxx	4012	DIAMETER_CREDIT_LIMIT_REACHED
4xxx	4999	SPLIT_BRAIN_SUSPECTED
5xxx	5001	DIAMETER_AVP_UNSUPPORTED
5xxx	5003	DIAMETER_AUTHORIZATION_REJECTED
5xxx	5004	DIAMETER_INVALID_AVP_VALUE
5xxx	5005	DIAMETER_MISSING_AVP
5xxx	5006	DIAMETER_RESOURCES_EXCEEDED

Table 3-7 (Cont.) Supported Diameter Result Codes

Series	Code	Name
5xxx	5007	DIAMETER_RESOURCES_EXCEEDED
5xxx	5008	DIAMETER_AVP_NOT_ALLOWED
5xxx	5009	DIAMETER_AVP_OCCURS_TOO_MANY_TIMES
5xxx	5010	DIAMETER_NO_COMMON_APPLICATION
5xxx	5011	DIAMETER_UNSUPPORTED_VERSION
5xxx	5012	DIAMETER_UNABLE_TO_COMPLY
5xxx	5013	DIAMETER_INVALID_BIT_IN_HEADER
5xxx	5014	DIAMETER_INVALID_AVP_LENGTH
5xxx	5015	DIAMETER_INVALID_MESSAGE_LENGTH
5xxx	5016	DIAMETER_INVALID_AVP_BIT_COMBO
5xxx	5017	DIAMETER_NO_COMMON_SECURITY
5xxx	5030	DIAMETER_USER_UNKNOWN
5xxx	5031	DIAMETER_RATING_FAILED
5xxx	5032	DIAMETER_ERROR_BYTE

Note

If the congestion response code other than the above supported diameter result codes is configured, then the default response code 3004 shall be assigned, and below logs message will be logged:

Log line => Configured 1234: {} does not belongs to 3xxx, 4xxx or 5xxx series. using default response code: 3004

The following table lists the supported Experimental result code:

Table 3-8 Experimental Result Codes

Series	Code	Name	Vendor Id	Vendor Name
5xxx	5011	DIAMETER_ERROR_FEATURE_UNSUPPORTED	10415	3GPP
5xxx	5997	NO_CAPACITY	21274	Telelec
5xxx	5453	DIAMETER_ERROR_LATE_OVERLAPPING_REQUEST	10415	3GPP

Table 3-8 (Cont.) Experimental Result Codes

Series	Code	Name	Vendor Id	Vendor Name
5xxx	5454	DIAMETER_ERRO R_TIMED_OUT_R EQUEST	10415	3GPP

3.6 Stale Binding Detection Audit, Report and Recover

Service disruption due to network storm, system overload, database latency, and other events can impact signaling between PCF and BSF. This service disruption can affect session binding between PCF and BSF.

BSF revalidates the binding information of a PDU session and checks if there is any missing binding information in BSF due to service disruption between PCF and BSF.

When BSF receives a binding session revalidation request from PCF, BSF processes the request, determines if the session is available in BSF. Existence of the binding association for the PDU session in BSF confirms the binding association being valid in BSF. If the binding association is missing in BSF, it is restored by creating the association in BSF.

When the session binding revalidation is enabled in BSF

When BSF receives a PDU session binding revalidation request from PCF through Ingress Gateway:

- BSF checks if session binding revalidation is enabled in BSF and the revalidation request includes x-oc-binding-revalidation header.

Note

Session binding revalidation can be enabled in BSF, either using **Enable Binding Revalidation** field in **Management Service** page on CNC Console or using `enableBindingRevalidation` parameter in BSF Management Service API. For more details, see [Management Service](#).

- If BSF identifies that the revalidation request is a colliding or a duplicate request, and the if the following parameters in the request matches with the stored binding object in BSF:
 - SUPI/GPSI
 - UE IP(IPV4/IPV6),
 - IPDomain(if IPV4 was present for UE IP)
 - dnn
 - snssai
 - diamHost/fqdn
 - pcfFqdn

A binding revalidation request is considered as a duplicate request if the request is for the same subscriber (SUPI, DNN, SNSSAI, IPV4/IPV6 prefix, or IpDomain) from same PCF instances (ID/pcfDiamHost/pcfFqdn or pcfIpEndpoints). If the parameters in the request matches with the stored binding object in BSF, BSF responds to PCF with the given `binding_id` and does not perform any database update.

- If the parameters do not match, or if it is not a duplicate request, BSF creates the binding record and responds to PCF with a 201 response to PCF including a location HTTP header field that contains the URI of the created binding information.

Example:

```
{apiRoot}/oc-bsf-configuration/v1/services/management-service/{bindingId}
```

where, {bindingId} is the new binding identifier for the session.

When the session binding revalidation is enabled in BSF

If the session binding revalidation is enabled in BSF, when BSF receives a PDU session revalidation request from PCF, BSF checks if the binding information exists in BSF. If the record created by the same PCF already exists, BSF deletes the existing record and creates a new record for the PDU session.

When BSF receives a PDU session revalidation request from PCF:


- BSF verifies if the parameters for late arrival handling is configured.
- If the request has not arrived late, BSF checks if **Enable Collision Detection** field is enabled.

If the parameters in the request do not match with the stored binding information in BSF, BSF creates the binding record for the pcfd and responds to PCF with 201 response code with a new binding_id.

Call Flow

The following diagram depicts a sample call flow for revalidating the binding information for a PDU session in BSF when the binding revalidation is enabled in BSF and the revalidation request includes x-oc-binding-revalidation header:

Figure 3-6 Binding revalidation in BSF



1. BSF receives a binding revalidation request for a PDU session from PCF.
2. BSF checks if the session binding revalidation is enabled in BSF and the revalidation request includes x-oc-binding-revalidation header.
3. If the revalidation request is a duplicate or a colliding request and if the binding information for the PDU session already exists in BSF, BSF sends a 201 response code with the existing binding_id.
4. If the binding information stored in BSF differs from the details received in the request or if the binding information for the PDU session does not exist in BSF, BSF recreates the session binding record (re-registers) for the PDU session. BSF sends a 201 response code with the new binding_id.

Managing Stale Binding Detection Audit, Report and Recover

Enable

This feature can be enabled using CNC Console for BSF as well as using REST API.

Enable using CNC Console

To enable this feature using CNC Console for BSF, configure **Enable Binding Revalidation** field on **Management Service** page under **Service Configurations** in CNC Console for BSF.

For more information, see [Management Service](#).

Enable using REST API

To enable this feature using REST API, configure the EnableBindingrevalidate parameter under Management Service API for BSF: {apiRoot}/oc-bsf-configuration/v1/services/management-service.

For more information, see *Management Service* section in *Oracle Communications Cloud Native Core, Binding Support Function REST Specification Guide*.

Observability

Metrics

The following BSF Management service metrics are used to monitor this feature in BSF:

- ocbsf_binding_revalidation_request_total
- ocbsf_binding_revalidation_response_total
- ocbsf_binding_revalidation_pcfBinding_missing_total

For more information, see [BSF Management Service](#).

Alerts

The following BSF alerts are used for this feature:

- BINDING_REVALIDATION_PCF_BINDING_MISSING_MINOR
- BINDING_REVALIDATION_PCF_BINDING_MISSING_MAJOR
- BINDING_REVALIDATION_PCF_BINDING_MISSING_CRITICAL

For more information, see [List of Alerts](#).

Logging

The following logs with a specific marker "marker":{"name":"BINDING"}, "Binding Audit:<additional log text>" are generated for this feature:

- For Revalidation request, if the binding association is not present in BSF, BSF updates its logs with the failure for the revalidation request. The updated logs include:
 - SUPI
 - DNN
 - S-NSSAI
 - UE-IP (IPv6/IPv4)
 - IPDomain if available
 - PCF FQDN
 - error code and cause at "WARN" level
- For revalidation request, if BSF can not reestablish or restore the BSF binding after all session retries, BSF updates its logs with the failure for the restoration request.
 - SUPI
 - DNN
 - S-NSSAI
 - UE-IP (IPv6/IPv4)
 - IPDomain if available
 - PCF FQDN
 - error code and cause at "WARN" level

3.7 Traffic Segregation

This feature provides end-to-end traffic segregation to BSF based on traffic types. Within a Kubernetes cluster, traffic segregation can divide applications or workloads into distinct sections such as OAM, SBI, Kubernetes control traffic, etc. The Multus CNI container network interface (CNI) plugin for Kubernetes enables attaching multiple network interfaces to pods to help segregate traffic from each BSF microservice.

This feature addresses the challenge of logically separating IP traffic of different profiles, which are typically handled through a single network (Kubernetes overlay). The new functionality ensures that critical networks are not cross-connected or sharing the same routes, thereby preventing network congestion.

With traffic segregation, operators can segregate traffic to external feeds and applications more effectively. Previously, all external traffic was routed through the same external network, but now, egress traffic from the BSF pods can be directed through non-default networks to third-party applications. This separation is achieved by leveraging cloud-native infrastructure and the load balancing algorithms in CNE.

The feature supports the configuration of separate networks, Network Attachment Definitions (NADs), and the Cloud Native Load Balancer (CNLB). These configurations are crucial for enabling cloud native load balancing, facilitating ingress-egress traffic separation, and optimizing load distribution within BSF.

Prerequisites

The CNLB feature is only available in BSF if CNE is installed with CNLB and Multus.

Cloud Native Load Balancer (CNLB)

CNE provides Cloud Native Load Balancer (CNLB) for managing the ingress and egress network as an alternate to the existing LBVM, lb-controller, and egress-controller solutions. You can enable or disable this feature only during a fresh CNE installation. When this feature is enabled, CNE automatically uses CNLB to control ingress traffic. To manage the egress traffic, you must preconfigure the egress network details in the `cnlb.ini` file before installing CNE.

Note

CNLB is supported only for IPv4 stack.

For more information about enabling and configuring CNLB, see *Oracle Communications Cloud Native Core, Cloud Native Environment User Guide*, and *Oracle Communications Cloud Native Core, Cloud Native Environment Installation, Upgrade, and Fault Recovery Guide*.

Network Attachment Definitions for CNLB

A Network Attachment Definition (NAD) is a resource used to set up a network attachment, in this case, a secondary network interface to a pod. BSF supports two types of CNLB NADs:

1. Ingress Network Attachment Definitions

Ingress NADs are used to handle inbound traffic only. This traffic enters the CNLB application through an external interface service IP address and is routed internally using interfaces within CNLB networks.

- **Naming Convention:** `nf-<service_network_name>-int`

2. Egress Only Network Attachment Definitions

Egress Only NADs enable outbound traffic only. An NF pod can initiate traffic and route it through a CNLB application, translating the source IP address to an external egress IP address. An egress NAD contains network information to create interfaces for NF pods and routes to external subnets.

- **Requirements:** Destination (egress) subnet addresses are known beforehand and defined under the `cnlb.ini` file's `egress_dest` variable to generate NADs.
- **Naming Convention:** `nf-<service_network_name>-egr`

3. Ingress/Egress Network Attachment Definitions

Ingress/Egress Network Attachment Definitions enable inbound/outbound traffic. An NF pod can initiate traffic and route it through a CNLB app, translating source IP address to an external egress IP address (defined under `cnlb.ini` file `egress_addr` variable). An Ingress/Egress Network Attachment Definition contains network information to create interfaces for NF pods and routes to external subnets. Even though an Ingress/Egress Network Attachment Definition enables outbound traffic, it also handles inbound traffic, so if inbound/outbound traffic is needed an Ingress/Egress Network Attachment Definition should be used.

- **Requirements:** Source (ingress) and destination (egress) subnet addresses are known beforehand and defined under `cnlb.ini` file `egress_dest` variable to generate Network Attachment Definitions.
- **Naming Convention:** `nf-<service_network_name>-ie`

Managing Ingress and Egress Traffic Segregation

Enable:

This feature is disabled by default. To enable this feature, you must configure the network attachment annotations in the custom values file.

Configuration

For more information about Traffic Segregation configuration, see "Configuring Traffic Segregation" section in *Oracle Communications Cloud Native Core, Converged Binding Support Function Installation, Upgrade, and Fault Recovery Guide*.

Observe

There are no Metrics, KPIs, or Alerts available for this feature.

Maintain

To resolve any alerts at the system or application level, see [BSF Alerts](#) section. If the alerts persist, perform the following:

1. **Collect the logs:** For more information on how to collect logs, see *Oracle Communications Cloud Native Core, Binding Support Function Troubleshooting Guide*.
2. **Raise a service request:** See [My Oracle Support](#) for more information on how to raise a service request.

3.8 Support for Dual Stack

Using the dual stack mechanism, applications or NFs can establish connections with pods and services in a Kubernetes cluster using either IPv4, IPv6, or both, depending on the configured deployment mode. Dual stack provides:

- Coexistence strategy that allows hosts to reach IPv4 and IPv6.
- IP address allocation to the Kubernetes clusters during cluster creation. This allocation is applicable for all Kubernetes resources unless explicitly specified during cluster creation.

① Note

- The dual stack feature is not supported on Oracle Cloud Native Environment (CNE) but is supported on third-party CNEs.
- Operator needs to redeploy (Fresh install) the BSF to enable IP Dual Stack functionality.

On a dual stack deployment, externally communicating services of type LoadBalancer such as Ingress Gateway and Diameter Gateway supports dual stack where as internally communicating services such as cnDBTier, CM, and so on, uses single stack based on the preference. Please refer the table below.

Table 3-9 Preference of Deployment Mode

Services	IPv6_Pv4	IPv4_IPv6	IPv6	IPv4
Internal communicating services	IPv6 (Single Stack)	IPv4 (Single Stack)	IPv6 (Single Stack)	IPv4 (Single Stack)
External communicating services	IPv6 preferred (Dual Stack)	IPv4 preferred (Dual Stack)	IPv6 (Single Stack)	IPv4 (Single Stack)

Note

Communication between services and CoreDNS, as well as the Kubernetes API, depends on the cluster's network preference.

IP Address Allocation to Pods

IP address allocation to pods depends on the IP address preference set in the Kubernetes cluster. Pods do not have the privilege to choose an IP address. Consider the following example of a pod deployed in an IPv4 preferred infrastructure. Here, if the Kubernetes cluster has IPv4 preferred configuration, both IPv4 and IPv6 are allocated to the pod, but the primary IP address is IPv4. Example:

```
IP:          10.xxx.xxx.xxx
IPs:
  IP:        10.xxx.xxx.xxx
  IP:        fd00::1:cxxx:bxxx:8xxx:xxxx
```

IP Address Allocation to Services

IP address allocation to all the BSF services, depends on the DeploymentMode Helm parameter configuration. This Helm parameter automatically configures *IP Family Policy* and *IP Families* attributes for the BSF services. For more information about DeploymentMode, see the "Customizing BSF" section in the *Oracle Communications Cloud Native Core, Binding Support Function Installation, Upgrade, and Fault Recovery Guide*.

You can customize the IP address allocation to services based on the DeploymentMode Helm parameter. Services route the traffic to the destination endpoints based on this configuration. If the DeploymentMode Helm parameter is set to IPv4, then IPv4 is allocated to services, and services use IPv4 pod IPs to send the traffic to endpoints.

The following table describes how IP address allocation, IP Family Policy, and IP Families vary based on the DeploymentMode Helm parameter configuration for services:

Table 3-10 IP Address Allocation

Infrastructure Preference	Application Preference (DeploymentMode Helm Parameter)	IP Family Policy Attribute	IP Families Attribute	Pod IP	Service IP	Endpoints
IPv4 Preferred	IPv4	SingleStack	IPv4	IPv4,IPv6	IPv4	IPv4
IPv6 Preferred	IPv4	SingleStack	IPv4	IPv6,IPv4	IPv4	IPv4

Table 3-10 (Cont.) IP Address Allocation

Infrastructure Preference	Application Preference (DeploymentMode Helm Parameter)	IP Family Policy Attribute	IP Families Attribute	Pod IP	Service IP	Endpoints
IPv4 Preferred	IPv6	SingleStack	IPv6	IPv4,IPv6	IPv6	IPv6
IPv6 Preferred	IPv6	SingleStack	IPv6	IPv6,IPv4	IPv6	IPv6
IPv4 Preferred	IPv4_IPv6 (IPv4Preferred)	RequiredDualStack	IPv4 Preferred	IPv4,IPv6	IPv4,IPv6	IPv4
IPv6 Preferred	IPv4_IPv6 (IPv4Preferred)	RequiredDualStack	IPv4 Preferred	IPv6,IPv4	IPv6,IPv4	IPv4
IPv4 Preferred	IPv6_IPv4 (IPv6Preferred)	RequiredDualStack	IPv6 Preferred	IPv4,IPv6	IPv4,IPv6	IPv6
IPv6 Preferred	IPv6_IPv4 (IPv6Preferred)	RequiredDualStack	IPv6 Preferred	IPv6,IPv4	IPv6,IPv4	IPv6

Enable

Dual Stack support can be enabled by setting the `global.deploymentMode` parameter to `IPv4_IPv6` or `IPv6_IPv4` depending on the preference needed. For more information, see the Oracle Communications Cloud Native Core, Binding Support Function Installation, Upgrade, and Fault Recovery Guide.

Configure**Note**

The `diam-gateway.dualStackPreferredDnsResolutionIpAddressType` parameter is replaced by `global.dgwRoutingMode`. The `diamgateway.envSupportedIpAddressType` parameter is no longer applicable.

You can configure dual stack using the following parameters:

- `global.deploymentMode`
- `global.egressRoutingMode`
- `global.dgwRoutingMode` controls the IP address selection for Diameter Gateway connections when DNS resolves a destination to multiple IP address types. The supported values are `IPv4`, `IPv6`, `IPv4_IPv6`, and `IPv6_IPv4`. The default value is `IPv4_IPv6`.

For more information on these parameters, see "Configuration Parameters for Dual Stack" in the *Oracle Communications Cloud Native Core, Binding Support Function Installation, Upgrade, and Fault Recovery Guide*.

Table 3-11 Gateway Routing Mode

dgwRoutingMode/ egressRoutingMode	Description	DNSQuery/k8s endpoints returns IPv4 address only	DNSQuery/k8s endpoints returns IPv6 address only	DNSQuery/k8s endpoints returns both IPv4 and IPv6 address
IPv4	Strict IPv4 only	Gateway attempts connection towards the destination	Gateway will not initiate connection towards the destination	Gateway will pick up the IPv4 address and attempts connection toward the destination
IPv6	Strict IPv6 only	Gateway will not initiate connection towards the destination	Gateway attempts connection towards the destination	Gateway will pick up the IPv6 address and attempts connection toward the destination
IPv4_IPv6	Supports both with IPv4 preferred	Gateway attempts connection towards the destination	Gateway attempts connection towards the destination	Gateway will pick up the IPv4 addresses only and attempts connection toward the destination
IPv6_IPv4	Supports both with IPv6 preferred	Gateway attempts connection towards the destination	Gateway attempts connection towards the destination	Gateway will pick up the IPv6 addresses only and attempts connection toward the destination
None (Applicable to egressRoutingMode only)	No changes. Egress Gateway will rely on library to select IPs.	NA	NA	NA

Table 3-12 Recommended Configurations for Dual Stack and Single Stack

Deployment		Deployment Mode	
global.deploymentMode		global.dgwRoutingMode	
IPv4		IPv4	

global.deploymentMode	global.dgwRoutingMode

Table 3-12 (Cont.) Recommended Configurations for Dual Stack and Single Stack

Deployment Mode		Deployment Mode	
global.deploymentMode		global.dgwRoutingMode	
IPv6		IPv6	

global.deploymentMode	global.dgwRoutingMode

Table 3-12 (Cont.) Recommended Configurations for Dual Stack and Single Stack

global.deploymentMode	global.dgwRoutingMode
IPv4_IPv6	IPv4_IPv6

Deployment Options	
global.deploymentMode	global.dgwRoutingMode

Table 3-12 (Cont.) Recommended Configurations for Dual Stack and Single Stack

global.deploymentMode	global.dgwRoutingMode
IPv6_IPv4	IPv6_IPv4

Table 3-12 (Cont.) Recommended Configurations for Dual Stack and Single Stack

global.deploymentMode	global.dgwRoutingMode

Table 3-12 (Cont.) Recommended Configurations for Dual Stack and Single Stack

global.deploymentMode	global.dgwRoutingMode
ClusterPreferred	IPv4_IPv6

[illegible]

Note

- Any combination of dual stack or IPv6 configurations other than recommended values shall be considered as invalid and not supported.
- Egress routing for automatic fallback between IPv4 and IPv6 is not supported. You should connect using their preferred IP address type. If that IP type is unavailable, connection to the peer is considered unavailable.

Upgrade details

During upgrade, remove the retired parameter and set `global.dgwRoutingMode` to the value that matches the required IP-family preference. If the previous value was both, select either `IPv4_IPv6` or `IPv6_IPv4` explicitly. Do not use `diamgateway.envSupportedIpAddressType`; it is no longer applicable.

Table 3-13 Upgrade details

<code>diam-gateway.dualStackPreferredDnsResolutionIpAddressType</code>	<code>global.dgwRoutingMode</code>
IPv4	IPv4
IPv6	IPv6
IPv4_IPv6	IPv4_IPv6
IPv6_IPv4	IPv6_IPv4
both	IPv4_IPv6 or IPv6_IPv4 depending on the preference

Observe

There are no Metrics, KPIs, or Alerts available for this feature.

3.9 Logging Support for Error Response

Error handling framework was introduced in earlier release of BSF as a general purpose error handling tool. This feature used the error handling framework to add more details to the error using 3GPP error response format.

BSF sends error responses to consumer NFs due to some exceptions, such as signaling, validations, and internal errors. These error responses have payloads containing the problem title, status, details, and cause of the error that are used to investigate the error. BSF has been enhanced to support logs for the error responses.

Currently, BSF supports the following log levels:

- ERROR:** Error is at a service or component level and not per subscriber or call flow level.
- WARN:** An error is not expected in a typical call flow, and there is a need to know if it occurred in production when troubleshooting an issue.
- INFO:** Informational logs showing what is happening in each service or component during a call flow. All in-out messages to a service should be logged at **INFO** level. It is also beneficial to log messages once during processing by a component within a service.
- DEBUG:** Provides detailed logs that helps to troubleshoot an issue.
- TRACE:** Includes detailed logging of messages in and out of each component or method.

When sending any error response triggered by an HTTP Request, this error response format will be mapped into the following *general log format*:

```
{
  "errorStatus": "Value", (ProblemDetails status field)
  "errorTitle": "Value", (ProblemDetails title field)
  "errorDetails": "Value", (ProblemDetails detail field)
  "errorCause": "Value" (ProblemDetails cause field)
  "sender": "Value", (nfType-nfInstanceId)
  "subscriberId": "Value" (UE ID associated with event if present)
}
```

Table 3-14 General Log Format

Parameter	Description	Example
errorStatus	Specifies the status code of the error.	404, 500, etc.
errorTitle	Specifies the title of the error.	Required parameter in binding data is missing
errorDetails	Specifies the error detail produced by error handling framework in case of BSF as producer. Note: For errorDetails field to be populated in the required format, error handling framework should be enabled in CNC Console for the required service.	ocbsf1-2-api-gateway.bsf1-2.svc.atlantic.morrisville.us.lab.oracle.com:BSF_MGMT:Mandatory parameter is missing in request :EC-0BSF-BSF_MGMT-REQVLD-EI-05-02-400-00010-01-02
errorCause	Specifies the cause of the error.	MANDATORY_IE_MISSING
sender	Specifies the sender which is composed by the nfType plus the instanceId. Note: In case of BSF as producer of error, the sender field will be BSF plus the instanceId of the BSF.	BSF-fe7d992b-0541-4c7d-ab84-c6d70b1b0666
subscriberId	Specifies the subscriberId which can be SUPI or GPSI associated with the event.	imsi-65008100001061

In the case of subscriberID, if UE identifier is present and shows UE identifier flag as *enabled*, the logging happens in the following hierarchy:

1. If SUPI and GPSI both are present, use SUPI.
2. If SUPI is not present, and GPSI is present, use GPSI.

Note

Sender-contains the *nfType-instanceId* from the sender of the HTTP Request. Where *nfType* is a value which contains the type of Network Function set as *BSF* by default, and *nfInstanceId* is a unique identifier for a specific instance of a Network Function. This value is set in the application configuration for BSF

Support for Burst rate limit modification and customization for logs

When the enhanced error response logging feature is enabled, due to certain error scenarios volume of error logs can become huge. In order to control the volume of logs, BSF allows to add a burst filter for logging against logs for all the microservices.

The Burst filter provides a mechanism to control the rate at which LogEvents are processed by silently discarding events after the maximum limit has been reached. It enables to control the frequency and volume of warning messages generated. This ensures that excessive logging does not overwhelm the system or obscure critical information.

Configuring `BurstFilter` property in `log4j` includes the following parameters:

- **Level:**
Determines the minimum log level at which the filter becomes active. Log events below this level are not affected.
Possible values are:
 - OFF
 - TRACE
 - DEBUG
 - INFO
 - WARN
 - ERROR
 - FATAL
 - ALLDefault value is `DEBUG`.
- **Rate:**
Defines the average number of log events allowed per second. Events beyond this rate are suppressed unless allowed by `maxBurst`.
It accepts any positive integer or decimal value.
Default value is 500 logs per second.
- **maxBurst:**
Specifies the maximum number of log events that can be logged in a sudden spike before suppression starts. Helps tolerate short bursts without dropping all messages.
It accepts any positive integer value.
Default value is 3000.
- **onMismatch:**
Defines the filter's behavior when a log event does not match the filter conditions such as below the level or exceeding the rate/burst limit.
Possible values are:
 - **DENY:** The event is suppressed and not passed to the appender.
 - **NEUTRAL:** The event is passed to the next filter in the chain for further evaluation.
 - **ACCEPT:** The event is always logged regardless of conditions.Default value is `DENY`.

Managing Enhancements to Logging Support for Error Response

This section explains the procedure to enable and configure the feature.

Enable

By default, this feature is disabled. The operator can enable this feature through the CNC Console configurations.

① Note

Oracle does not recommend enabling this feature. When this feature is enabled, sensitive information in the form of UE identifier gets published. Users can secure the logs in their environment using appropriate configuration.

Configure

You can configure logging support for error response using the **Enable Enhanced Logging** and **Enable UE Identifier Information** toggle buttons. These toggle buttons are available in **General Settings** under **General Configurations** on CNC Console for BSF. For information about how to configure for BSF Management Service in CNC Console, see [General Settings](#).

Observe

There are no new metrics in BSF Management Service for this feature.

3.10 Support for TLS

BSF uses Hypertext Transfer Protocol Secure (HTTPS) and Diameter Gateway to establish secured connections with consumer NFs and producer NFs, respectively. These communication protocols are encrypted using Transport Layer Security (TLS).

TLS comprises the following components:

- **Handshake Protocol:** Exchanges the security parameters of a connection.
- **Record Protocol:** Receives the messages to be transmitted, fragments the data into multiple blocks, secures the records, and then transmits the result. Received data is delivered to higher-level peers.

This feature enables the support for TLS 1.3 to all consumer NFs, producer NFs, the Data Director, SBI Interfaces, and any interfaces previously supporting TLS 1.2. Support for TLS 1.2 will remain available.


① Note

If both TLS 1.2 and TLS 1.3 are supported, TLS 1.3 is given priority.

TLS 1.2 and TLS 1.3 Handshake

This section describes the differences between TLS 1.3 and TLS 1.2, as well as the advantages of TLS 1.3 over TLS 1.2 and earlier versions.

Figure 3-7 TLS 1.2 and TLS 1.3 Handshake

**TLS 1.2**

Step 1: The connection or handshake starts when the client sends a "client hello" message to the server. This message consists of cryptographic information such as supported protocols and cipher suites. It also contains a random value or random byte string.

Step 2: To respond to the "client hello" message, the server sends a 'server hello' message. This message contains the CipherSuite that the server has selected from the options provided by the client. The server also sends its certificate, along with the session ID and another random value.

Step 3: The client verifies the certificate sent by the server. When the verification is complete, it sends a byte string encrypted using the public key of the server's certificate.

Step 4: When the server receives the secret, both the client and server generate a master key along with session keys (ephemeral keys). These session keys are used for symmetrically encrypting the data.

Step 5: The client sends an "HTTP Request" message to the server to enable the server to transition to symmetric encryption using the session keys.

Step 6: To respond to the client's "HTTP Request" message, the server does the same and switches its security state to symmetric encryption. The server concludes the handshake by sending an HTTP response.

Step 7: The client-server handshake is completed in two round trips.

TLS 1.3

Step 1: The connection or handshake begins when the client sends a "client hello" message to the server, which includes the list of supported cipher suites and the client's key share for the specific key agreement protocol.

Step 2: To respond to the "client hello" message, the server sends the key agreement protocol that it has chosen. The "Server Hello" message includes the server key share, server certificate, and the "Server Finished" message.

Step 3: The client verifies the server certificate, generates keys since it has the server's key share, and then sends the "Client Finished" message along with an HTTP request.

Step 4: The server completes the handshake by sending an HTTP response.

Note

ASM and TLS are not supported together.

The following table provides comparison of TLS 1.2 with TLS 1.3:

Table 3-15 Comparison of TLS 1.2 with TLS 1.3

Feature	TLS 1.2	TLS 1.3
TLS Handshake	This is less efficient as it requires more round-trips to complete the handshake process.	This is more efficient as it requires less round-trips to complete the handshake process.
Cipher Suites	This has less secured Cipher Suites.	This has more secured Cipher Suites. They support the following ciphers: • TLS_CHACHA20_POLY1305_SHA256 • TLS_AES_128_GCM_SHA256 • TLS_AES_256_GCM_SHA384 • TLS_AES_128_CCM_8_SHA256: This Cipher is not supported by Java library. • TLS_AES_128_CCM_SHA256: This Cipher is not supported by Java library.
Round-Trip Time (RTT)	This has higher RTT during TLS handshake.	This has low RTT.
Performance	This has higher latency during TLS handshake.	This has low latency during TLS handshake.

Note

- BSF does not prioritize cipher suites on the basis of priorities. To select cipher on the basis of priorities, you must write the cipher suites in the decreasing order of priority.
- BSF does not prioritize supported groups on the basis of priorities. To select supported group on the basis of priorities, you must write the supported group values in the decreasing order of priority.
- If you want to provide values for the `signature_algorithms` extension using the `clientSignatureSchemes` parameter, the following comma-separated values must be provided to deploy the pods:
 - `rsa_pkcs1_sha512`
 - `rsa_pkcs1_sha384`
 - `rsa_pkcs1_sha256`
- The mandatory extensions as listed in RFC 8446 cannot be disabled on the client or server side. The following is the list of the extensions that cannot be disabled:
 - `supported_versions`
 - `key_share`
 - `supported_groups`
 - `signature_algorithms`
 - `pre_shared_key`

The following digital signature algorithms of TLS 1.2 and TLS 1.3 are supported in TLS handshake:

Table 3-16 Digital Signature Algorithms

Algorithm	Key Size (Bits)	Elliptic Curve (EC)
RS256 (RSA)	2048	NA
	4096 This is the recommended value.	NA
ES256 (ECDSA)	NA	SECP384r1 This is the recommended value.

Note

The following functionalities from TLS 1.3 specifications are not supported:

- Zero round-trip time (0-RTT) mode.
- Pre-Shared Key (PSK) exchange.

3.10.1 Support for TLS Using HTTPs

BSF uses Hypertext Transfer Protocol Secure (HTTPS) to establish secured connections with consumer NFs and producer NFs, respectively. These communication protocols are encrypted using Transport Layer Security (TLS).

Enable

This feature is enabled by default at the time of Gateway Services deployment by completing the required Helm configurations.

Configure

You can configure this feature using Helm. For information about Helm configurations, see "Global Parameters" in *Oracle Communications Cloud Native Core, Binding Support Function Installation, Upgrade, and Fault Recovery Guide*.

Observe

The following metrics are available for this feature:

- oc_ingressgateway_incoming_tls_connections
- oc_egressgateway_outgoing_tls_connections
- security_cert_x509_expiration_seconds

For more information about metrics, see [TLS Metrics](#) section.

The following alerts are available for this feature:

- CERTIFICATE_EXPIRY
- BSF_CONNECTION_FAILURE

For more information about alerts, see [Configuring BSF Alerts](#) section.

Maintain

If you encounter alerts at system or application levels, see [Configuring BSF Alerts](#) section for resolution steps.

In case the alert still persists, perform the following:

1. **Collect the logs and Troubleshooting Scenarios:** For more information on how to collect logs and troubleshooting information, see *Oracle Communications Cloud Native Core, Binding Support Function Troubleshooting Guide*.
2. **Raise a service request:** See [My Oracle Support](#) for more information on how to raise a service request.

3.10.2 Support for TLS in Diameter Gateway

BSF uses Diameter Gateway to establish secured connections with consumer NFs and producer NFs, respectively. These communication protocols are encrypted using Transport Layer Security (TLS).

Enable

This feature is disabled by default. It can be enabled using TLS_ENABLED parameter using Helm configurations. For information about Helm configurations, see "Global Parameters" in

Oracle Communications Cloud Native Core, Binding Support Function Installation, Upgrade, and Fault Recovery Guide.

Configure

You can configure this feature using Helm. For information about Helm configurations, see "Global Parameters" in *Oracle Communications Cloud Native Core, Binding Support Function Installation, Upgrade, and Fault Recovery Guide*.

Peer level support for TLS versions can be configured using CNC Console. For more information, see [Peer Nodes](#) section.

Observe

The following metrics are available for this feature:

- `diam_conn_network`
- `diam_failed_conn_network`
- `diam_conn_network_responder`
- `dgw_tls_cert_expiration_seconds`

For more information about metrics, see [TLS Metrics](#) section.

The following alerts are available for this feature:

- `DGW_TLS_CONNECTION_FAILURE`
- `DIAM_GATEWAY_CERTIFICATE_EXPIRY_MINOR`
- `DIAM_GATEWAY_CERTIFICATE_EXPIRY_MAJOR`
- `DIAM_GATEWAY_CERTIFICATE_EXPIRY_CRITICAL`

For more information about alerts, see [Configuring BSF Alerts](#) section.

Following extensions are available for TLS in Diameter Gateway:

- `status_request (0x0005)`
- `status_request_v2 (0x0011)`
- `supported_groups (0x000A)`
- `ec_point_formats (0x000B)`
- `extended_master_secret (0x0017)`
- `session_ticket (0x0023)`
- `signature_algorithms (0x000D)`
- `signature_algorithms_cert (0x0032)`
- `supported_versions (0x002B)`
- `psk_key_exchange_modes (0x002D)`
- `key_share (0x0033)`
- `renegotiation_info (0xFF01)`

Maintain

If you encounter alerts at system or application levels, see [Configuring BSF Alerts](#) section for resolution steps.

In case the alert still persists, perform the following:

1. **Collect the logs and Troubleshooting Scenarios:** For more information on how to collect logs and troubleshooting information, see *Oracle Communications Cloud Native Core, Binding Support Function Troubleshooting Guide*.
2. **Raise a service request:** See [My Oracle Support](#) for more information on how to raise a service request.

3.10.3 TLS 1.3 Support for Kubernetes API Server

In a Kubernetes-based 5G Core deployment, BSF regularly interacts with the Kubernetes API server (Kube-API-Server) to retrieve the secrets and configmap information.

With the implementation of this feature, BSF supports TLSv1.3 in addition to TLSv1.2 for establishing secure communication with the Kubernetes API server.

Managing TLS 1.3 Support for Kubernetes API Server

Enable and Configure

This feature can be enabled or disabled using `kubeApiServerTlsVersion` flag. If the flag is enabled, user can choose TLS 1.3 or TLS 1.2 for communication between each service and Kubernetes API server. For information, see *Oracle Communications Cloud Native Core, Binding Support Function Installation, Upgrade, and Fault Recovery Guide*.

Observe

There are no alerts and metrics for this feature.

3.11 Enhancements to Error Response

BSF sends error responses to consumer NFs due to some exceptions, such as signaling, validations, and internal errors. These error responses have payloads containing the problem title, status, details, and cause of the error that are used to investigate the error. The details section is now enhanced with application error IDs.

The error handling module gives provision to configure the error response dynamically and the same is responded when BSF is producer of the call flow.

With the enhanced error response mechanism, BSF sends additional information such as server FQDN, micro-service ID, error category, and application error ID in the `detail` attribute of the `ProblemDetails`. This enhancement provides more information about the error and troubleshoot them.

Application error ID follows the below format.

[EC] [NF ID] [Microservice ID] [Category] [Error ID]

An error code dictionary will be provided to identify the cause and possible solution of the error. For more details of the error code dictionaries for BSF management service, Egress Gateway, Ingress Gateway, and NRF Client, see [Error Code Dictionary](#).

Managing Enhancements to Error Response

This section explains the procedure to enable and configure the feature.

Enable

By default, this feature is disabled. The operator can enable this feature through the CNC Console configurations.

Configure

You can configure error handling functionality under **Error Handling** on CNC Console for BSF. For information about how to configure for BSF Management Service in CNC Console, see [Error Configurations](#).

Observe

The following metrics have been added in BSF Management Service for this feature:

- error_handler_exec_total
- error_handler_in_total
- error_handler_out_total

For more information, see [BSF Management Service](#).

3.12 Validating Destination-Realm Attribute-Value Pair (AVP) Received in AAR-I Message

The destination-realm Attribute-Value Pair (AVP) received in the AAR-I message from an AF must be validated at the BSF Diameter Gateway before processing and forwarding the AAR-I message to a corresponding PCF instance.

Note

This validation applies exclusively to the AAR-I message within the BSF. For subsequent AAR-U and STR messages, it is assumed that the AF will send the correct Destination-Host and the destination-realm Attribute-Value Pair (AVP) values based on the AAA response received.

Figure 3-8 Diameter Gateway for BSF





Figure 3-9 destination-realm Validation Process

1. When Diameter Gateway receives an AAR-I message from an AF, it checks if the validation realm flag is enabled.
2. If the validation realm flag is enabled, it validates the destination-realm AVP received in the AAR-I message against the BSF Diameter Gateway realm.

If the validation is successful, it proceeds with the message processing and the Diameter Gateway sends a Binding Discovery message to BSF Management Service.

If the validation fails, the Diameter Gateway rejects the AAR-I message with error code 3003 (REALM_NOT_SERVED).

Note

The Diameter Gateway error code configuration for BSF is applicable and the default result code 3003 can be modified to any other error codes.

3. If the validation realm flag is disabled, the Diameter Gateway sends a Binding Discovery message to BSF Management Service without validating the destination-realm AVP.
4. After receiving a successful response to Binding Discovery request from BSF Management service, the Diameter Gateway checks if the validate binding realm flag is enabled.

If the validate binding realm flag is enabled, the Diameter Gateway validates the destination-realm against the pcfBinding realm.

If the validation is successful, the Diameter Gateway forwards the message for further processing.

If the validation fails, the Diameter Gateway rejects the AAR-I message with error code 3003 (REALM_NOT_SERVED).

If the validate binding flag is disabled, the Diameter Gateway forwards the message for further processing without validating the destination-realm against the pcfBinding info.

5. If the Binding Discovery request fails, the Diameter Gateway rejects the AAR-I message with error code 5012 (UNABLE_TO_COMPLY).
6. After the successful validation of the destination-realm against pcfBinding info, the Diameter Gateway checks if overwrite diam-realm flag is enabled. If this flag is enabled, it overwrites the destination-realm received in the AAR-I message with the pcfBinding info received in the binding discovery response.

Note

If Overwrite Realm configuration is enabled, the `diam_response_network` metric for AAA message will show the new realm that was overwritten as the 'reqDestRealm', and not the original realm that was sent in the AAR message.

7. If overwrite diam-realm flag is disabled, there is no modification made to the destination-realm.
8. The message is forwarded to PCF.

Note

Fake AVP validation will not be performed if routing table is configured. That is, if the above mentioned fake AVP parameters are configured and the routing tables are also configured, AVP validation will not be performed. The call will be processed as per the routing table configuration.

Managing Validation of destination-realm in AAR-I Message

The following Advanced Settings for Diameter Gateway are used to enable and validate the destination-realm in AAR-I message:

- DIAMETER.Enable.Validate.Realm
- DIAMETER.BSF.Enable.Validate.Binding.Realm
- DIAMETER.BSF.Enable.Overwrite.Realm

For more details on the above mentioned advanced settings keys, see [Settings](#).

Observability**Metrics**

ocbsf_diam_realm_validation_failed_total metric is used to count the number of failed destination-realm validation at Diameter Gateway for BSF. For more details, see [Diameter Gateway Metrics](#).

Alerts

The following alerts are used for Validating destination-realm Received in AAR-I Message feature:

- [DIAM_RESPONSE_REALM_VALIDATION_ERROR_MINOR](#)
- [DIAM_RESPONSE_REALM_VALIDATION_ERROR_MAJOR](#)
- [DIAM_RESPONSE_REALM_VALIDATION_ERROR_CRITICAL](#)

3.13 Support for Automated Certificate Lifecycle Management

Public Key Interface (PKI) is the set of elements such as public/private keys, certificate signing request, and certificates that are required to handle secure communications and transactions. BSF uses secure protocols for its communications, such as HTTPS and Secure Socket Layer (SSL) / Transport Layer Security (TLS) technologies to handle these secure communications. This is achieved with the use of Public and Private Keys, and the presence of trusted authorities, also known as Certificate Authorities (CA), which create and issue certificates. These certificates have a determined validity period. These certificates must be renewed before expiry. They can also be revoked when the CA or its keys are compromised. These certificates must be recreated when required.

This feature enables BSF to support automation of certificate lifecycle management in integration with Oracle Communications Cloud Native Core, Certificate Manager (OCCM).

OCCM provides the option to automatically create, renew, and delete certificates for a given CA, with the possibility to track previously created certificates and renew/delete them when required.


The certificate lifecycle management includes:

- Certificate Creation,
- Certificate Deletion,
- Certificate Monitoring (including the ones that were created using a different tool from OCCM),
- Certificate Renewal.

Note

OCCM does not support OAuth token generation and distribution (currently handled by NRF) used for SBI signalling.

Figure 3-10 BSF Integration with OCCM



There is no direct communication between OCCM and BSF. All the communications are handled using Kubernetes Secrets.

All the required certificates are configured using OCCM.

After OCCM creates these Kubernetes Secrets, or monitors the already existing ones, the Ingress and Egress Gateways monitor these Secrets and keep track of their current status:

- **VALID:** A Kubernetes Secret which holds a certificate that has not expired and it is properly signed
- **EXPIRED:** A Kubernetes Secret which holds a certificate that has met its expiration date (the value determined in its notAfter value)
- **MISSING:** A Kubernetes Secret which has its certificate missing, or any other essential file for the TLS/SSL bundle

- **CORRUPT:** A Kubernetes Secret which has its certificate corrupt, either invalid file, invalid signature, or invalid format

Managing the keys and certificates

Install Guide Considerations

- Upgrade: When BSF is deployed with OCCM, follow the specific upgrade sequence as mentioned in the *Oracle Communications, Cloud Native Core Solution Upgrade Guide*.
- Rollback: You can remove Kubernetes secrets if the current version of BSF does not use that secret by checking the `ocbsf_custom_values.yaml` file. Before deleting, please make sure that there is no plan to rollback to the BSF version which uses these secrets. Otherwise Rollback will fail. For more information on migrating the secrets from BSF to OCCM and removal of Kubernetes secrets from the yaml file, see *Upgrade Strategy in Oracle Communications Cloud Native Core, Binding Support Function Installation, Upgrade, and Fault Recovery Guide*.

Configure

To configure HTTPS in ingress-gateway, the following parameters must be configured in `custom-value.yaml` file in the ingress-gateway section:

- `ingress-gateway.enableIncomingHttps`
- `ingress-gateway.service.ssl.privateKey.k8SecretName`
- `ingress-gateway.service.ssl.privateKey.k8Namespace`
- `ingress-gateway.service.ssl.privateKey.rsa.fileName`
- `ingress-gateway.service.ssl.certificate.k8SecretName`
- `ingress-gateway.service.ssl.certificate.k8Namespace`
- `ingress-gateway.service.ssl.certificate.rsa.fileName`
- `ingress-gateway.service.ssl.caBundle.k8SecretName`
- `ingress-gateway.service.ssl.caBundle.k8Namespace`
- `ingress-gateway.service.ssl.caBundle.fileName`
- `ingress-gateway.service.ssl.keyStorePassword.k8SecretName`
- `ingress-gateway.service.ssl.keyStorePassword.k8Namespace`
- `ingress-gateway.service.ssl.keyStorePassword.fileName`
- `ingress-gateway.service.ssl.trustStorePassword.k8SecretName`
- `ingress-gateway.service.ssl.trustStorePassword.k8Namespace`
- `ingress-gateway.service.ssl.trustStorePassword.fileName`

For more information, see *Basic Configurations in Ingress Gateway* section in *Oracle Communications Cloud Native Core, Binding Support Function Installation, Upgrade, and Fault Recovery Guide*.

To configure HTTPS in egress-gateway, configure the following parameters under `egress-gateway` section in `custom-value.yaml` file:

- `egress-gateway.enableOutgoingHttps`
- `egress-gateway.egressGwCertReloadEnabled`
- `egress-gateway.egressGwCertReloadPath`
- `egress-gateway.service.ssl.privateKey.k8SecretName`

- egress-gateway.service.ssl.privateKey.k8NameSpace
- egress-gateway.service.ssl.privateKey.rsa.fileName
- egress-gateway.service.ssl.privateKey.ecdsa.fileName
- egress-gateway.service.ssl.certificate.k8SecretName
- egress-gateway.service.ssl.certificate.k8NameSpace
- egress-gateway.service.ssl.certificate.rsa.fileName
- egress-gateway.service.ssl.certificate.ecdsa.fileName
- egress-gateway.service.ssl.caBundle.k8SecretName
- egress-gateway.service.ssl.caBundle.k8NameSpace
- egress-gateway.service.ssl.caBundle.fileName
- egress-gateway.service.ssl.keyStorePassword.k8SecretName
- egress-gateway.service.ssl.keyStorePassword.k8NameSpace
- egress-gateway.service.ssl.keyStorePassword.fileName
- egress-gateway.service.ssl.trustStorePassword.k8SecretName
- egress-gateway.service.ssl.trustStorePassword.k8NameSpace
- egress-gateway.service.ssl.trustStorePassword.fileName

For more information, see *Basic Configurations in Egress Gateway* section in *Oracle Communications Cloud Native Core, Binding Support Function Installation, Upgrade, and Fault Recovery Guide*.

Observability

Monitoring the keys and certificates

BSF supports monitoring and automatic renewal of its' TLS certificates in integration with OCCM.

It is validated that the renewed certificate and key are picked up for any new TLS connections.

Also, the existing TLS connections using the previous key and certificate are gracefully brought down.

Clean up of the certificates are also handled through OCCM.

For information about enabling HTTPS, see *Configuring Secrets for Enabling HTTPS* in *Oracle Communications Cloud Native Core, Binding Support Function Installation, Upgrade, and Fault Recovery Guide*.

.

Metrics

The `oc_certificatemanagement_tls_certificate_info` metric is used to support automated certificate lifecycle management.

For more information, see [Metrics for Automated Certificate Lifecycle Management](#).

3.14 Support for cnDBTier APIs in CNC Console

With the implementation of this feature, cnDBTier APIs are integrated into the CNC Console, and BSF users can view specific cnDBTier functions, such as checking the cnDBTier version, status of cnDBTier clusters, and georeplication status in the CNC Console.

Note

This **cnDBTier** APIs can be accessed only through CNC Console.

The following cnDBTier APIs are read only and can be viewed on the CNC Console:

- **Backup List:** This API displays the details of stored backups, such as the ID and size of the backup.
- **cnDBTier Backup Status:** This API checks the cnDBTier backup status.
- **Database Statistics Report:** This API displays the number of available database.
- **Georeplication Status:**
 - **Real Time Overall Replication Status:** This API displays the overall replication status in multisite deployments. For example, in a four-site deployment, it provides the replication status between the following sites: site1-site2, site1-site3, site1-site4, site2-site3, site2-site4, and site2-site1. This is applicable for all other sites.
 - **Site Specific Real Time Replication Status:** This API displays the site-specific replication status.
- **Georeplication Status Across All Sites:** This API is used to retrieve the georeplication status across all the configured sites, offering a consolidated view of the overall replication status.
- **cnDBTier Health:** This API displays the health status of the following services:
 - **Replication Health Status:** This API displays the health status of the replication service. It checks the following:
 - * if the replication service is up or not
 - * if the replication service can connect to database or not
 - **Monitor Health Status:** This API displays the health status of the monitor service. It checks the following:
 - * if the monitor service is up or not
 - * if the service can connect to database or not
 - * if the metrics are fetched or not (the metrics are fetched when the service is up and vice versa)

NDB Health Status: This API displays the health status of the NDB service pods like (data pods, sql pods, app-my-sql pods, mgmt pods). It checks the following:

- if the pod is connected to PVC or not
- if the pods status is up or not

Note

PVC Health Status attribute is set to NA when some of the database pods are not connected to the PVC.

- Backup Manager Health Status: This API displays the health status of the backup manager service. It checks the following:
 - * if the backup manager service is up or not
 - * if the service can connect to database or not

- HeartBeat Status: This API displays the connectivity status between the local site and the remote site to which BSF is connected.
- IP Family Switch Over: This API is used to configure or update the preferred IP family (IPv4 or IPv6) used for communication between the replication channels.
- Georeplication Recovery: This API provides the recovery status of georeplication for the cnDBTier cluster:
 - Update Cluster As Failed: This API is used to mark a disrupted cluster as failed.
 - Start Georeplication Recovery: This API is used to start the georeplication recovery process.
 - Georeplication Recovery Status: This API is used to monitor the recovery status of georeplication for both FAILED and ACTIVE cnDBTier sites.
- Local Cluster Status: This API displays the status of the local cluster.
- On Demand Backup: This API provides options to initiate as well as the display the status of the on-demand backup. It also displays the status of initiated on-demand backups.
- Version: This API displays the cnDBTier version.
- Gracefully Start Georeplication: This API is used to initiate the georeplication between a primary site and one or more remote sites.
- Gracefully Stop Georeplication: This API is used to stop georeplication between a primary site and specified remote sites.

Managing cnDBTier Functionalities at CNC Console

Enable

This feature is enabled automatically when cnDBTier is configured as an instance during the CNC Console deployment. For more information about integrating cnDBTier APIs in CNC Console, see *Oracle Communications Cloud Native Core, cnDBTier User Guide*.

Configure

You can view cnDBTier APIs in CNC Console in the [Viewing cnDBTier APIs in CNC Console](#) section.

Maintain

If you encounter alerts at the system level, see the [BSF Alerts](#) section for resolution steps.

In case the alerts persist, perform the following tasks:

1. **Collect the logs:** For information about how to collect logs, see *Oracle Communications Cloud Native Core, Binding Support Function Troubleshooting Guide*.

2. **Raise a service request:** For information about how to raise a service request, see [My Oracle Support](#).

3.15 Diameter Session Retry

BSF Diameter gateway sends Authorization-Authentication Request (AAR) and Session Termination Request (STR) messages to the PCF Diameter gateway, and Re-Authorization Request (RAR), and Abort Session Request (ASR) messages to P-CSCF (AF). These requests are acknowledged by corresponding response messages such as Authorization-Authentication Answer (AAA), Session Termination Answer (STA), Re-Authorization Answer (RAA), and Abort Session Answer (ASA), indicating either success or failure.

If BSF Diameter gateway receives messages such as Authorization-Authentication Answer (AAA), Re-Authorization Answer (RAA), (Session Termination Answer) STA or Abort Session Answer (ASA) with errors like 5012 (DIAMETER_UNABLE_TO_COMPLY) or any other error code or a session timeout then the CNC console has configurations that can be used to resend this failed message to a different/alternate PCF Diameter gateway, or to resend the message to the same destination through an alternate routing path (for example, via a different DRA peer). This enables two retry methods: retry to an alternate destination peer or retry to the same destination using an alternate intermediate peer.

The BSF Diameter gateway, on receiving a failed message from the Diameter Routing Agent (DRA), captures the error context and details. These error details are passed to the error handling framework implemented within the Diameter gateway. The framework determines the appropriate action such as retrying the message on an alternate route or peer based on configurations defined by the user in the CNC Console.

The Diameter message retry feature supports both direct connectivity and routing through Diameter Routing Agent (DRA) peers. When BSF is directly connected to PCF / P-CSCF, retries can be performed to alternate destination peers. When routing is performed through DRA peers (using FQDN), retries can be performed to the same destination through an alternate DRA peer.

If the Error Handler was previously enabled for AAR, it will also be enabled for RAR, ASR and STR after an upgrade. It can be disabled through the configuration screen by going to the specific message. Diameter message retry is enabled through the Error Mapping Framework feature in BSF for Rx interface messages, including AAR, STR, RAR, and ASR. This framework resolves application errors and takes necessary actions based on the error context. It attempts to find alternate solutions using the configurations provided in the CNC Console. If the error is resolved, a success response is returned to the caller; otherwise, the framework retries the request based on the configured maximum number of attempts or stops processing by forwarding the last known error.

The operator should have configured a host and realm in the diameter routing table to retry sending the failed diameter messages. Diameter gateway finds the alternate peer from the routing table. If the diameter routing table is not configured, then there is no retry behavior.

When selecting an alternate peer for a retry, the Diameter Gateway evaluates matching Diameter Routing Table entries based on the configured **Priority** value. The route entry with the lowest numerical priority value is considered first.

By default, the diameter message retry behavior is disabled for Rx interface. The operator can enable this feature through the CNC Console configurations.

Note

Default number of retry attempts is 1. The number of retries can be any positive integer (≥ 1).

BSF retries re-sending failed Diameter messages (Rx AAR, Rx STR, Rx RAR, and Rx ASR) for configured Diameter result codes, except 2xxx result codes. The 2xxx result code series represents successful responses and is not supported for retry configuration. Retry is also supported for response timeout scenarios.

Retry Attempts

The user configures the number of retries to be performed for Rx AAR, Rx STR, Rx RAR and Rx ASR diameter messages in the CNC Console. The retry attempt happens only when the alternate peers are available. If alternate peers are not available, then there is no retry attempt made. The value for number of retries can be any positive integer (≥ 1). There is no explicit maximum limit. It can go up to the value of Integer.MAX_VALUE. If a value less than 1 is configured, the system treats it as 1.

The number of retries is set through the advance settings configurations, using the following advanced settings keys:

- `DIAMETER.ErrorHandler.MaxRetryCount.Rx.AAR`
- `DIAMETER.ErrorHandler.MaxRetryCount.Rx.STR`
- `DIAMETER.ErrorHandler.MaxRetryCount.Rx.RAR`
- `DIAMETER.ErrorHandler.MaxRetryCount.Rx.ASR`

If the above advanced settings configurations are not provided in the CM-GUI, then the default value for the number of retries is set to 1.

Peer Cycle Back Retry

In case the configured number of retry count is more than the total available alternate peers, the user can configure to cycle back the alternate peers. The user configures this in the CNC Console by setting the `retry peer cycle back` field to true. This field value can be either true or false.

This peer cycle back retry configuration is set through the advanced settings configurations using the following advanced setting keys:

- `DIAMETER.ErrorHandler.CycleBackRetry.Rx.AAR`
- `DIAMETER.ErrorHandler.CycleBackRetry.Rx.STR`
- `DIAMETER.ErrorHandler.CycleBackRetry.Rx.RAR`
- `DIAMETER.ErrorHandler.CycleBackRetry.Rx.ASR`

If advanced settings configuration are not supported in the CNC Console, then the default peer cycle back retry is false.

For Example: Number of configured retry = 2 and only 2 PCF (PCF1, PCF2) Diameter gateway are configured as alternate peers in BSF.

If Rx AAR message was sent BSF - PCF1 Diameter gateway and the response has failed error code such as timeout/3002/3004.

Then the first retry uses BSF - PCF2-Diam-Gateway and the response has failed error code such as timeout/3002/3004.

Then the second retry uses BSF - PCF1-Diam-Gateway and thus uses the peer cycle back retry mechanism.

Similarly, the same peer cycle back retry behavior is applicable for Rx STR, Rx RAR, and Rx ASR messages when retries are configured for those message types.

Error Originator Peer

The Error Originator Peer indicates as to where the failed message error occurred/originated when sending or retrying the Diameter messages. The user can customize the error origination peer by using Error Response Originator field in the CNC Console and the customizing options are based on:

- The error received from an intermediate peer (INTERMEDIATE PEER).
- The error received from the destination peer, which is not an intermediate peer (DESTINATION PEER).
- The error received from any peer (ANY).


By default the error originator peer option is any peer.

In response timeout cases the error originator option is not available since the origination of error cannot be found.

Call Flows in Diameter Session Retry

Call Flow of Rx AAR Existing Success or Failed Case - When Binding is Found


The call flow describes the scenario where BSF Diameter Gateway sends AAR request toward PCF1 Diameter Gateway on receiving the binding response from BSF Management service. On receiving the AAR request the PCF1 Diameter Gateway may send a successful or a failed AAA response toward BSF Diameter Gateway. The failed response is not handled by the BSF Diameter Gateway.

Figure 3-11 Rx AAR Existing Success/Failed Case - When Binding is Found

1. Application Function (AF) sends AAR request to BSF Diameter Gateway.
2. BSF Diameter Gateway sends binding request to BSF Management Service.
3. BSF Management service sends back binding response.
4. BSF Diameter Gateway sends AAR request to PCF1 Diameter Gateway.
5. PCF1 Diameter Gateway may send either a successful or failed AAA response and BSF Diameter Gateway do not handle the failed response.

Call Flow of Rx AAR No Retry Case - When Binding is Not Found


The call flow describes the scenario where BSF Diameter Gateway do not receive the binding response from BSF Management service.

Figure 3-12 Rx AAR No Retry Case - When Binding is Not Found

1. Application Function (AF) sends AAR request to BSF Diameter Gateway.
2. BSF Diameter Gateway sends binding request to BSF Management Service.
3. BSF Management service sends "Binding not found" response.
4. BSF Diameter Gateway responds to AF with error code 5065 (IP-CAN_SESSION_NOT_AVAILABLE).

Call Flow of Rx AAR Retry Case - When No Binding Request


The call flow describes the scenario where BSF Diameter Gateway sends AAR request toward PCF1 Diameter Gateway on not receiving the binding response from BSF Management service. On receiving the AAR request the PCF1 Diameter Gateway may send a successful or a failed AAA response toward BSF Diameter Gateway. The failed response is handled by the BSF Diameter Gateway.

Figure 3-13 Rx AAR Retry Case - When No Binding Request

1. Application Function (AF) sends AAR request to BSF Diameter Gateway.
2. BSF Diameter Gateway sends AAR request to PCF1 Diameter Gateway.
3. PCF1 Diameter Gateway sends a failed AAA response to BSF Diameter Gateway.
4. BSF Diameter Gateway retries the AAR message to PCF2 Diameter Gateway.
5. PCF2 Diameter Gateway sends a successful AAA response to BSF Diameter Gateway.
6. BSF Diameter Gateway sends this AAA message to AF.

Call Flow of RX-AAR / RX-STR Retry - Success Case


This call flow describes the Diameter AAR/STR retry and failover handling mechanism between AF, BSF, DRA, and PCF components. When the initial AAR or STR request sent through DRA1 fails, the Diameter Gateway automatically retries the request through an alternate DRA instance, DRA2. The retry request is successfully forwarded to the PCF, and the successful AAA or STA response is returned back to the AF.

Figure 3-14 RX-AAR / RX-STR Retry - Success Case

1. AF sends an AAR or STR request to the Diameter Gateway.
2. The Diameter Gateway forwards the AAR or STR request to DRA1.
3. DRA1 fails to process the AAR or STR request.
4. DRA1 returns a failed AAA or STA response to the Diameter Gateway.
5. The Diameter Gateway retries the AAR or STR request through DRA2.
6. DRA2 forwards the AAR or STR request to the Diameter Gateway on the PCF side.
7. The PCF-side Diameter Gateway processes the request successfully and returns a successful AAA or STA response to DRA2.
8. DRA2 sends the successful AAA or STA response to the Diameter Gateway.
9. The Diameter Gateway returns the successful AAA or STA response to the AF.

Call Flow of RX-AAR / RX-STR Retry - Inter-Pod Routing Case

This call flow describes how the Diameter Gateway handles AAR or STR request failures when direct connectivity to a target DRA is unavailable. The request is first rerouted between Diameter Gateway pods using inter-pod routing. After the initial request sent through DRA1 fails, the Diameter Gateway retries the request through DRA2. The retry succeeds, and the successful AAA or STA response is returned back to the AF.

Figure 3-15 RX-AAR / RX-STR Retry - Inter-Pod Routing Case


1. AF sends an AAR or STR request to Diameter Gateway Pod-1.
2. Diameter Gateway Pod-1 detects that direct connectivity is unavailable.
3. Diameter Gateway Pod-1 forwards the request to Diameter Gateway Pod-2 using inter-pod routing.
4. Diameter Gateway Pod-2 forwards the AAR or STR request to DRA1.
5. DRA1 fails to process the AAR or STR request.
6. DRA1 returns a failed AAA or STA response to Diameter Gateway Pod-2.
7. Diameter Gateway Pod-2 forwards the failed AAA or STA response to Diameter Gateway Pod-1.
8. Diameter Gateway Pod-1 retries the AAR or STR request through DRA2.
9. DRA2 forwards the AAR or STR request to the Diameter Gateway on the PCF side.
10. The PCF-side Diameter Gateway processes the request successfully and returns a successful AAA or STA response to DRA2.
11. DRA2 forwards the successful AAA or STA response to Diameter Gateway Pod-1.
12. Diameter Gateway Pod-1 returns the successful AAA or STA response to the AF.

Call Flow of RX-RAR / RX-ASR Existing - Failed Case

This call flow describes the handling of failed RAR or ASR requests between the PCF, BSF, DRA, and AF components. The PCF-side Diameter Gateway initiates a RAR or ASR request toward the AF through the BSF-side Diameter Gateway and DRA1. When the request

processing fails, the failure response is sent back through the signaling path to the originating Diameter Gateway.


Figure 3-16 RX-RAR / RX-ASR Existing - Failed Case



1. The PCF-side Diameter Gateway sends a RAR or ASR request to the BSF-side Diameter Gateway.
2. The BSF-side Diameter Gateway forwards the RAR or ASR request to DRA1.
3. DRA1 fails to process the RAR or ASR request.
4. DRA1 returns a failed RAA or ASA response to the BSF-side Diameter Gateway.
5. The BSF-side Diameter Gateway forwards the failed RAA or ASA response to the PCF-side Diameter Gateway.

Call Flow of RX-RAR / RX-ASR Retry - Success Case


This call flow describes the retry and failover handling mechanism for Diameter RAR or ASR requests between the PCF, BSF, DRA, and AF components. When the initial RAR or ASR request sent through DRA1 fails, the BSF-side Diameter Gateway retries the request through DRA2. The retry request successfully reaches the AF, and the successful RAA or ASA response is sent back to the PCF-side Diameter Gateway.

Figure 3-17 RX-RAR / RX-ASR Retry - Success Case

1. The PCF-side Diameter Gateway sends a RAR or ASR request to the BSF-side Diameter Gateway.
2. The BSF-side Diameter Gateway forwards the RAR or ASR request to DRA1.
3. DRA1 fails to process the RAR or ASR request.
4. DRA1 returns a failed RAA or ASA response to the BSF-side Diameter Gateway.
5. The BSF-side Diameter Gateway retries the RAR or ASR request through DRA2.
6. DRA2 forwards the RAR or ASR request to the AF.
7. The AF processes the request successfully and returns a successful RAA or ASA response to DRA2.
8. DRA2 forwards the successful RAA or ASA response to the BSF-side Diameter Gateway.
9. The BSF-side Diameter Gateway forwards the successful RAA or ASA response to the PCF-side Diameter Gateway.

Call Flow of RX-RAR / RX-ASR Retry - Inter-Pod Routing Case


This call flow describes the retry and failover handling mechanism for Diameter RAR or ASR requests when the initial request sent through DRA1 fails and the selected alternate DRA peer, DRA2, is not directly reachable from Diameter Gateway Pod-1. In this case, Diameter Gateway Pod-1 uses inter-pod routing to forward the retry request to Diameter Gateway Pod-2, which can reach DRA2. The retry request is then sent through DRA2, successfully reaches the AF, and the successful RAA or ASA response is sent back to the PCF-side Diameter Gateway.

Figure 3-18 RX-RAR / RX-ASR Retry - Inter-Pod Routing Case

1. The PCF-side Diameter Gateway sends a RAR or ASR request to Diameter Gateway Pod-1.
2. Diameter Gateway Pod-1 forwards the RAR or ASR request to DRA1.
3. DRA1 fails to process the RAR or ASR request.
4. DRA1 returns a failed RAA or ASA response to Diameter Gateway Pod-1.
5. After the failed RAA or ASA response from DRA1, Diameter Gateway Pod-1 selects DRA2 as the alternate DRA peer. Since Diameter Gateway Pod-1 cannot directly reach DRA2, it initiates inter-pod routing to Diameter Gateway Pod-2.
6. Diameter Gateway Pod-2 retries the RAR or ASR request through DRA2.
7. DRA2 forwards the RAR or ASR request to the AF.
8. The AF processes the request successfully and returns a successful RAA or ASA response to DRA2.
9. DRA2 forwards the successful RAA or ASA response to Diameter Gateway Pod-2.
10. Diameter Gateway Pod-2 forwards the successful RAA or ASA response to Diameter Gateway Pod-1.
11. Diameter Gateway Pod-1 forwards the successful RAA or ASA response to the PCF-side Diameter Gateway.

Call Flow of RX-AAR / RX-STR Retry without FQDN Peers - BSF Directly Connected to PCFs

This call flow describes the retry handling mechanism for Diameter AAR or STR requests between AF, BSF, and PCF components. The initial request sent to PCF1 fails and returns an unsuccessful AAA or STA response. The BSF-side Diameter Gateway retries the request to an alternate PCF instance, PCF2. The retry succeeds, and the successful AAA or STA response is returned to the AF.

Figure 3-19 RX-AAR / RX-STR Retry without FQDN Peers - BSF Directly Connected to PCFs


1. AF sends an AAR or STR request to the BSF-side Diameter Gateway.
2. The BSF-side Diameter Gateway forwards the AAR or STR request to PCF1.
3. PCF1 returns a failed AAA or STA response to the BSF-side Diameter Gateway.
4. The BSF-side Diameter Gateway retries the AAR or STR request to PCF2.
5. PCF2 processes the request successfully and returns a successful AAA or STA response to the BSF-side Diameter Gateway.
6. The BSF-side Diameter Gateway forwards the successful AAA or STA response to the AF.

Call Flow of RX-RAR / RX-ASR Retry without FQDN Peers - BSF Directly Connected to AFs

This call flow describes the retry handling mechanism for Diameter RAR or ASR requests between PCF, BSF, and AF components. The PCF-side Diameter Gateway initially sends the request to AF1, but AF1 returns a failed RAA or ASA response. The BSF-side Diameter

Gateway retries the request to an alternate AF instance, AF2. The retry succeeds, and the successful RAA or ASA response is returned back to the PCF-side Diameter Gateway.

Figure 3-20 RX-RAR / RX-ASR Retry without FQDN Peers - BSF Directly Connected to AFs



1. The PCF-side Diameter Gateway sends a RAR or ASR request to the BSF-side Diameter Gateway.
2. The BSF-side Diameter Gateway forwards the RAR or ASR request to AF1.
3. AF1 returns a failed RAA or ASA response to the BSF-side Diameter Gateway.
4. The BSF-side Diameter Gateway retries the RAR or ASR request to AF2.
5. AF2 processes the request successfully and returns a successful RAA or ASA response to the BSF-side Diameter Gateway.
6. The BSF-side Diameter Gateway forwards the successful RAA or ASA response to the PCF-side Diameter Gateway.

Default Error Handling Configuration

BSF provides the default error handling configuration to retry on all error codes (except diameter result code 2xxx) and timeout for Rx AAR, Rx-STR, Rx-RAR and Rx-ASR failed diameter messages. When the diameter message retry feature is enabled on Rx interface, these default error handling configurations get applied by default. The user has an option to enable/disable these default configurations through the CNC Console edit configurations.

For all default error handling configurations, the value for retry attempt is 1 and the default action is ONE_RETRY_TO_ALTERNATE_PEER. The default value for peer cycle back retry is false.

Upgrade Impact

1. If the Error Handler Configuration was enabled in a previous version, then the default error handling configurations for Rx RAR, Rx STR, and Rx ASR messages are automatically enabled as part of the current version, since the feature is enabled. If the user wants to disable retry for Rx RAR, Rx STR, and Rx ASR messages, the user can do so manually by setting the Error Rule Enabled field to false for the respective error rule.
2. The system restricts the configuration of 2xxx series error codes, as these represent success status codes. The error code must be a 4-digit combination of numbers and wildcard (*) characters, and 2xxx series codes are not allowed. If a 2xxx series error code was configured for AAR messages in a previous version, it is retained after upgrade. However, if such a configuration is exported and then imported, the import operation results in a PARTIAL_SUCCESS status, as 2xxx series codes are not supported during import. No other validation is applied to the status codes.

Managing Diameter Session Retry

This section explains the procedure to enable and configure the feature.

Enable

By default, Diameter Message Retry behavior is disabled for Rx interface and operator can enable this feature through the CNC Console configurations.

Configure Using CNC Console

Perform the feature configurations in CNC Console as described in [Error Configurations](#) section.

To enable over-writing of destination host on retry message, *DIAMETER.ErrorHandler.Enable.UpdateDestinationHost* key must be set to true in the Advanced Settings. For more information, see [Settings](#).

Configure Using REST API

Perform the export/import error configurations as described in "Error Configurations" section in *Oracle Communications Cloud Native Core, Binding Support Function REST Specification Document*.

Observe

Observability

Metrics

Following metrics were updated with retry and retryReason dimensions in the [Diameter Gateway Metrics](#) section:

- ocbsf_diam_request_network_total
- ocbsf_diam_request_inter_total

Alerts

Following alerts are used by this feature:

- [AAA_RX_FAIL_COUNT_EXCEEDS_CRITICAL_THRESHOLD](#)

- [AAA_RX_FAIL_COUNT_EXCEEDS_MAJOR_THRESHOLD](#)
- [AAA_RX_FAIL_COUNT_EXCEEDS_MINOR_THRESHOLD](#)
- [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](#)

Maintain

If you encounter alerts at system or application levels, see [BSF Alerts](#) section for resolution steps.

In case the alerts still persist, perform the following:

- Collect the logs: For more information on how to collect logs, see *Oracle Communications Cloud Native Core, Binding Support Function Troubleshooting Guide*.
- Raise a service request: See [My Oracle Support](#) for more information on how to raise a service request.

3.16 Support for BSF Status on NRF on CNC Console

CNC Console for BSF shows health status information of BSF and other producer NF instances related to BSF. It provides a consolidated status of the BSF instances registered with NRF.

BSF CNC Console has a new **BSF NRF Status** page, added under **Status and Query** page. This page provides the consolidated status of BSF instances registered with NRF. The user can also see the health status of primary and secondary NRF instances, or the alternate NRF.

Managing Support for BSF status on NRF on CNC Console

Enable

BSF Status on NRF feature uses NRF Client service. Hence ensure that the NRF Client service is enabled by setting the value `global.nrfClientNfManagementEnable` as `true`.

For more information on NRF Client configurations, see *Configuring NRF Client* section in *Oracle Communications Cloud Native Core, Binding Support Function Installation, Upgrade, and Fault Recovery Guide*.

Configure Using CNC Console

To view the health status using CNC Console, see [BSF NRF Status](#).

Configure Using REST API

Perform the feature configurations as described in "NRF Status" and "NRF Client" sections in *Oracle Communications Cloud Native Core, Binding Support REST Specification Guide*

Observability

Metrics:

Following metrics were updated in the [NRF Client Metrics](#) section.

- nrfclient_perf_info_nf_profile_load
- nrfclient_current_nf_status
- nrfclient_nf_status_with_nrf
- nrfclient_nrf_operative_status
- nrfclient_nrf_status_total
- nrfclient_nrf_successive_healthy_count
- nrfclient_nrf_successive_unhealthy_count
- nrfclient_on_demand_conn_in_request_total
- nrfclient_on_demand_conn_out_response_total
- nrfclient_on_demand_processing_latency_ms
- ocpm_nrf_tracing_request_timeout_total
- nrfclient_nw_conn_out_request_total
- nrfclient_nw_conn_in_response_total
- nrfclient_nw_conn_in_notify_request_total
- nrfclient_nw_conn_out_notify_response_total
- nrfclient_network_message_processing_latency

Maintain

If you encounter alerts at system or application levels, see [BSF Alerts](#) section for resolution steps.

In case the alerts still persist, perform the following:

- Collect the logs: For more information on how to collect logs, see *Oracle Communications Cloud Native Core, Binding Support Function Troubleshooting Guide*.
- Raise a service request: See [My Oracle Support](#) for more information on how to raise a service request.

3.17 Network Policies

Network Policies are an application-centric construct that allows you to specify how a pod communicates with various network entities. It creates pod-level rules to control communication between the cluster's pods and services, and to determine which pods and services can access one another inside a cluster.

Previously, the pods under BSF deployment could be contacted by any other pods in the Kubernetes cluster without any restrictions. Now, Network Policies provide namespace-level isolation, which allows secured communications to and from BSF with rules defined in respective Network Policies. The Network Policies enforce access restrictions for all the applicable data flows except communication from Kubernetes node to pod for invoking container probe. For example, BSF internal microservices cannot be contacted directly by any other pods.

The following table lists the different access policies to be used by BSF traffic flows.

Note

This list is not exhaustive but tries to represent all the traffic flows supported by BSF.

Microservice	Direction	Client/Server	Port	Access Policy
Configuration Svc	Egress	- Database - K8s API server for K8s secret	3306, K8s API Server Port	K8s Network Policies
Configuration Svc	Egress	Jaeger Agent	6831	K8s Network Policies
Configuration Svc	Ingress	- Console - Egress Gateway for configuration - Ingress Gateway for configuration - Perf-info for configuration - App-info for configuration - ATS - ARS - NrfClient	8081	K8s Network Policies
Configuration Svc	Ingress	- Prometheus - Liveness - Readiness	9000	K8s Network Policies
Ingress Gateway	Egress	Jaeger Agent	6831	K8s Network Policies
Ingress Gateway	Egress	- Database - K8s API Server for K8s Secret	3306, K8s API Server Port	K8s Network Policies
Ingress Gateway	Egress	Coherence	8000, 7	K8s Network Policies
Ingress Gateway	Ingress	Perf Info	8080	K8s Network Policies
Ingress Gateway	Ingress	SBI Peer	80, 443	3GPP-defined Access Policies
Ingress Gateway	Ingress	Coherence	8000, 8095, 8096	K8s Network Policies
Ingress Gateway	Ingress	- Prometheus - Liveness - Readiness	9000	K8s Network Policies
Egress Gateway	Egress	Jaeger Agent	6831	No Access Policy due to SBI Egress*
Egress Gateway	Egress	- Database - K8s API Server for K8s Secret	3306, K8s API Server Port	No Access Policy due to SBI Egress*

Microservice	Direction	Client/Server	Port	Access Policy
Egress Gateway	Egress	Coherence	8000, 8095, 8096	No Access Policy due to SBI Egress*
Egress Gateway	Egress	SBI Peer	Decided at run-time	3GPP-defined Access Policies
Egress Gateway	Egress	ARS	ARS Port	K8s Network Policies
Egress Gateway	Ingress	Registration	8080	K8s Network Policies
Egress Gateway	Ingress	Egress Gateway for coherence	8000	K8s Network Policies
Egress Gateway	Ingress	- Prometheus - Liveness - Readiness	9000	K8s Network Policies
Egress Gateway	Ingress	Coherence	8000, 8095, 8096	K8s Network Policies
Audit	Ingress	- Prometheus - Liveness - Readiness	9000	K8s Network Policies
Audit	Egress	- Database - K8s API Server for K8s Secret	3306, K8s API Server Port	K8s Network Policies
App Info	Ingress	- Registration - Subscription - Auditor	5906	K8s Network Policies
App Info	Ingress	- Prometheus - Liveness - Readiness	9000	K8s Network Policies
diam-gateway	Egress	Jaeger Agent	6831	K8s Network Policies
diam-gateway	Ingress	- Prometheus - Liveness - Readiness	9000	K8s Network Policies
diam-gateway	Egress	- Database - K8s API Server for K8s Secret	3306, K8s API Server Port	K8s Network Policies
diam-gateway	Ingress	Peer	3868	K8s Network Policies
Bsf-Management	Ingress	- Prometheus - Liveness - Readiness	9000	K8s Network Policies
Bsf-Management	Egress	Jaeger Agent	6831	K8s Network Policies
Bsf-Management	Egress	- Database - K8s API Server for K8s Secret	3306, K8s API Server Port	K8s Network Policies

Microservice	Direction	Client/Server	Port	Access Policy
NRF-Client	Ingress	- Prometheus - Liveness - Readiness	9000	K8s Network Policies
NRF-Client	Egress	Jaeger Agent	6831	K8s Network Policies
NRF-Client	Egress	- Database - K8s API Server for K8s Secret	3306, K8s API Server Port	K8s Network Policies
Query-Svc	Ingress	- Prometheus - Liveness - Readiness	9000	K8s Network Policies
Query-Svc	Egress	Jaeger Agent	6831	K8s Network Policies
Query-Svc	Egress	- Database - K8s API Server for K8s Secret	3306, K8s API Server Port	K8s Network Policies

As an assumption when deploying BSF, the following labels are set by default:

Table 3-17 Default Labels

Pod	Label
All BSF Pods	app.kubernetes.io/part-of: ocbsf
Ingress Gateway Pod	app.kubernetes.io/name: ocbsf-ingress-gateway
Egress Gateway Pod	app.kubernetes.io/name: ocbsf-egress-gateway
Diam-gateway	app.kubernetes.io/name: diam-gateway
CM-Service	app.kubernetes.io/name: cm-service

BSF Security Policies:

- **deny-ingress-all:** To block all ingress traffic of pods presents in a BSF deployment.
- **allow-ingress-sbi:** To allow traffic on the Ingress Gateway Pods on container ports 8000 and 9443 to allow sbi traffic.
- **allow-diam-gateway:** To allow traffic on the Diameter-Gateway on port 3868.
- **allow-ingress-prometheus:** To allow the traffic flow from Prometheus service to the BSF with default ports (These ports can be changed by the customer).
- **allow-ingress-from-bsf-pods:** To allow ingress communication between the different microservices of the BSF.
- **allow-ingress-from-console:** To allow ingress communication between CNCC-Core and CM-Service on port 8081.
- **deny-egress-all-except-egw:** To block all egress traffic of pods present in a BSF deployment, except for Egress-Gateway and Diameter-Gateway.
- **allow-egress-database:** To allow the traffic flow from BSF to db sql port and db monitoring port with default ports (These ports can be changed by the customer).

- **allow-egress-k8s-api:** To allow the traffic flow from BSF to Kubernetes API server port (These ports can be changed by the customer).
- **allow-egress-jaeger:** To allow the traffic flow from BSF to Jaeger agent port and DNS service with default ports (These ports can be changed by the customer).
- **allow-egress-traffic-and-dns:** To allow the traffic flow from BSF to k8s DNS service with default ports (These ports can be changed by the customer).
- **allow-egress-to-bsf-pods:** To allow egress communication between the different microservices of BSF.

① Note

The default Network Policies to be applied for BSF are the recommended even though they are not very granular but they keep operational overhead to the minimum and still achieve access control security.

If a Network Policy is installed to restrict Prometheus escaping the metrics from the PODs, then a restart of the PODs is required. If the NP is installed before the PODs are up, then there is no need to restart the PODs.

Managing Network Policies

Enable

To use this feature, Network Policies need to be applied to the namespace wherein BSF is applied.

Configure

You can configure this feature using Helm. For information about configuring network policy for BSF deployment, see *Configuring Network Policy* section in *Oracle Communications Cloud Native Core, Binding Support Function Installation, Upgrade, and Fault Recovery Guide*.

Observe


There are no specific metrics and alerts required for the Network Policies feature.

3.18 Monitoring the Availability of SCP using HTTP2 OPTIONS

BSF determines the availability and reachability status of all SCPs irrespective of the configuration types.

This feature is an enhancement to the existing SBI routing functionality. Egress Gateway microservice interacts with SCP on their health API endpoints using HTTP2 OPTIONS method. It monitors the health of configured SCP peers to ensure that the traffic is routed directly to the healthy peers. This enhancement avoids routing or rerouting towards unhealthy peers, thus minimizing the latency time.

Egress Gateway microservice maintains the health status of all available and unavailable SCPs. It maintains the latest health of SCPs by periodically monitoring and uses this data to route egress traffic to the most preferred healthy SCP.

Figure 3-21 New SCP Selection Mechanism

Once `peerconfiguration`, `peersetconfiguration`, `routesconfiguration`, and `peermonitoringconfiguration` parameters are configured at Egress Gateway microservice, and all SCPs (after Alternate Route Service (ARS) resolution, if any vFQDN is configured) are marked initially as healthy. The peers attached to the associated peerset are scheduled to run health API checks and update the health status continuously.

During the installation, the value of the parameter `peermonitoringconfiguration` is set to false by default. Since, this feature is an add-on to the existing SBI Routing feature and will be activated if the `sbirouteconfig` feature is enabled. To enable this feature, perform the following:

- configure `peerconfiguration` with `healthApiPath`
- configure `peersetconfiguration`
- configure `sbiroutingerroractionsets`
- configure `sbiroutingerroractionsets`

- configure routesconfiguration
- enable peermonitoring

If SBI Routing feature is enabled before upgrading, the healthApi in peerconfiguration should be attached manually to existing configured peers. If the operator tries to enable peermonitoringconfiguration and the targeted peers do not have the healthApiPath then an appropriate error response is sent.

Managing Monitoring the Availability of SCP Using SCP Health APIs

This section explains the procedure to enable and configure the feature.

Configure

You can configure the Monitoring the Availability of SCP using the REST API.

Configure Using REST API: Perform the following feature configurations as described in *Oracle Communications Cloud Native Core, Binding Support Function REST Specification Document*:

- create or update peer Peer Configuration with health status endpoint details.
- create or update the peerset peersetconfiguration to assign these peers
- enable the feature using the below peermonitoring configuration peermonitoringconfiguration.

Note

Health Monitoring of the peer will start only after the feature is enabled and the corresponding peerset is used in sbirouteconfig.

Observe

Following metrics are added in the Metrics in [Egress Gateway Metrics for SCP](#) section:

- oc_egressgateway_peer_health_status
- oc_egressgateway_peer_health_ping_request_total
- oc_egressgateway_peer_health_ping_response_total
- oc_egressgateway_peer_health_status_transitions_total
- oc_egressgateway_peer_count
- oc_egressgateway_peer_available_count

Alert

Following alerts are added in the Alert section:

- [SCP_PEER_UNAVAILABLE](#)
- [SCP_PEER_UNAVAILABLE](#)

3.19 Supports 3gpp-Sbi-Correlation-Info Header

The 3gpp-Sbi-Correlation-Info header may be used to contain correlation information such as UE identity, that may be used by an operator in various offline network management,

performance analysis and troubleshooting tools/applications to identify messages (requests, responses, subscriptions, notifications) related to a particular subscriber.

By supporting this feature, BSF as a service consumer or as a service producer generates, forwards and sends the UE identity in 3gpp-Sbi-Correlation-Info header, to identify the UE related to the HTTP request or response.

BSF provides a global configurations page on CNC Console GUI to enable or disable the correlation-info header feature. On enabling,

- BSF receives the correlation-info header and forwards them to the producer NFs.
- BSF does not receive the correlation-info header, then BSF generates and forwards them to the producer NFs.

In BSF, generation of new correlation-info header is managed by the **NF Communication Profile** configuration page. This allows enable or disable of header generation along with the flexibility of selecting correlation type to use for the header. The correlation-types such as SUPI, GPSI, or both are supported for this release.

The generated or received headers can only be forwarded when the setting *Send Correlation-Info Header* as part of **NF Communication Profile** is enabled. **NF Communication Profile** must be attached to **PCF Communication Profile** under the **Audit** section of the **Management Service** page

3gpp-Sbi-Correlation-Info

The header contains correlation information such as UE identifier related to the HTTP request or response.

Note

1. The possibility to include more than 1 correlationinfo parameter in the 3gpp-Sbi-Correlation-Info header is kept for future extensibility.
correlationinfo = ctype "-" cvalue

ctype = "imsi" / "impi" / "suci" / "nai" / "gci" / "gli" / "impu" / "msisdn" / "extid" / "imei" / "imeisv" / "mac" / "eui" / token
2. The token is defined for future extensibility.
The token of ctype shall not use the dash ("-") character.

cvalue = 1*tchar

Table 3-18 The format of cvalue shall comply with the data type description.

ctype	Description
SUPI	VarUeld format defined for IMSI and starting after the string "imsi-"
GPSI	VarUeld format defined for MSISDN and starting after the string "msisdn-"

Table 3-19 3GPP defined Custom HTTP Headers

Header	Description	Example
3gpp-sbi-correlation-info	This header may be used to contain correlation information such as UE identity, that may be used by an operator in various offline network management, performance analysis and troubleshooting tools/applications to identify messages (requests, responses, subscriptions, notifications) related to a particular subscriber.	EXAMPLE 1: When UE identifier used is SUPI and SUPI type is an IMSI: 3gpp-Sbi-Correlation-Info: imsi-345012123123123 EXAMPLE 2: When UE identifier used is GPSI and GPSI type is an MSISDN:3gpp-Sbi-Correlation-Info: msisdn-1234567890 EXAMPLE 3: When UE identifiers used are SUPI and GPSI where SUPI type is an IMSI and GPSI type is an MSISDN:3gpp-Sbi-Correlation-Info: imsi-345012123123123; msisdn-1234567890

Managing SBI messages correlation using Subscriber Identity

This section explains the procedure to enable and configure the feature.

Configure

In CNC Console, enable this feature in the **NF Communication Profile** page.

For more details on enabling or disabling the correlation-info header on GUI, see [Service Configurations](#).

Configure Using REST API

For configuring parameters for SBI messages correlation using Subscriber Identity feature using REST APIs, see *Oracle Communications Cloud Native Core, Binding Support Function REST API Specification Guide*.

Observe

Following metrics are added in [Correlation-Info Header Metrics](#) section:

- ocbsf_correlation_info_header_received
- ocbsf_correlation_info_header_forwarded
- ocbsf_correlation_info_header_generated

3.20 Configurations for Pre and Post Upgrade/Install Validations

This feature applies validation checks that are required on the application, databases, and its related tables before and after the upgrade/installation of BSF application.

On enabling this mandatory pre-flight and post-flight validation checks, for successful upgrade/installation following are validated:

- does the related database exists
- does all the required tables exist

- does the required table schema exist for all the required tables
- does all the required infrastructure exists

This pre-flight and post-flight checks ensures that all the dependent databases, tables, schema, applications are in right order for performing successful update/installation.

For more information on how to set the parameter value for pre and post flight checks, see *Upgrade Hardening, Pre and Post Flight Checks* section in *Oracle Communications Cloud Native core, Binding Support Function (BSF) Installation, Upgrade and Fault Recovery Guide*.

3.21 Detection and Handling of Late Arrival Requests

BSF receives requests from Ingress Gateway with the 3GPP headers. These requests help in detecting the response time for the BSF Management Service. The Ingress Gateway receives the following headers:

- 3gpp-Sbi-Origination-Timestamp- It contains the timestamp when the originating entity initiates the request
- 3gpp-Sbi-Max-Rsp-Time- The header indicates the duration (expressed in milliseconds) during which the HTTP client waits for a response.


If the configuration for determining the late arriving requests is enabled and BSF does not receive the required headers, the requests are valid and processed. The call flow will continue as normal.

BSF must be able to read the 3gpp-Sbi-Origination-Timestamp and 3gpp-Sbi-Max-Rsp-Time headers to respond to the request. Consider the following scenarios:

- If the sum of 3gpp-Sbi-Origination-Timestamp and 3gpp-Sbi-Max-Rsp-Time is less than the current time, the PCF service rejects the message with a 504 HTTP code and sends the message "TIMED_OUT_REQUEST".
- If the request does not include either 3gpp-Sbi-Origination-Timestamp or 3gpp-Sbi-Sender-Timestamp, or 3gpp-Sbi-Max-Rsp-Time headers, then the request is accepted. There are no changes in the call flow with the inclusion of collision detection.
- If 3gpp-Sbi-Max-Rsp-Time receives a negative value, the header is considered invalid. In this case, the service fallbacks to the default behavior and accepts the request irrespective of the 3gpp-Sbi-Origination-Timestamp or 3gpp-Sbi-Sender-Timestamp value.

The following diagram illustrates the request timeout in BSF:

Figure 3-22 Request Timeout in BSF



For handling and detection of Late Arrival and Collision Detection functionality in BSF, it is necessary that these headers and the new Custom header are transmitted from PCF to BSF.

Custom Header Enhancement adds the Time Stamp as a default value in the request and response custom headers with specific time formats and time zones. A valid function definition for timestamp in the configuration is: `func: currentTime(time-format,time-zone)`.

Note: Here, only the GMT, IST, PST, and UTC time zones are to be considered for all 5G timestamps.

The only date format supported for the headers is RFC 7231: `EEE, dd MMM yyyy HH:mm:ss.SSS zzz` (Sun, 04 Aug 2019 08:49:37.845 GMT).

Any other format will result in a parse error which leads to not using the header for the feature (taking it as null value).

Collision Detection

BSF may encounter a colliding or duplicate request for a binding registration from a different PCF. BSF shall consider a request as colliding or duplicate binding registration when two registrations are for the same subscriber (SUPI), DNN, SNSSAI, IPV4/IPV6 prefix, or IpDomain (in case IPV4Address is present) but from different PCF instances (ID/pcfDiamHost/pcfFqdn or pcfIpEndpoints).

In case support for timer headers is enabled:

- If 3gpp-Sbi-Origination-Timestamp header support is enabled and the colliding or the duplicate requests contain the header with an appropriate value, compare the values. If the values are different, it will consider the request with the more recent timestamp in the header.
- If the value of 3gpp-Sbi-Origination-Timestamp in both the requests is the same, then if Custom-Sbi-Sender-Timestamp header support is enabled and the colliding or the duplicate requests contain the header with an appropriate value, compare the values. It will consider the request with the more recent timestamp in the header.

In case support for timer headers is disabled:

- If 3gpp-Sbi-Sender-Timestamp header support is enabled and the colliding or the duplicate requests contain the header with an appropriate value, compare the values. If the values are different, it will consider the request with the more recent timestamp in the header.
- If the value of 3gpp-Sbi-Sender-Timestamp in both requests is the same, it will consider the later request.

If the incoming colliding or duplicate request has an older timestamp than the ongoing or existing request or record, it will reject the request with an HTTP "403 Forbidden" status code. The condition name for the error in the GUI is "Existing binding information found in DB" on the SBI Error Codes configuration page. The BSF function to detect and resolve collision requests is configurable.

If Late Arrival functionality is not enabled, the headers must be propagated to bsf-management-service if collision detection is required or enabled. For this, Ingress Gateway should populate three new collision headers following the required criteria:

Table 3-20 Collision Headers


At Ingress Gateway	At bsf-management-service
3gpp-Sbi-Origination-Timestamp	collision-3gpp-origination-timestamp
Custom-Sbi-Sender-Timestamp	collision-custom-sender-timestamp
3gpp-Sbi-Sender-Timestamp	collision-3gpp-sender-timestamp

Note

These collision headers implement collision detection at BSF.

Enable

You can enable or disable the collision detection in the General Settings of the General Configuration window of the CNC Console.

Figure 3-23 Enabling Collision Detection

You can enable or disable the collision detection feature using the REST API for BSF. Use the General Settings REST API to enable or disable this feature.

REST API Path: /oc-bsf-configuration/v1/general

For more information, see the *Oracle Communications Cloud Native Core Binding Support Function User Guide*.

Configuration

To create the collision headers that are used for collision detection in BSF it is necessary to add the following configuration to BSF Ingress Gateway.

```
routesConfig:
  - id: bsf_management_register
    uri: http://{ template "service-name-bsf-management" . } } :
  { { .Values.global.servicePorts.bsfManagementServiceHttp } }
    path: /nbsf-management/**
    order: 1
    method: POST
    filters:
      customReqHeaderEntryFilter:
        headers:
          - methods:
              - POST
            headersList:
              - headerName: 3gpp-Sbi-Message-Priority
                defaultVal: 24
                source: incomingReq
                sourceHeader: 3gpp-Sbi-Message-Priority
                override: false
              - headerName: collision-3gpp-origination-timestamp
```

```
source: incomingReq
sourceHeader: 3gpp-Sbi-Origination-Timestamp
override: false
- headerName: collision-custom-sender-timestamp
source: incomingReq
sourceHeader: Custom-Sbi-Sender-Timestamp
override: false

- headerName: collision-3gpp-sender-timestamp
source: incomingReq
sourceHeader: 3gpp-Sbi-Sender-Timestamp
override: false
```

Observe

To observe the collision detection functionality, you can use metrics that are specific to BSF management service. For information, see [BSF Metrics](#).

3.22 Support for Timer Configuration

BSF supports the configuration of Diameter interfaces or Rx interfaces timers for all the applicable Diameter messages. The timer configuration needs to be configured for AAR, RAR, STR, and ASR messages. The timer configuration is an option to configure the diameter response timeout value for diameter messages. If the timer value is not configured, then the BSF Diameter gateway works with the default value of 4000 milliseconds for AAR, RAR, STR, and ASR messages.

The Diameter response timeout or the timer value is configured per Diameter interface level. This value can also be configured per message level of the Diameter interface. The timer configuration needs to be configured for Authentication Request (AAR), Re-Auth-Request (RAR), Session-Termination-Request (STR), and Abort-Session-Request (ASR) messages.

For BSF, currently Diameter Rx interface is applicable. The response timeout value is configured using the Rx application level and its messages (AAR, RAR, STR, and ASR) level.


Enable

The timer configuration is a functionality supported by Binding Support Function configurations. You do not need to enable or disable this feature.

Configure


The `Binding discovery request timeout` configuration parameter is used for request timeout value for the discovery request sent by the BSF Diameter Gateway towards the BSF Management Service.

Figure 3-24 General Settings



The Enhanced Timer Configuration is available on the settings of the Diameter Configuration page.

Figure 3-25 Enhanced Timer Configuration



This configuration accepts interface level Diameter response timeout value. If the Diameter interface level timeout is configured, then the user is able to configure message level response timeout value for the corresponding Diameter interface. The response timeout value range for interface and message level is 3000 milliseconds to 2147483647 milliseconds.

In the CNC Console, you can specify the Command Code Response TimeOut value for the AAR, RAR, STR, and ASR message types.

Figure 3-26 Edit Enhanced Timer Configuration

Edit Enhanced Timer Configuration

Application Name:
Rx

Application Response Timeout (milliseconds):
8,000

✓ **Command Code Response TimeOut**

AAR (milliseconds):
6,000

STR (milliseconds):
5,000

RAR (milliseconds):
RAR (milliseconds)

Save Cancel

You can customize the configurations related to this feature using the CNC Console or REST APIs for BSF.

- **Configure using CNC Console:** Perform the feature configurations on the General Configurations and Diameter Configurations page. For more information about the configurations, see [General Settings](#) and [Diameter Configurations](#).
- **Configure using REST API:** Perform the configurations using GET and PUT operations. For more information about REST API configuration, see *Oracle Communications Cloud Native Core, Binding Support Function REST Specification Document*.

Observe

To observe the timer configuration functionality, you can use metrics that are specific to BSF management service. For information, see [BSF Metrics](#).

3.23 Support for Session Retry and Alternate Route Service

In previous releases, Binding Support Function was configured with primary and secondary Network Repository Function (NRF) statically and limiting to a specific number. Starting with Release 1.8.0 of BSF, session retry enables the alternate recovery mechanisms to mitigate the impact of any unavailable resource.

This feature allows you to configure virtual FQDNs and perform DNS SRV Lookup to retrieve alternate failover NRF which can be maintained dynamically at the DNS Server.

Managing Session Retry and Alternate Route Service

Enable

To enable the Session Retry and Alternate Routing functionality, set the value of `enableVirtualNrfResolution` to true in the `custom-values.yaml` file for BSF. For more information on setting the parameter value, see section "Configuring NRF Client" in *Oracle Communications Cloud Native Binding Support Function Installation Guide*.

Configure

You can configure the parameter for Session Retry and Alternate Routing functionality by updating the `custom-values.yaml` file for BSF.

- To configure the retry functionality, see "Configuring NRF Client" in *Oracle Communications Cloud Native Binding Support Function Installation Guide*.
- To configure alternate routing, see section "Alternate Route Service Configuration" in *Oracle Communications Cloud Native Binding Support Function Installation Guide*.

3.24 XFCC Header Validation

Overview

With XFCC Header Validation feature, Binding Support Function (BSF) as a producer, checks if the SCP that is sending the HTTP request is the same SCP that is configured in the BSF. BSF performs this check by comparing the FQDN of the SCP present in the "x-forwarded-client-cert" (XFCC) of http2 header with the list of FQDN of the SCPs configured in the PCF. This configured list contains all the host FQDNs resolved successfully via DNS-SRV as well as static SCPs. The header validation can be enabled at global as well as at the route level.

Note

This feature is applicable only when SCP is deployed in the network topology.

Configuring SCPs at BSF

To configure SCP, you need to customize `custom.yaml` at the time of deploying BSF.

In the earlier releases, users could only configure SCPs statically as shown in the following snippet:

```
xfccHeaderValidation:
  validation:

    enabled: false
    nfList:
```

- scp.com
- smf.com
- amf.com

However, in BSF release 22.1.0 or later, users can configure single or multiple virtual FQDNs for the SCP along with the static configuration as shown in the following snippet:

```
global:
  xfccHeaderValidation:
    validation:
      enabled: false
      peerList:
        - name: scp.com
        - name: smf.com
        - name: amf.com
        - name: scp1.com
          enabled: true
        - name: scp2.com
        - name: scp3.com
          enabled: false
        - name: xyz.test.com
          enabled: true
          scheme: http
          type: virtual
        - name: abc.test.com
          enabled: true
          scheme: https
          type: virtual
        - name: xfcc.test.com
          enabled: false
          scheme: http
          type: virtual
```

Static SCP: To define an SCP instance statically, add the name and set enabled parameter to true in the peerList. If the enabled parameter is set to false for an instance, then it is not included in the list of configured FQDNs. If you do not specify enabled parameter then by default it is considered as true.

Virtual SCP: To define an SCP with virtual FQDN, add the name, scheme as http or https, type as virtual, and set enabled parameter to true. If the enabled parameter is set to false for an instance, then it is not included in the list of configured FQDNs.

Resolving FQDNs to find Authorized SCPs

During the bootup of Ingress Gateway, it tries to resolve the configured virtual FQDN via Alternate Route service using the following helm configuration:

```
dnsSrv:
  port: *svcAlternateRouteServiceHttp #Alternate-route port for scheme
  'http'. Change is required if the scheme below changes.
  scheme: http
```

If Alternate Route service is unable to resolve the configured virtual host, Ingress Gateway stores it in the list of failed FQDNs and reattempts the request at 300 s (default value configured for **dnsResolutionInterval**).

The following metric is used when the request to resolve configured virtual FQDNs is unsuccessful:

- `oc_ingressgateway_dns_resolution`: This metric is pegged when DNS resolution for a given FQDN fails.
- `oc_ingressgateway_dns_resolution_failure`: This is a gauge metric that is triggered when DNS resolution for a given FQDN fails.

Handling Traffic Flow

The XFCC header is validated when:

- a single XFCC header present in the incoming request to IGW
- multiple XFCC headers are present in the incoming request to IGW

Validating single XFCC Header

The following figure describes the call flow for validation of a single XFCC header:

Figure 3-27 Call Flow Validation of Single XFCC Header




Table 3-21 Single XFCC Header Configuration

Scenario	Condition	Action
Ensure that the XFCC Header validation parameter is enabled for the corresponding route match and the matchCerts count is configured correctly. Given XFCC Header validation parameter is enabled for the corresponding route match and matchCerts count correctly configured.	When the matchField parameter of the client certificate field in XFCC header matches with one of the configured NF FQDNs .	Forwards the request to a back-end microservice and receives a corresponding response.

Table 3-21 (Cont.) Single XFCC Header Configuration

Scenario	Condition	Action
Given XFCC header validation parameter is enabled for the corresponding route match and matchCerts count correctly configured.	When matchField parameter of client certificate field in XFCC header does not match with the configured NF FQDNs.	Return a 400 Bad Request response from Ingress Gateway. For more information about error codes, see SBI Error Codes Configurations
Given XFCC header validation parameter disabled for the corresponding route match.	NA	Forwards the request to the back-end microservice and receives a corresponding response.

Example of a single XFCC header request:

```
x-forwarded-client-cert: By=http://
router1.blr.com;Hash=468ed33be74eee6556d90c0149c1309e9ba61d6425303443c0748a02dd8d
e68; Subject="/C=US/ST=CA/L=San Francisco/OU=Lyft/CN=Test Client"; URI=http://
testenv1.blr.com;DNS=blr.com; DNS=www.blr.com
```

Validating multiple XFCC Headers

The following figure describes the call flow for validation of multiple XFCC headers for the following scenarios:


Figure 3-28 Call Flow Validation of Multiple XFCC Headers

Table 3-22 Multiple XFCC Headers Configuration

Scenario	Condition	Action
Given XFCC header validation parameter is enabled and matchCerts count correctly configured to validate across XFCC header certificates from the right most entry.	When matchField parameter of the corresponding client certificate field being validated against currently in the corresponding XFCC header matches with the NF FQDN's configured at Ingress Gateway.	Consider the request as a valid request and forward the request to the back-end micro-service and receive a corresponding response.
Given XFCC header validation parameter is enabled and matchCerts count correctly configured to validate across XFCC header certificates from the right most entry.	When matchField parameter of client certificate field in corresponding XFCC headers do not match with the NF FQDN's configured at Ingress Gateway for the corresponding matchCerts count.	Consider the request as an invalid request and return a 400 Bad Request response from IGW. For more information about error codes, see SBI Error Codes Configurations
Given XFCC header validation parameter is enabled and matchCerts count -1.	NA	Consider the request as valid request and match against the corresponding match field in all XFCC headers, if validation successful then forward the request else return 400 BAD Request.
Given XFCC header validation parameter is disabled.	NA	Forward the request to back-end microservice and receive a corresponding response.

Example of multiple XFCC header request:

```
x-forwarded-client-cert:By=http://
router1.blr.com;Hash=468ed33be74eee6556d90c0149c1309e9ba61d6425303443c0748a02dd8d
e68; Subject="/C=US/ST=CA/L=San Francisco/OU=Lyft/CN=nf1.com"; URI=http://
testenv1.blr.com; DNS=nf8.com; DNS=nf1.com; DNS=nf6.com, By=http://
router1.blr.com;Hash=468ed33be74eee6556d90c0149c1309e9ba61d6425303443c0748a02dd8d
e68; Subject="/C=US/ST=CA/L=San Francisco/OU=Lyft/CN=nf10.com"; URI=http://
testenv1.blr.com; DNS=nf10.com; DNS=nf8.com; DNS=nf9.com, By=http://
routexr1.blr.com;Hash=468ed33be74eee6556d90c0149c1309e9ba61d6425303443c0748a02dd8
de68; Subject="/C=US/ST=CA/L=San Francisco/OU=Lyft/CN=nf4.com"; URI=http://
testenv1.blr.com; DNS=nf9.com; DNS=nf4.com;DNS=nf1.com
```

Managing XFCC Header Validation

- **Global Level:** To enable or disable the XFCC header validation feature, set the value of the `ingress-gateway.global.xfccHeaderValidation.validation.enabled` to true or false respectively.
- **Route Level:** To enable or disable the XFCC header validation feature at route level, set the value of the `xfccHeaderValidation.validationEnabled` under `routesConfig` to true or false respectively.

Note

If the `xfccHeaderValidation.validationEnabled` parameter is defined at route level, then the configuration takes precedence over global configuration.

For instance, if you want to enable XFCC header validation for selected routes, then set the global parameter as false and make route specific configuration to true.

```
global:
  xfccHeaderValidation:
    validation:
      enabled: false

routesConfig:
  - id: reverse_bsf_service
    uri: http://{ template "service-prefix" . }-bsf-management:
  {{ .Values.global.servicePorts.bsfManagementServiceHttp }}
    path: /nbsf-management/**
    order: 1
    metadata:
      xfccHeaderValidation:
        validationEnabled: true
```

For more information about setting the parameter values, see section "XFCC Header Validation Configuration" in *Oracle Communications Cloud Native Binding Support Function Installation Guide*.

Configure

You can configure the parameter for XFCC Header Validation by updating the custom-values.yaml file for BSF. For more information about configuring the parameter value, see the "XFCC Header Validation Configuration" section in *Oracle Communications Cloud Native Binding Support Function Installation Guide*.

Observe

To observe the XFCC header validation functionality, you can use metrics that are specific to Ingress Gateway. For information, see [Ingress Gateway Metrics](#).

Configuring Error Codes

When the XFCC header validation feature is enabled and SCP FQDN in the incoming header does not match the configured FQDN in PCF, XFCC header is not present, or XFCC header is invalid, then PCF may return error in the response. Users have the ability to customize the error code returned in the response using the following helm configuration:

```
errorTrigger:
  - exceptionType: XFCC_HEADER_INVALID
    errorCode: '401'
    errorCause: xfcc header is invalid
    errorTitle: 'Invalid XFCC Header'
    errorDescription: 'Invalid XFCC Header'
  - exceptionType: XFCC_MATCHCERTCOUNT_GREATER_THAN_CERTS_IN_HEADER
    errorCode: '402'
    errorCause: matchCerts count is greater than the certs in the
request
    errorTitle: ''
    errorDescription: ''
  - exceptionType: XFCC_HEADER_NOT_PRESENT_OR_EMPTY
    errorCode: '403'
    errorCause: xfcc header is not present or empty in the request
```

```
errorTitle: ''  
errorDescription: ''
```

If the configured error code in the `errorCodeOnValidationFailure` field lies in 3xx error series only then the values for `retryAfter` and `redirectUrl` if configured under XFCC Header Validation Configuration at Ingress Gateway are used to populate Retry-After and LOCATION headers correspondingly while sending error response from Ingress Gateway.

3.25 Georedundancy Support

The Cloud Native Core (CNC) architecture supports Geographically Redundant (Georedundant) BSF deployments to ensure high availability and redundancy. It offers two, three, or four sites georedundancy to ensure service availability when one of the BSF sites is down.


The specifications for georedundancy feature are as follows:

- All the georedundant sites must have Helm and REST based configurations except for NF InstanceId, BSF Endpoint, and port.
- The georedundant BSF sites must be reachable from NFs or Peers on all the sites.
- The same NFs or Peers must not communicate to other georedundant BSF sites simultaneously for the same session.
- All the sites must register with NRF independently and work in an active state.
- All BSF instances share the **Session State** data by using the DB Tier replication service. This enables service continuity during the failure of any of the sites.
- The NFs in a given site can discover BSF instances through NRF. However, local configurations such as DNS SRV or static configuration are required to determine the primary and secondary or alternate BSF configuration. When the primary instance is available, the NFs send service requests to the primary instance.
- NRF always reflects current functional status of a given BSF instance. Thus, during the failure of a given BSF instance, the value of `NfStatus` is updated to **SUSPENDED** by either NRF or BSF instance. Therefore, when NF detects failure of primary instance due to error response or status notification from NRF, the NF redirects its traffic to the secondary instance, until the primary instance becomes available again.

BSF supports the following types of georedundant deployment:

Two-Site Georedundancy Deployment


The following diagram depicts the topology for two-site georedundant BSF deployment:

Figure 3-29 Two-Site Georedundancy

After the second site instance of the cnDBTier is created, you can establish the two site georedundant connections that provide bi-directional data replication between both sites. Therefore, when the records are updated at one site, these changes are replicated to the other remote site in real-time. These updates can be creating, changing, or deleting a record.

Three-Site Georedundancy Deployment

The following diagram depicts the topology for three site georedundant BSF deployment:

Figure 3-30 Three-Site Georedundancy

In case of three site georedundancy, bi-directional replication is established from each site to the other two sites. The database updates from each site are replicated to the other two sites over the replication channel.

The advantages of three-site georedundancy are:


- The service remains available even when 2 sites become unavailable.
- In case of a single site failure, the remaining two sites keep establishing the bi-directional replication.
- No action is required in case of a site failure.
- Requires 4 SQL pods and 2 rep-svc pods at each site

When the records are updated at one site, these changes are replicated to the other two remote sites in real-time. These updates can be creating, changing, or deleting a record.

Four-Site Georedundancy Deployment

The following diagram depicts the topology for four-site georedundant BSF deployment:

Figure 3-31 Four-Site Georedundancy Deployment



cnDBTier supports the four-site georedundancy deployment. In case of four-site georedundancy, each site participates in a 4-way replication. The database updates from each site are replicated to the other three-sites over the replication channels.

The advantages of four-site georedundancy is:

- In case of a single site failure, the remaining three-sites keep establishing the bi-directional replication.

- No action is required in case of a site failure.
- Requires 6 SQL pods and 3 db-rep-svc at each site.
- Each site uses two SQL nodes for active and standby replication channels for high availability of the replication channels.

When the four-sites are correctly replicated, then any update done at one site is replicated to the other three remote sites in real-time. The changes include creating, changing, or deleting a record.

Managing Georedundancy

Deploy

To deploy BSF in a georedundant environment:

1. Set up the replicated cnDBTier version 1.8.0.0.3 or above on two or three sites as required. For more information about installing cnDBTier, see "Installing cnDBTier" in *Oracle Communications cnDBTier Installation Guide*.
2. Deploy BSF over the replicated (two or three) cnDBTier sites. For more information about installing and deploying BSF, see *Oracle Communications Cloud Native Binding Support Function Installation Guide*.

Configure

To configure georedundancy:

You need to configure the georedundancy functionality while deploying the BSF instances on the replicated sites. The following parameters must be updated in the custom-values.yaml file for BSF:

Table 3-23 Georedundancy Parameters

Parameter	Description
global.envMysqlHost	The database instance for each site. BSF communicates to the database at the same site only.
global.nfInstanceId	The ID for the site
config-server.envMysqlDatabase	The database for the config server. The two sites must use different database names for config server
cm-service.envCommonConfigMysqlDatabase	The common configuration database. It must be different on the sites
nrf-client.configmapApplicationConfig.profile	Configuration data for nrf client. The <code>appProfile</code> and the <code>nfInstanceId</code> parameters must be aligned with <code>global.nfInstanceId</code>
nrf-client-nfdiscovery.dbConfig.dbName	The common configuration database. It must be different on the sites
nrf-client-nfmanagement.dbConfig.dbName	The common configuration database. It must be different on the sites
appinfo.dbConfig.dbName	The common configuration database. It must be different on the sites
perf-info.dbConfig.dbName	The common configuration database. It must be different on the sites

Table 3-23 (Cont.) Georedundancy Parameters

Parameter	Description
BSFds.envMySQLDatabaseConfigServer	The database for the config server. The two sites must use different database names for config server
ingress-gateway.dbConfig.dbName	The common configuration database for ingress gateway. It must be different on the sites
egress-gateway.dbConfig.dbName	The common configuration database for egress gateway. It must be different on the sites
alternate-route.dbConfig.dbName	The common configuration database for alternate route. It must be different on the sites

For more information about configuring the parameter value, see the "Alternate Route Service Configuration" section in *Oracle Communications Cloud Native Binding Support Function Installation Guide*.

Observe

cnDBTier generates critical alerts in case of application or database failure. For more information, see *Oracle Communications Cloud Native Core DBTier User Guide*.


Maintain

BSF allows you to monitor the georedundancy deployment through cnDBTier alerts. Access the Prometheus GUI to check for new App alerts.

3.26 Diameter Gateway Pod Congestion Control

The Diameter Gateway is a diameter proxy agent for Binding Support Function (BSF). It is a front-end microservice for diameter traffic for both Ingress and Egress traffic and can get congested due to higher traffic, higher CPU usage, and higher memory utilization. Thus, it is imperative to have suitable congestion control features in place for Diameter Gateway pods to avoid adverse impacts on latency and performance.

Another reason for the need for a congestion control mechanism for Diameter Gateway is the nature of diameter connections. An external LoadBalancer distributes these long-lived connections. As shown in the following image, when the LoadBalancer routes an incoming request from network to Diameter Gateway pod (indicated by green line), the balancer does not take health or load of the pod into consideration. As a result, uneven distribution of traffic can happen, and gateway pods experience congestion.



Congestion control helps to identify such conditions and invoke load shedding rules to address these situations when these load conditions persist. With Diameter Gateway service supporting the congestion control mechanism feature, it helps to protect its pods from congestion and ensures consistent service availability to its consumer.

In Binding Support Function, a congestion control mechanism is implemented at pod level that allows the system to perform the following tasks:

- Determine the pod congestion state
- Trigger Congestion Control

Determining Pod Congestion State

At any given point of time, a pod can be in any one of the following states:

- Normal
- Danger of Congestion (DOC)
- CONGESTION_L1
- CONGESTION_L2
- CONGESTED

Figure 3-32 Different Pod Congestion States



To decide the state of a pod, the following points are taken into consideration:

1. **CPU:** The CPU usage for congestion state is calculated by comparing the CPU usage of the container (monitored using cgroup parameter, `cpuacct.usage`, which provides current cpu usage in nanoseconds) with the configured threshold.

2. **Queue:** For the DOC, CONGESTION_L1, CONGESTION_L2, and CONGESTED pod states, compare the number of pending messages in the queue with the configured pending messages threshold.

To avoid switching between these states due to traffic pattern, it is required for the pod to be in a particular state for a given period before transitioning to another state. Periodically, the state of the pod's congestion gets determined. The following configurations are used to define the period that the pod has to be in a particular state for:

- **stateChangeSampleCount:** This REST API parameter can be configured to specify after how many continuous intervals, the pod state can be changed. The default value is 5. This value can range from 1 to 2147483647.
- **stateCalculationInterval:** This REST API parameter can be configured to specify the time duration or interval, after which the pod congestion state will be re-verified. The default value is 200ms. This interval is configured in milliseconds and can range from 50 to 2147483647.

For more information about these congestion control configurations, see [Congestion Control Settings](#) section in CNC Console.

CPU and Queue Threshold

A pod's state changes when its resources, namely the CPU and Queue buffer usage exceeds the predefined congestion threshold, checked at regular intervals. An active CPU and Queue default threshold profile can be viewed using CNC Console at *BSF → Overload and Congestion Control → Congestion Control → Thresholds* section. The user can create new threshold profile with the congestion states being in the order of CONGESTED → CONGESTION_L1 → CONGESTION_L2 → CONGESTED states, and make it active. The Diameter Gateway service pod's congestion states and their default congestion parameters, CPU, and Queue counts are provided in the following table:

Table 3-24 Diameter Gateway service Congestion States

Congestion States	CPU Count	Queue (Pending Request)
DANGER_OF_CONGESTION (DOC)	60	6300
CONGESTION_L1	70	7200
CONGESTION_L2	75	8100
CONGESTED	80	9000

For more information about these CPU and Queue count thresholds configurations, see [Congestion Control Thresholds](#) section in CNC Console.

Triggering Congestion Control

With Congestion Control feature enabled, every time a message is fetched for processing, the system checks the current congestion state of the pod. If the current state is either in DOC, CONGESTION_L1, CONGESTION_L2, or Congested state, then the congestion control mechanism is triggered. After verifying that the message type is a request, a request priority is assigned to it. If the request priority is greater than or equal to the discard priority, then the message is rejected, otherwise it is accepted.

Congestion Control Load Shedding using Discard and Request Priority

Diameter Gateway pods performs load shedding by considering the load rule configured for the current congestion state. The load shedding rule is based on message priority configuration. A default load shedding rule is active and the messages are discarded based on it. You can

configure these load rules with the congestion states being in the order of CONGESTED → CONGESTION_L1 → CONGESTION_L2 → CONGESTED states, using the Overload and Congestion Control menu in CNC Console. The Diameter Gateway discards the requests based on either:

- Priority
- Priority and Percentage

Table 3-25 Priority based Message Discards

Congestion State	Discard Priority
DANGER_OF_CONGESTION	27
CONGESTION_L1	19
CONGESTION_L2	17
CONGESTED	15

Table 3-26 Priority and Percentage based Discards

Congestion State	Discard Priority Range	Discard Percentage
DANGER_OF_CONGESTION	27-31	50
CONGESTION_L1	19-26	50
	27-31	75
CONGESTION_L2	17-18	50
	19-26	75
	27-31	100
CONGESTED	15-31	100

For more information about these Priority or Priority and Percentage discards configurations, see Congestion Control [Load Shedding Rules](#) page in CNC Console.

The following two AVP's are added in Pre-Defined AVP Conditions in Diameter Gateway Message Priority Rules configurations at *BSF → Overload and Congestion Control → Diameter → Message Priority Profiles* in CNC Console.

Table 3-27 Pre-Defined AVP Conditions

Pre-Defined AVP Conditions	Type Operator	Values
DRMP	• Equals • Not Equals • Less Than • Less Than Equals • Greater Than • Greater Than Equals	The Diameter Routing Message Priority (DRMP) AVP allows Diameter Gateway to indicate the priority of diameter messages. This will take numeric values. Allowed numeric value in the range of 0-15.
OC-Message-Priority	• Equals • Not Equals • Less Than • Less Than Equals • Greater Than • Greater Than Equals	The header attribute allows Diameter Gateway to indicate the priority of diameter messages. This will take numeric values. Allowed numeric value in the range of 0-31.

In the Message Priority Profiles for Diameter Gateway, if both DRMP Priority and OC-Message-Priority is selected, then OC-Message-Priority shall be considered as highest priority for message priority rule selection. The message priority range for OC-Message-Priority is 0-31, and for DRMP AVP it is 0-15.

For more information about configuring message priority profiles, see [Message Priority Profiles](#) section in CNC Console.

Congestion Response Code

After defining load rules Overload and Congestion Control menu, the response to reject the request can also be configured using the advanced settings in Diameter configurations. The following are the advanced settings key and value:

Table 3-28 Advanced Settings

Key	Value	Description
CONGESTION_DIAMETER_RESPONSE_CODE	3004	This is to configure the diameter result code for the congestion control discards and its value shall be numerical. The value can be set to either 3xxx, 4xxx, or 5xxx response codes. If this value is not configured, then default congestion result code shall be 3004 (<i>DIAMETER_TOO_BUSY</i>). Default Value: 3004 Examples: CONGESTION_DIAMETER_RESPONSE_CODE=5065 CONGESTION_DIAMETER_RESPONSE_CODE=3004
CONGESTION_DIAMETER_VENDOR_ID	0	This is to configure the diameter vendor id for the congestion control discards and its value shall be numerical. If the value is non-zero, then the <i>CONGESTION_DIAMETER_RESPONSE_CODE</i> shall be considered as the experimental result code. If value is zero or it is not configured, then <i>CONGESTION_DIAMETER_RESPONSE_CODE</i> shall be considered as the normal result code. Examples: CONGESTION_DIAMETER_VENDOR_ID=10415 CONGESTION_DIAMETER_VENDOR_ID=21274

If above advance settings keys are not configured, then Diameter Gateway will discard the messages with default diameter result code 3004 *DIAMETER_TOO_BUSY*.

For more information about different response codes that can be configured, see [Congestion Control Discard Message Response Code](#) section.

Call Flow

To best understand how the diameter congestion control feature works, consider a scenario where the sample message priority profiles rules are configured as described in the following table:


Table 3-29 Message Priority Profiles

Message	Priority
Default	6
Rx RAR	7
Gx CCR-I	10

Table 3-29 (Cont.) Message Priority Profiles

Message	Priority
Sy SLR	11

Figure 3-33 Call flow diagram for Diameter Congestion Control



Message will be rejected when message priority is greater than discard priority

- When Diameter Gateway 2, currently is either in DOC or Congested state, receives RAR request, the priority is compared against discard priority. As the message priority is greater than the discard priority, the message is rejected with result code – DIAMETER_UNABLE_TO_COMPLY.
- When Diameter Gateway 2, currently is either in DOC or Congested state, receives AAR-U request, the priority is compared against discard priority. As message priority is higher than discard priority, the message is accepted with result code – DIAMETER_SUCCESS.

Feature Upgrade Impact

With upgrade of BSF application from any of the previous release versions to 25.1.200, the Diameter Gateway pod Congestion Control feature can be updated to use the common Congestion Control mechanism available in 25.1.200. In order to do this, there are some considerations to address, especially due to the difference in configuration data. Diameter Gateway's pod congestion control configurations from the previous version should be migrated to use the current version configurations. The data migration process is manual and one-time activity that user has to perform, by using either the CNC Console or Congestion Control migration APIs.

Feature Transition

Diameter Gateway Congestion Control feature transition between older version to the current version's Congestion Control mechanism is based on migration status of the BSF application during the upgrade. The migration status can be either in "False" or "True".

False (Before Migration): After upgrading to the current BSF version, the Diameter Gateway pod congestion control feature still operates under the configurations and designs of the pre-upgrade version until the user actively migrates the setup. In this "Before Migration" state, the users will only see:

- **Old Configurations:** The Diameter Gateway Congestion Control configurations from the previous application version.
- **Previous CNC Console pages:** User will have access to the older user interface rather than the new screen designed for Congestion Control feature.
- **Legacy REST APIs:** Calls to APIs will still work with the previous version's endpoints.

To enable the common Congestion Control mechanism for Diameter Gateway, users will need to complete the data migration process. For more information related to Diameter Gateway Congestion Control configuration migration, see "Diameter Gateway Migration of Congestion Control Configurations" section in Congestion Control [Settings](#) page in CNC Console.

True (After Migration): After the upgrade to current BSF version, perform the data migration for Diameter Gateway pod Congestion Control feature. With this, the migration status will be "After Migration" and Diameter Gateway services will be able to use the updated Congestion Control mechanisms.

Managing Diameter Gateway Pod Congestion Control

Enable

By default, the Pod Congestion Control is disabled for Diameter Gateway. You can enable this feature using CNC Console or REST API for BSF.

Feature Migration

If the BSF application is upgraded to the current version, then you can enable this feature by performing the migration process. For more information how to perform the data migration, see "Diameter Gateway Migration of Congestion Control Configurations" section in [Settings](#) page using CNC Console.

Configure

You can customize the configurations related to this feature using the CNC Console or REST APIs for BSF.

- **Configure using CNC Console:**
To enable the feature using CNC Console, set the **Enable** parameter in **Settings** page under **Congestion Control** for **Overload and Congestion Control Configurations** for Diameter Gateway Service.

To configure Diameter Gateway service pod congestion control feature in CNC Console, see [Congestion Control](#) page in CNC Console.
- **Configure using REST API:**
To enable this feature using REST API, perform the feature configurations as described in "Congestion Control" APIs section in *Oracle Communications Cloud Native Core Binding Support Function REST API Specification Guide*.

If the BSF application is upgraded to 25.1.200, then you can enable this feature by performing the migration process. For more information how to perform the data migration using Congestion Control Migration APIs, see "Congestion Migration" APIs section in

*Oracle Communications Cloud Native Core Binding Support Function REST API
Specification Guide.*

Observability

Binding Support Function uses the pod congestion metrics for this feature. For more information, see the [BSF Metrics](#) section. Alerts are raised when the following metrics are pegged:

- `pod_congestion_state`
- `pod_resource_congestion_state`

Note

Prometheus automatically injects name of the pod name with the label "kubernetes_pod_name" to the metric. This information is further used for alerting purposes.

Alerts

BSF uses the following congestion control alerts for this feature:

- [POD_CPU_DOC](#)
- [POD_DOC](#)
- [POD_CPU_DOC](#)
- [POD_CONGESTED](#)
- [POD_CPU_CONGESTED](#)
- [POD_CONGESTION_L1](#)
- [POD_CPU_CONGESTION_L1](#)
- [POD_CONGESTION_L2](#)
- [POD_CPU_CONGESTION_L2](#)

For more information, see the [List of Alerts](#) section.

Maintain

Warning logs are generated to indicate the congestion level. Error logs are generated when the system is congested and the actions are needed to be taken to bring the system back to normal. However, no error logs are generated when messages are rejected to avoid additional resource usage to write error logs.

3.26.1 Support for Server Header

BSF handles various requests from consumer Network Functions (NFs) and other network entities over HTTP protocol. On receiving these requests, BSF validates and processes them before responding to these requests. In case, BSF sends an error response, then the consumer NFs need to know the source of the error to trouble shoot the error and take corrective measures. The integration of this feature at BSF helps to determine the originator of the error response.

This feature offers the support for Server Header in BSF responses, which contains information about the origin of an error response and the type of the error encountered. The Server Header

includes the type of NF as "NF Type", followed by a "-" and the identity of the NF or the network entity. It is expected to be present in all BSF responses in the following format:

<NF_Type>-<NF_Instance_Id>

Where,

- <NF Type> is the type of the NF generating the error.
- <NF Instance-Id> is the unique identifier of the NF instance generating the error response.

For example: BSF-54804518-4191-46b3-955c-ac631f953ed8

The inclusion of the Server header in the BSF response is configurable, and can be enabled or disabled using a flag. Also the error codes that are included as part of the Server header in the error response are also configurable. The configuration of these parameters are done through either with REST APIs that are exposed through configuration server or Helm Configurations.

The operation mode that is either REST or HELM for Server Header configuration is done using the below flag:

ingress-gateway:

serverHeaderConfigMode: REST # Possible values: HELM, REST. Based on this value, the feature flag for "server" header will need to be enabled either in Helm configuration or Rest configuration.

Note

Nf Type and Nf Instance Id are mandatory fields for Server Header to get included in the error response. If either of the fields Nf Type or Nf Instance Id are configured as empty, then the Server Header will not get included in the error response.

Managing Server Header

Enable

By default, this feature is disabled.

You can enable the Server Header feature using Helm or REST API configurations:

- **Helm:** To enable the server header feature using Helm configuration, set the value for parameter `serverHeaderConfigMode` to `HELM` in the `custom-values.yaml` file. Then, set the value for parameter `serverHeaderDetails.enabled` to `true` under `global` and `routesConfig` for `ingress-gateway`.
- **REST API:** To enable the server header feature using REST configuration, set the value for parameter `serverHeaderConfigMode` to `REST` in the `custom-values.yaml` file. Using REST API, set the enabled parameter to `true` in the following resource URI:
`{apiRoot}/BSF/nf-common-component/v1/igw/serverheaderdetails`

Configure

You can configure the server header feature using the REST API or CNC Console:

- **Configure using REST API:**
Perform the REST API configurations in the following sequence to configure this feature:
 1. Configure **serverheaderdetails** to enable the feature.

`{apiRoot}/BSF/nf-common-component/v1/igw/serverheaderdetails`

2. Configure **routesconfiguration** to map route ID and its corresponding route-level configuration.
`{apiRoot}/BSF/nf-common-component/v1/igw/routesconfiguration`
3. Configure **errorcodeserieslist** to update the **errorcodeserieslist** that are used to list the configurable exception or error for an error scenario in Ingress Gateway.
`{apiRoot}/BSF/nf-common-component/v1/{serviceName}/errorcodeserieslist`

Note

If you define server header configuration at both global and route levels, the route level configuration takes precedence over the global level configuration.

For more information, see the "Server Header at Ingress Gateway" section in *Oracle Communications Cloud Native Core, Binding Support Function REST Specification Guide*.

- **Configure using Helm:** When parameter `serverHeaderConfigMode` is set to HELM and to configure the Server Header at Ingress Gateway, you need to perform the helm configurations either at Global or at Route level.
Following Helm Configuration performed at Global Level:

All attributes under "serverHeaderDetails" will need to be configured only if "serverHeaderConfigMode" is set as "HELM"

```
serverHeaderDetails:
  enabled: true
  errorCodeSeriesId: E1
  configuration:
    nfType: BSF
    nfInstanceId: INS-1
```

Use below configuration to define `errorCodeSeries` list

`errorCodeSeriesList`:

Value of "id" attribute will need to be used for assigning "errorCodeSeriesId" either at Global or Route level conf for Server header.

```
- id: E1
  errorCodeSeries:
    - errorSet: 4xx
      errorCodes:
        - 400
        - 408
    - errorSet: 5xx
      errorCodes:
        - 500
        - 503
- id: E2
  errorCodeSeries:
    - errorSet: 4xx
      errorCodes:
        - -1
```

Following Helm Configuration performed at Route Level:

```

routesConfig:
- id: backend_ms1_route
  uri: https://backend-ms1:8440/
  path: /ms1/**
  order: 1
  metadata:
    # All attributes under "serverHeaderDetails" will need to be
    configured only if "serverHeaderConfigMode" is set as "HELM" and Route
    level configuration is required. If not defined, Global configurations
    will be used
    serverHeaderDetails:
      enabled: true      # Since this flag is set to true at Route level,
      "server" header configuration will be enabled for this Route with
      respective "errorCodeSeriesId" as E2
      errorCodeSeriesId: E2 # This attribute will need to be defined if
      "server" header configuration is enabled at Route level.
- id: backend_ms2_route
  uri: https://backend-ms2:8550/
  path: /ms2/**
  order: 2
  metadata:
    # All attributes under "serverHeaderDetails" will need to be
    configured only if "serverHeaderConfigMode" is set as "HELM" and Route
    level configuration is required. If not defined, Global configurations
    will be used
    serverHeaderDetails:
      enabled: false     # Since this flag is set to false at Route level,
      "server" header configuration will be disabled for this Route altogether.

```

Note

If you define server header configuration at both global and route levels, the route level configuration takes precedence over the global level configuration.

For more information, see the "Server Header Configurations" section in the *Oracle Communications Cloud Native Core, Binding Support Function Installation and Upgrade Guide*.

- **Configure using CNC Console:** This feature can be configured using **NF Communication Profile** and **PCF Communication Profile** under the **Audit** section of the **Management Service** page.

3.27 Overload Control

Overload means when 100% of the planned capacity is exhausted. It can be due to uneven distribution of traffic towards a given policy service instance, network fluctuations leading to traffic bursts or unexpected high traffic volume at any given point of time.

During overload conditions, the service response times may grow to unacceptable levels, and exhaustion of resources can result in downtime or services exhibiting unexpected behavior. Overload management is a critical requirement for any telecom node, server, and service to

protect against downtime and ensure serviceability during extreme overload conditions. Thus, overload management aims to prevent service performance from degrading in an uncontrolled manner under heavy loads. When BSF service starts approaching its saturation or planned limit, response times typically grow high and throughput may degrade substantially. Under such conditions, it is desirable to shed load based on the user's configuration, instead of causing all messages and signaling flows to experience unacceptable response times, failures, or downtime.

BSF allows to configure a percentage of messages to be rejected. That is, messages are discarded based on configured percentage. This enables system's overload and congestion control to manage gauge system's load with better accuracy. Also, it allows the user to provide less rejections instead of providing 100% rejections.

Percentage of message rejections for each load level is configurable. Also, the rejection percentage for each message priority can be configured.

For example, if the discard value for CCR-I messages is 50% , when the system is under load, 50% of the messages are accepted and 50% of them are rejected.

Note

All CCR-Ts are accepted.

Enable

To enable the overload control functionality, set value for the following parameter to true in the `custom-values.yaml` file for BSF:

- **perf-info.overloadManager.enabled**

Then, configure the values for the following parameters in the `custom-values.yaml` file:

```
perf-info:
  envMySQLDatabase: ''
  overloadManager:
    enabled: false
    ingressGatewaySvcName: occnp-ingress-gateway
    ingressGatewayPort: *svcIngressGatewayHttp
    # nfType is used to query configuration from common cfg server
    nfType: BSF
    # diam Gateway overload management feature configurations
    diamGWPort: *svcDiamGatewayHttp
```

For more information about setting the parameter values, see *Overload Manager Configurations* in *Oracle Communications Cloud Native Core, Binding Support Function Installation, Upgrade, and Fault Recovery Guide*.

Configure

You configure the overload control feature either using CNC Console, or through REST API.

- **Configure using REST API:** BSF provides `overloadLevelThreshold` and `overloadLevelThresholdProfiles` API end points to configure overload control feature. You can perform the POST, PUT, or GET operations to configure the feature. For more information about REST API configuration, see *Overload Level Threshold Overload Level*

Threshold sections in *Oracle Communications Cloud Native Core, Binding Support Function REST API Specification Guide*.

- **Configure using CNC Console:** Using CNC Console, you can configure the threshold values based on profiles. For more information, see [Overload Control Threshold](#).

Recommended Overload Threshold Values - BSF Management Services

This section describes the recommended default overload threshold values for BSF Management Services. To calculate threshold values, you must consider the resource values for microservices. The following table lists the default resource values for BSF Management Services:

Table 3-30 Default Resource Values

Resources	Values
CPU (Limits)	4
CPU (Requests)	3
Maximum CPU Usage	1.6
Maximum CPU Usage (%)	61
Maximum Replicas	8
Maximum TPS	1500
Maximum TPS (all replicas)	12000
Worst RTT (assumed)	250 ms
Maximum Pending Transactions	2100

Based on the values in the aforementioned table, you can calculate the onset and abatement values for load levels - L1, L2, and L3, as shown in the following table:

Table 3-31 Formulas to Calculate Default Overload Threshold Levels

Load Level	CPU (%)	Pending Message Count (Absolute Value)	Failure Count (Absolute Value)
L1 - Onset	$80\% * C$	$60\% * P$	$05\% * T$
L1 - Abatement	$75\% * C$	$50\% * P$	$03\% * T$
L2 - Onset	$90\% * C$	$75\% * P$	$10\% * T$
L2 - Abatement	$85\% * C$	$70\% * P$	$08\% * T$
L3 - Onset	$95\% * C$	$90\% * P$	$15\% * T$
L3 - Abatement	$91\% * C$	$85\% * P$	$12\% * T$

Abatement value is the lower range where as the onset value is higher range for that particular level.

Note:

C denotes the maximum CPU utilization per pod of core-service

P denotes the maximum pending transaction size value based on worst RTT and max TPS, that is, $25K/s * 2000ms = 50000$.

T denotes the maximum TPS of a given service.

Note

You can configure the memory for each of these services.

The following table lists the default overload threshold values for BSF Management Services:

Note

These are the recommended values. It can be modified as per the customer requirements.

Table 3-32 Default Overload Threshold Values - BSF Management Services

Load Level	CPU (%)	Pending Message Count	Failure Count
L1 - Onset	44	1260	420
L1 - Abatement	41	1050	252
L2 - Onset	49	1575	840
L2 - Abatement	47	1470	672
L3 - Onset	52	1890	1260
L3 - Abatement	50	1785	1008

Observe

BSF provides the following metrics specific to Overload Control feature:

- service_resource_stress
- service_resource_overload_level
- load_level
- system_overload_threshold_config_mode
- active_overload_threshold_fetch_failed

For more information, see [BSF Metrics](#) section.

Alerts

- BSF provides the following alerts for overload control feature on SBI interface:
 - ServiceOverloaded - This alert is raised whenever a given service is in overload state - L1, L2, and L3.
 - ServiceResourceOverLoaded - This alert is raised when a given service is in overload state - L1, L2, or L3 due to resource types such as memory, CPU, pending count, and failure count.
- BSF provides PERF_INFO_ACTIVE_OVERLOAD_THRESHOLD_FETCH_FAILED alert for overload control threshold configuration. This alert is raised when the service is unable to fetch the current active overload threshold data.

Maintain

Error logs are generated when the system is overloaded and the actions taken to bring the system back to normal. Warning logs are generated to indicate the change in load level.

3.27.1 Overload Control - Diameter

For Diameter Gateway, BSF provides the following means for overload management:

- Pre-defined threshold load levels.
- Tracks number of pending and failure messages from Diameter Gateway.
- Tracks CPU and memory usage of Diameter Gateway.
- Enforce load shredding during various overload levels based on priority and percentage discard value for each priority. The priority and percentage discard value are configurable.

Configure

To configure the threshold values, discard priority, and error codes for the defined overload control levels, you may use CNC Console as well as REST API.

Note

Currently, threshold values can be configured using REST API only.

- **Configure using CNC Console:** Perform the feature configurations on the Load Shedding Profiles and Message Priority Profiles page. For more information about the configurations, see [Load Shedding Profiles](#).
- **Configure using REST API:** BSF provides the following REST API for configuring Overload Control feature on Diameter Gateway:
Load Shedding Profiles: {apiRoot}/oc-bsf-configuration/v1/diameter/loadsheddingprofiles
Message Priority Profiles: {apiRoot}/oc-bsf-configuration/v1/diameter/messagepriorityprofiles

You can perform the POST, PUT, or GET operations to configure the feature. For more information about REST API configuration, see *Oracle Communications Cloud Native Core, Binding Support Function REST API Specification Guide*.

The following are the recommended configurations for load shedding profile and message priority profile respectively for overload control:

```
{
  "name": "default_overload_control_load_shedding_profile",
  "type": "Overload Control",
  "overloadLoadSheddingRules": [{
    "level": "L1",
    "discardPriority": 13,
    "ansWithResultCode": "DIAMETER_T00_BUSY"
  }, {
    "level": "L2",
    "discardPriority": 11,
    "ansWithResultCode": "DIAMETER_T00_BUSY"
  }, {
    "level": "L3",
    "discardPriority": 6,
    "ansWithResultCode": "DIAMETER_T00_BUSY"
  }
}
```

```

    }}
  }

  {
    "name": "default_msg_priority_profile",
    "priorityRules": [{
      "ruleName": "Rx_AAR_I",
      "messagePriority": 13,
      "rulePriority": 1,
      "enableDRMPPriority": true,
      "conditions": {
        "application": "Rx",
        "message": "AAR",
        "preDefinedAVPConditions": [{
          "conditionName": "Rx-Request-Type",
          "conditionRxRTValue": ["INITIAL_REQUEST"]
        }]
      }
    }, {
      "ruleName": "Rx_STR",
      "messagePriority": 7,
      "rulePriority": 1,
      "enableDRMPPriority": true,
      "conditions": {
        "application": "Rx",
        "message": "STR",
        "preDefinedAVPConditions": []
      }
    }, {
      "ruleName": "Rx_AAR_U",
      "messagePriority": 11,
      "rulePriority": 1,
      "enableDRMPPriority": true,
      "conditions": {
        "application": "Rx",
        "message": "AAR",
        "preDefinedAVPConditions": [{
          "conditionName": "Rx-Request-Type",
          "conditionRxRTValue": ["UPDATE_REQUEST"]
        }]
      }
    }
  ]
}

```

Observe

BSF provides the following metric specific to Overload Control feature:

- `diam_overload_message_reject_total`

For more information, see [BSF Metrics](#) section.

3.27.2 Overload Control - SBI

For HTTP signaling, BSF provides the following means for overload management:

- Pre-defined threshold levels
- Tracks number of pending messages for BSF service
- Tracks number of failed responses (configurable as error code) generated BSF service
- Tracks CPU and memory usage of BSF services
- Determines the overload level of the system using data collected from all the above mentioned points against planned threshold levels.
- Enforce load shedding at various overload levels

Configure

To configure the discard policies, Discard Policy mapping, and Error Code Profiles for overload control, you may use CNC Console as well as REST API.

Note

Currently, threshold values can be configured using REST API only.

- **Configure using CNC Console:** Perform the feature configurations on the Discard Policy Mapping, Discard Policy, and Error Code Profiles pages. For more information about the configurations, see [Overload and Congestion Control Configurations](#).
- **Configure using REST API:** BSF provides the following REST APIs for configuring overload control feature on SBI interface:
 - **OC Policy Mapping:** {apiRoot}/BSF/nf-common-component/v1/igw/ocpolicymapping
 - **OC Discard Policies:** {apiRoot}/BSF/nf-common-component/v1/igw/ocdiscardpolicies
 - **Error Code Profiles:** {apiRoot}/BSF/nf-common-component/v1/igw/errorcodeprofiles

You can perform the GET, PUT, or PATCH operations to configure the feature. For more information about REST API configuration, see *Oracle Communications Cloud Native Core, Binding Support Function REST API Guide*.

The following are the recommended configurations for default message priority values for overload control:

Table 3-33 Default Message Priority Values

Message Type	Priority
bsf_management_register	24
bsf_management_deregister	18
bsf_management_discovery	24

3.28 Rate Limiting - SBI

With the support for rate limiting, Ingress Gateway screens all configured routes and their respective rate limit configurations. Within the configured sampling period (configurable using CNC Console and REST APIs), Ingress Gateway calculates the rate for the required route along with the HTTP method. For BSF, the following routes and HTTP methods are supported:

- BSF Management Register - POST method
- BSF Management Deregister - DELETE method

- BSF Management Discovery - GET method

Then, it notifies the route level rate limiter with the calculated rate at the end of the sampling period. If the feature is enabled, any request with the sbi-priority header value greater than the configured value is discarded, and Ingress Gateway returns the error response with configured `errorCode`.

Note

Ingress Gateway determines the number of messages being dropped or rejected in the current sampling period based on extra unrejected messages received in the previous sampling period.

Managing Rate Limiting

Enable

Perform the following configurations to enable the rate limiting feature at Ingress Gateway:

- **CNC Console:** By default, this feature is disabled. To enable the rate limiting feature using CNC Console, set the **Enable Rate Limiting** parameter to `true` on the Rate Limiting Policy page.
- **REST API:** By default, this feature is disabled. To enable the rate limiting feature using REST API, set the `enabled` parameter to `true` in the following resource URI:

Define rate limit: `{apiRoot}/BSF/nf-common-component/v1/igw/routelevelratelimiting`

Define rate limit at route level: `{apiRoot}/BSF/nf-common-component/v1/igw/routesconfiguration`

For more information, see *Rate Limiting at Ingress Gateway* section in *Oracle Communications Cloud Native Core, Binding Support Function REST API Guide*.

Configure

To configure the rate limiting policy, route level mapping, and error code profiles for rate limiting, you may use CNC Console as well as REST API.

- **Configure using CNC Console:** Perform the feature configurations on the Rate Limiting Policy, Route Level Mapping, and Error Code Profiles pages. For more information about the configurations, see [Overload and Congestion Control Configurations](#).

- **Configure using REST API:** BSF provides the following REST API for configuring Overload Control feature on SBI interface:

Define error code profiles: `{apiRoot}/BSF/nf-common-component/v1/igw/errorcodeprofiles`

Define rate limit: `{apiRoot}/BSF/nf-common-component/v1/igw/routelevelratelimiting`

Define rate limit at route level: `{apiRoot}/BSF/nf-common-component/v1/igw/routesconfiguration`

You can perform the GET, PUT, or PATCH operations to configure the feature. For more information about REST API configuration, see *Oracle Communications Cloud Native Core, Binding Support Function REST API Specification Guide*.

3.29 Pod Protection at Ingress Gateway

This section describes how to protect the Ingress Gateway pods when they are overloaded with numerous incoming requests.

The Ingress Gateway pods are not protected against any incoming traffic congestion. As a result, the pods are overloaded and congested. This impacts system latency and performance. It also leads to stability issues due to uneven distribution of connections and traffic on Ingress Gateway pods. As a front end microservice for HTTP traffic, it is important for Ingress Gateway to have pod protection implemented.

To configure pod protection on Ingress Gateway, you can define threshold limit for DoC and Congested state through REST:

Table 3-34 Configuring Threshold Limit

Level	Resource
DoC	• CPU • Memory • Pending Message
Congested	• CPU • Memory • Pending Message

Configure

You need to perform the following configurations for pod protection feature:

- **Configure using REST API:** BSF provides the following REST API:
{apiRoot}/BSF/nf-common-component/v1/igw/podprotection

You can perform the GET, PATCH, or PUT operation to configure the feature. For more information about REST API configuration, see *Oracle Communications Binding Support Function REST API Specification Guide*.

3.30 Service Mesh for Intra-NF Communication

BSF leverages the Istio or Envoy service mesh (Aspen Service Mesh) for all internal and external communications. The service mesh integration provides inter-NF communication and allows API gateway co-working with service mesh. The service mesh integration supports the services by deploying a sidecar proxy in the environment to intercept all network communications between microservices.

The Aspen Service Mesh (ASM) configurations are classified into:

- **Control Plane:** It involves adding labels or annotations to inject sidecar.
- **Data Plane:** It helps in traffic management such as handling NF call flows by adding Service Entries (SE), Destination Rules (DR), Envoy Filters (EF), and other resource changes such as apiVersion change between versions. This is done manually depending on each NF requirement and ASM deployment.

Managing Service mesh for intra-NF Communication

Enable

To enable Aspen Service Mesh, configure the following parameters under `nrf-client-nfdiscovery`, `ingress-gateway`, `egress-gateway`, and `alternate-route` sections in the custom values file for BSF:

- `serviceMeshCheck`
- `istioSidecarQuitUrl`
- `istioSidecarReadyUrl`

For more information on enabling the parameter value, see "Aspen Service Mesh Configurations" in *Cloud Native Binding Support Function Installation and Upgrade Guide*.

Configure

The Aspen Service Mesh (ASM) configurations are classified into:

- **Control Plane:** For information on configuring the parameter value, see "Aspen Service Mesh Configurations" section in *Oracle Communications Cloud Native Binding Service Function Installation Guide*.
- **Data Plane:** For information about Data plane configurations, see "Aspen Service Mesh Data Plane Configurations" in *Cloud Native Binding Support Function Installation and Upgrade Guide*.

When ASM is enabled, the ASM sidecar can modify the Server header in an Audit Notification response received by BSF. To preserve the header value unchanged, configure an `EnvoyFilter` with `server_header_transformation: PASS_THROUGH` for the BSF workload.

Use the following `EnvoyFilter` configuration. Update the namespace value to the namespace of the target BSF deployment before applying it:

```
apiVersion: networking.istio.io/v1alpha3
kind: EnvoyFilter
metadata:
  name: preserve-server-header
  namespace: <BSF_NAMESPACE>
spec:
  workloadSelector:
    labels:
      application: ocbsf
  configPatches:
    - applyTo: NETWORK_FILTER
      match:
        listener:
          filterChain:
            filter:
              name: "envoy.filters.network.http_connection_manager"
      patch:
        operation: MERGE
        value:
          name: "envoy.filters.network.http_connection_manager"
          typed_config:
            "@type": type.googleapis.com/
envoy.extensions.filters.network.http_connection_manager.v3.HttpConnectionManager
      server_header_transformation: PASS_THROUGH
```

3.31 Turning off AccessToken signature Validation

OAuth access tokens grant an NF service consumer access to the services of an NF producer of a particular NFType, for a specific period. With this feature, BSF can turn off AccessToken signature validation at the application layer. For example when Aspen service mesh is integrated with BSF, the service mesh can perform the AccessToken validation. In such cases, operator may want BSF to skip validating the AccessToken signature. BSF checks audience and scope fields only, and sends 403 - Forbidden response code when any of the values do not match. In addition, when BSF receives a request without the AccessToken, it sends a 401 - Unauthorized response code. To turn off the AccessToken signature validation at BSF application, the user must perform configurations as described in the *Oracle Communications Cloud Native Core Binding Support Function Installation Guide*.

3.32 Automated Test Suite Support

BSF provides Automated Test Suite (ATS) for validating the functionalities. ATS allows you to run BSF test cases using an automated testing tool and compares the actual results with the expected or predicted results. The ATS requires no user intervention. For more information, see *Oracle Communications Cloud Native Core Automated Testing Suite Guide*.

3.33 SBI Error Codes

Oracle Communications Cloud Native Core Binding Support Function (BSF) can handle Protocol or Application errors and a few other additional defined errors for various scenarios. When BSF encounters an error in processing a request, it sends error codes in the response message to the request. With this enhanced functionality, BSF allows users to configure error codes by adding customized values, for a defined condition, for the following fields:

- Error Description
- HTTP Status Code
- Application Error Code

Configure

To configure error codes for BSF, users can use any of the following ways:

- **CNC Console:** Perform the configurations on the SBI Error Codes page. For more information, see [SBI Error Codes Configurations](#).
- **REST API:** Perform the configurations using POST, PUT, or GET operations. For more information about REST API configuration, see *Oracle Communications Cloud Native Core, Binding Support Function REST Specification Document*.

Observe

When BSF generates error codes, it also increments the associated metric, `ocbsf_ingress_response_total`. To support the SBI error codes feature, a new dimension, `application_error_code` is added to the metrics to enable the user to easily identify the failed message or procedure and the associated error code. Metrics also include the consumer (IP or FQDN) to whom the error would be sent. For more details, see [BSF Metrics](#).

On generating error codes, logs are updated as well.

The following is a sample error log when optional parameter in binding data is invalid:

```
{
  "instant": {
    "epochSecond": 1628086282,
    "nanoOfSecond": 772064777
  },
  "thread": "XNIO-1 task-1",
  "level": "DEBUG",
  "loggerName": "ocpm.bsf.api.management.metrics.BsfMetrics",
  "message": "Pegging the Ingress Response Metric for Operation Type :
register, Response Code ProblemDetails [type=about:blank, title=Optional
parameter in binding data is invalid, status=411,
cause=OptionalInformationElementIsIncorrect, instance=http://b-ocbsf-ingress-
gateway.b/nbsf-management/v1/pcfBindings, details=Parameters in the request
is not correct , invalidParams=[InvalidParam {\n    param: pcfDiamRealm\n
reason: must match \"^([A-Za-z0-9]+([-A-Za-z0-9]+)\\.)+[a-z]{2,}$\"\\n}]]",
  "endOfBatch": false,
  "loggerFqn": "org.apache.logging.slf4j.Log4jLogger",
  "threadId": 31,
  "threadPriority": 5,
  "messageTimestamp": "2021-08-04T14:11:22.772+0000"
}
```

Result: As mentioned in the message field of the sample log, the metric `ocbsf_ingress_response_total` is pegged.

3.34 Handling Stale Session in BSF

There are chances of certain transactions failing due to situations like congestion or overload. To manage such failed transactions, BSF provides a database audit mechanism that monitors the stale records and cleans them up to not exhaust the database memory. The audit mechanism also notifies the microservice about the stale records so that the service can trigger signaling messages.

A session binding on BSF is considered as stale when the PCF binding is successfully registered on BSF but has no corresponding session on PCF. When you enable stale session handling feature, the Audit service detects stale sessions automatically at regular intervals.

During the audit, if the audit service finds records with a binding age greater than the configured value (default value is 3600 minutes), it marks such records as suspected stale and initiates a request to notify the BSF management service. Depending on the feature configurations, BSF may query PCF to confirm if the `pcfBinding` records are stale. If PCF confirms the record as stale, then BSF removes it from its local database. However, if PCF sends 2xx in the response, then BSF updates the last access time for the associated `pcfBinding` record.


If the audit service finds a `pcfBinding` record with binding age greater than the maximum binding age (default value is 7200 minutes), it notifies the BSF Management service. On receiving the notification, BSF deletes the specified record from its database.

For georedundant BSF deployments, binding records are replicated to other redundant sites after the audit cycle is complete. Logs are published when audit service detects and removes stale sessions.

Feature Design

The following diagram provides a high level design for stale session handling through session state Audit service:

Figure 3-34 High Level Design for Stale Session Handling using Audit Service



As shown in the above diagram:

1. BSF registers as a consumer with the Audit service and starts auditing the service database.
2. When a stale record is detected, the Audit service takes any of the following actions as requested by the BSF during the registration:
 - Deletes the stale records from the database.
 - Sends a notification to the service about the stale records.
 - Deletes the stale records from the database and notify the service.
3. The Audit service implements a minimum wait time between consecutive audits and the consecutive notifications for the same record.

Note


- The Audit service stores the service registration details in the DB to be recovered after the pod restart or upgrade.
- The audit service supports the local time zone.

Call Flow

Timeline View of Audit Service

The following figure shows a time line view of the Audit service for a record:

Figure 3-35 Timeline view of Audit Service



BSF identifies and deletes stale sessions in the following scenarios:

Query PCF to confirm a record as stale

Whenever Audit Service in BSF identifies a session which is soon to go stale (minimum Time-to-Live (minTTL) reached), Audit sends an audit notification request towards BSF Management service.

If **Query to PCF** is enabled, BSF sends a query to PCF to verify the corresponding session in PCF.

If PCF answers with a 404 (NOT FOUND) then BSF deletes the session.


To query PCF for confirming stale pcfBinding records, BSF supports receiving vendor-specific-attribute in the binding register request. So, when PCF sends a register request to BSF for binding creation, it includes the notification URL in the vendor-specific-attribute, as shown below:

```
"vendorSpecific-000111": {
  "version": 1,
  "notificationUrl": "<PCF Notification Url>",
  "createBindingTime": "<Timestamp>",
}
```

BSF, in turn, uses this notification URI to send audit notification towards Binding service on PCF. When PCF receives the notification, it checks whether any binding record exists by sending a query to the Query service using contextId (SmPolicyAssociationId). Based on the findings, PCF sends 2xx response if session exists and 404 if the session does not exist.

A sample call flow for this scenario is shown below:

Figure 3-36 Query PCF to confirm a record as stale




1. Audit service identifies sessions which are soon to become stale (with minTTL reached) and sends a notification to BSF Management Service.
2. If **Query to PCF** is enabled, BSF Management Service sends a query to PCF through Egress Gateway to check if the sessions exist in PCF.
3. If the queried sessions are not present in PCF, PCF responds with 404 NOT FOUND.
4. BSF considers the session as stale and deletes them.
5. BSF Management Service responds to Audit service with 200 OK.

Max TTL and minimum audit attempts reached

Whenever Audit Service in BSF identifies a session for which the Max Time-to-Live (maxTTL) is reached and has reached its minimum Audit Attempts too, Audit sends an audit notification requests towards BSF Management Service with its body value **maxTtlReached** as true. Whenever this request is received at BSF Management service, BSF deletes the session.

A sample call flow for this scenario is shown below:

Figure 3-37 Max TTL and minimum audit attempts reached


1. Audit Service identifies the sessions that are soon to become stale and sends a notification to BSF Management Service with `maxTTLReached=false`.
2. BSF Management Service sends a query to PCF through Egress Gateway to verify the corresponding sessions in PCF.
3. If PCF responds with `503 SERVICE UNAVAILABLE`, BSF Management Service increments the audit attempts for the audited session and responds to Audit Service with `200 OK`.
4. After exhausting the minimum audit attempts, when the `maxTTL` is reached (in the given example, `minAuditAttempts=10` and `maxTTL=5 mins`), Audit Service sends a notification to BSF Management Service with `maxTTLReached=true`.
5. BSF Management Service deletes the corresponding PCF Binding for the session from its database and responds to Audit Service with `200 OK`.

Stale Session Notification from Diameter Gateway

Whenever BSF sends an AAR (can be either AAR-I or AAR-U) to Diameter Gateway to check for the existence of a session in the bound PCF, and PCF responds with an error, such as `DIAMETER_UNABLE_TO_COMPLY`, Diameter Gateway sends a delete request towards BSF through the `/audit/notify` API.

Whenever this request is received at BSF Management Service, BSF deletes the session.

A sample call flow for this scenario is shown below:

Figure 3-38 Stale Session Notification from Diameter Gateway

Enable

After performing the required Helm configurations, you can enable the handling stale sessions in BSF using the CNC Console or REST API.

- **Enable using Helm:** To enable this feature during BSF deployment, set the value of `auditServiceEnable` parameter as `true` in the `ocbsf-22.2.0-custom-values.yaml` file. Once the Audit service is enabled, set appropriate values for the following parameters:

- `global.servicePorts.auditServiceHttp`
- `global.containerPorts.auditServiceHttp`
- `audit-service.envMySQLDatabase`

For more information on how to customize these parameters, see *Oracle Communications Cloud Native Core Binding Support Function Installation and Upgrade Guide*.

- **Enable using CNC Console:** To enable this feature, on the Management Service page, under the **Audit** group, set the value of **Enabled** field as `true`. For more information about enabling the feature through CNC Console, see [Service Configurations](#).
- **Enable using REST API:** Set the `audit.enable` parameter value to `true` in the Management Service configuration API. For more information about enabling the feature through REST API, see the "Management Service" section in *Oracle Communications Core Binding Support Function REST Specification Guide*.

Configure

You can configure the Stale Session Handling functionality for BSF using the CNC Console or REST API.

- **Configure using CNC Console:** Perform the feature configurations on the **Management Service** page. For more information, see [Service Configurations](#).
- **Configure using REST API:** BSF provides the following REST API for Stale Session Handling configuration:
API: `{apiRoot}/oc-bsf-configuration/v1/services/managementservice`

You can perform the GET and PUT operations to configure this feature. For more information about REST API configuration, see the "Management Service" section in *Oracle Communications Core Binding Support Function REST Specification Guide*.

Observe

BSF provides the following metrics for stale session handling feature:

- ocbsf_audit_notif_request_count_total
- ocbsf_audit_notif_response_count_total
- ocbsf_audit_delete_records_count_total
- ocbsf_diamgw_notification_request_count_total
- ocbsf_diamgw_notification_response_count_total
- ocbsf_query_request_count_total
- ocbsf_query_response_count_total
- ocbsf_bindingQuery_request_total
- ocbsf_bindingQuery_response_total
- ocbsf_bindingDelete_request_total
- ocbsf_bindingDelete_response_total

For more information, see the **BSF Management service** and **Audit service** metrics in the [BSF Metrics](#) section.

Maintain

BSF generates logs when audit service detects and removes stale sessions automatically. In addition, logs are printed for query requests towards PCF and the associated responses.

Stale Session Audit Retry

An audit query or notification request from BSF to PCF may fail due to reasons such as:

- Internal send failures
- SCP being unable to route a request to PCF
- PCF databases being out of sync, returning a 404 error

This may lead to unreliable or incorrect results regarding the existence of a PDU session. To prevent this, BSF allows for session retries within the same audit cycle. It uses Server Headers to identify the origin of an audit request failure and take appropriate action. Using NRF Client's On-Demand Discovery, it selects an alternate route and retries the audit request against a different PCF or SCP. The parameters **Minimum required 404 responses** and **Expect 404 from different PCF** can be set to decide the number of 404 responses, and if they must be from different PCF instances for a binding to be considered as stale.


BSF continues retrying till these conditions are met or until enough evidence is collected to mark a binding as stale. If all retries fail, it stops for that transaction and the binding is handled according to the audit decision logic.

Call Flows

Note


The following diagrams are part of a single call flow which has been divided into separate sections.

Figure 3-39 BSF Audit Notification Flow with EGW Rerouting and Alternate NF Discovery



1. BSF Management sends the POST audit request notification to Egress Gateway with the `oc-alternate-route=0` header.
2. Egress Gateway forwards the request to SCP1 and sets the `3gpp-Sbi-Target-apiRoot` header to PCF1.
3. SCP1 forwards the `/npcf-custom-binding` request to PCF1.
4. PCF1 returns a non-404 error response to SCP1.
5. SCP1 does not retry the request and forwards the non-404 error response back to the Egress Gateway.
6. The Egress Gateway detects that SbiRouting retry is enabled and selects SCP2 as the alternate SCP for retry.
7. The Egress Gateway sends a POST audit retry request to SCP2.
8. SCP2 forwards the `/npcf-custom-binding` request to PCF1 again.
9. PCF1 returns a 404 error response to SCP2.
10. SCP2 does not retry the request and forwards the 404 error response back to the Egress Gateway.
11. The Egress Gateway determines that all rerouting attempts are complete and forwards the 404 error response to BSF Management.
12. BSF Management checks whether retry handling is configured for 404 errors through alternate NF discovery.
13. BSF Management initiates NRF discovery using `SetId` and `NFType` to locate an alternate NF.
14. If BSF receives a non-404 error response or timeout without the `oc-egw-route` header during the initial request, BSF performs an internal retry based on the configured retry policy.

Figure 3-40 On-Demand NF Discovery Flow with NRF Client Caching




Scenario 1: On-demand NF discovery with NRF Client caching enabled

1. BSF Management requests NF discovery from the NRF Client.
2. The NRF Client checks its local cache for valid NF profiles.
3. If valid cached NF profiles are available, the NRF Client returns the cached NF discovery response to BSF Management.
4. If the cache entry is expired or unavailable, the NRF Client sends a fresh NF discovery request to the NRF.
5. The NRF returns the updated NF profiles to the NRF Client.
6. The NRF Client updates its local cache with the latest NF profiles.
7. The NRF Client returns the updated NF discovery response to BSF Management.

Scenario 2: On-demand NF discovery with NRF Client caching disabled or ForceDiscovery enabled

1. BSF Management sends an NRF discovery request to NRF Client.
2. The NRF Client initiates on-demand discovery using SetId and NFType.
3. The NRF Client forwards the NRF discovery request to the NRF.
4. NRF performs on-demand discovery using SetId and NFType.
5. NRF returns the NRF discovery response to the NRF Client.
6. NRF Client forwards the NF profiles received from the NRF to BSF Management.
7. BSF Management receives the NRF discovery response containing NF profiles from the NRF.

Figure 3-41 BSF Audit Retry Flow with Alternate NF Selection and 404 Error Handling**Scenario 1: Alternate NF exists in the same NF set**

1. BSF Management sends a POST audit retry request to the Egress Gateway for PCF2 with the oc-alternate-route=1 header.
2. Egress Gateway forwards the POST audit retry request to SCP2.
3. SCP2 forwards the /npcf-custom-binding request to PCF2.
4. PCF2 returns a 404 error response to SCP2.
5. SCP2 does not retry the request and forwards the 404 error response to the Egress Gateway.
6. Egress Gateway forwards the 404 error response to BSF Management because rerouting is not enabled for 404 errors.
7. BSF Management determines that the retry attempt is complete and that the configured total404Count limit has been reached.
8. BSF Management treats the 404 error as the final response and deletes the PCF binding record from the database.

Scenario 2: Alternate NF does not exist in the same NF set

1. BSF Management sends a POST audit retry request to the Egress Gateway for PCF2 with the oc-alternate-route=1 header.
2. Egress Gateway forwards the POST audit retry request to SCP2.
3. SCP2 forwards the /npcf-custom-binding request to PCF1.
4. PCF1 returns a 404 error response to SCP2.
5. SCP2 does not retry the request and forwards the 404 error response to the Egress Gateway.
6. Egress Gateway forwards the 404 error response to BSF Management as rerouting is not enabled for 404 errors.

Enable

To enable BSF to discover PCF instances, the PCF must register with the NRF and include BSF in the `allowedNFTypes` parameter of its NF Profile. This allows BSF, as a consumer NF, to discover the registered PCF instances.

A template NF Profile is provided and can be modified as required. The following attributes must be configured:

- NF FQDN
- Service FQDN and Port
- Service IP Endpoints
- Service API versions and schemes (HTTP/HTTPS)

Note

For the BSF audit notification feature, the PCF must register the `Npcf_SMPolicyControl` service with the NRF.

You can enable Stale Session Audit Retry in BSF using the CNC Console or REST API.

- **Enable using CNC Console:** To enable this feature, under Management Service page, under the **Audit** section, set the values of **Enabled** and **Query to PCF** fields as true. Additionally, set relevant values for the fields **Minimum required 404 responses**, **Expect 404 from Different PCF** and **PCF Communication Profile**. For more information about enabling the feature through CNC Console, see [Service Configurations](#).
- **Enable using REST API:** In the Management Service configuration API, set the values of parameters `audit.enable` and `queryToPcf` to **true**. Additionally, set relevant values for the parameters `total404Count`, `expect404FromDifferentPCF` and `pcfNfCommProfile`.

Configure

You can configure the Session Retry functionality for BSF using the CNC Console or REST API.

- **Configure using CNC Console:** Perform the feature configurations on the following pages:
 - Retry Profile
 - NF Communication Profile
 - NRF Agent Service
 - Management Service

For more information, see [Service Configurations](#).

- **Configure using REST API:** BSF provides the following REST APIs for Stale Session Audit Retry configuration:
 - API: `{apiRoot}/oc-bsf-configuration/v1/services/retryprofiles`
 - API: `{apiRoot}/oc-bsf-configuration/v1/services/nfcommprofiles`
 - API: `{apiRoot}/oc-bsf-configuration/v1/services/nrfAgentServiceConfig`
 - API: `{apiRoot}/oc-bsf-configuration/v1/services/managementservice`

Note

- If the BSF attempts to route the request to PCF1 and receives an error, it retries using the alternate route to PCF2, which returns a successful response. However, if the 3gpp-Sbi-Target-apiRoot header is not included in the response, the PCF notification URL will not be updated in the PcfBinding record. As a result, during the subsequent audit, the BSF will again make the initial attempt to PCF1.
- In the case of an alternate routing attempt, the 3gpp-Sbi-Routing-Binding header will not be updated with the selected NF profile's NF Instance ID. Instead, the header will retain the same value that was used in the original request.
- Enabling the Retry and Reroute feature requires the nfdiscovery application to be deployed as part of the BSF deployment to support NRF on-demand discovery calls from NRF-Client to NRF. This will increase the overall resource consumption of the BSF deployment, and the application resource's requirement may need to be updated accordingly. Additionally, since retry is not implemented between the BSF Management service and NRF-Client's NFDISCOVERY interface, any failure during on-demand discovery will cause the retry procedure for the audit request to fail, and the response from the initial attempt will be treated as the final response.
- The feature of SCP selection depends on Monitoring the Availability of SCP using HTTP2 OPTIONS (Health check API).
- When this feature is enabled, SCP retry/reselection should be disabled. SCP connected to the PCF has the capability to re-select Target NFs based on NFSet. It cannot re-select another SCP and forward the message to an alternate SCP. It can only re-select another producer from the NF Set.
- By default, the service name in the NF On-Demand Discovery (ODD) request is populated from the binding header. If the binding header does not contain the service name, the default service name is used. Since this discovery parameter is optional, it can be disabled by setting the advanced setting `ENABLE_DISCOVERYPARAM_SERVICE_NAMES` to false.
- When enabled in the NfCommProfile, the callback header is sent with the default value `Nbsf_Management_myNotification`. This value can be overridden by configuring the advanced setting `SBI_CALLBACK_HEADER_VALUE` with the desired callback header value.

Observe

BSF provides the following metrics for session retry for stale session audit:

- `ocbsf_ar_request_total`
- `ocbsf_ar_response_total`
- `ocbsf_ar_failure_total`
- `ocpm_nfDiscovery_request_total`
- `ocpm_nfDiscovery_response_total`
- `ocbsf_query_request_count_total`
- `ocbsf_query_response_count_total`

3.34.1 Enhanced Logging of BSF Stale Binding Deletion

BSF generates logs when a stale session is deleted. The log includes:

- Reason for deleting the stale session:
 - 404 Response received from PCF
 - Max TTL Reached
 - Stale Session Notification received from Diameter Gateway
- Policy binding session details such as:
 - bindingId
 - SUPI
 - GPSI
 - DNN
 - UE IP address:
 - * IPv4
 - * IP Domain
 - * IPv6 Prefix
 - Policy FQDN
 - Policy Instance ID
 - Policy Set ID
 - SNSSAI
 - Policy Diameter Host
 - Policy Diameter Realm
 - Created Date Time
 - Last Audited Date Time
 - Audit Attempts

For example:

```
{
  bindingId: 671df805-8a48-4711-ba23-fdbf198b3fa7,
  supi: imsi-311480611009999,
  gpsi: msisdn-19544398580,
  dnn: abcinternet,
  ipv4Addr: 140.0.0.1,
  ipv6Prefix: 2607:f160:3e:6d:0:0:0:0/64,
  pcFqdn: oc-diam-gateway,
  pcId: fe7d992b-0541-4c7d-ab84-c6d70b1b0420,
  snssai: 1-000007,
  pcDiamHost: oc-diam-gateway,
  pcDiamRealm: oracle.com,
  createdDateTime: 2025-06-11T07:25:56.701652Z,
  lastAuditedDateTime: 2025-06-11 07:25:57.0,
```

```
    auditAttempts: 0  
}
```

Note

Details with null values are not included in logs.

Managing the Enhanced Logging of BSF Stale Binding Deletion

Enable

Enhanced logging for BSF stale binding deletion can be enabled either using CNC Console or using REST API.

Enable using CNC Console:

Configure **Log Level for Stale Session Deletions** field under **Enhanced Logging Configuration** section on **General Settings** page for **General Configurations** in CNC Console for BSF.

The **Log Level for Stale Session Deletions** field specifies the log level for stale session deletions.

Figure 3-42 Log Level for Stale Session Deletions

Log Level for Stale Session Deletions:	Log Level for Stale Session Deletions
Enable UE Identifier Information:	INFO
Enable UE Identifier Information in Stale Session Deletion Logs:	TRACE
	DEBUG
	INFO
	WARN
	ERROR

The available log levels for stale session deletions are:

- TRACE
- DEBUG
- INFO
- WARN
- ERROR

The logs for BSF stale binding deletion are generated at the specified level.

Note

The value of this field does not modify in any way the log level that BSF Management service has. The root log level of BSF Management service must be either the same log level defined in the selected or lower for these deletion logs to appear. For example, if the stale session deletions log level is at INFO, then BSF Management service log level must be either TRACE, DEBUG, or INFO.

For more information, see [General Settings](#).

Enable using REST API

Configure `enhancedLogging.staleSessionDeletionLogLevel` parameter for General Configurations API.

The General Configurations API can be accessed using `{apiRoot}/oc-bsf-configuration/v1/general/export`.

For more information, see *General Configurations* section in *Oracle Communications Core Binding Support Function, REST Specification Guide*.

Configure

Enhanced logging for BSF stale binding deletion can be enabled either using CNC Console or using REST API.

Using CNC Console:

Contents of the logs can be configured by configuring the following fields under **Enhanced Logging Configuration** section on **General Settings** page for **General Configurations** in CNC Console for BSF.

- **Log Level for Stale Session Deletions:** Specifies the log level for stale session deletions. By default, the log level for stale session deletions is at INFO level.
- **Enable UE Identifier Information in Stale Session Deletion Logic:** Specifies whether to enable or disable UE Identifier information in stale session deletion logging. When this field is disabled, the UE Identifier information in the logs are masked. That is, the identifier information appears like 'xxx'. By default, this field is disabled.

Sample logs when **Enable UE Identifier Information in Stale Session Deletion Logic** is disabled:

```
{
  "bindingId": "b75b4c43-0636-485f-b9e4-26fb6eea97ca",
  "supi": "XXXX",
  "gpsi": "XXXX",
  "dnn": "dnn2",
  "ipv4Addr": "XXXX",
  "pcfFqdn": "pcf-smsservice. oracle .com",
  "snssai": {
    "sst": 11,
    "sd": "abc123"
  },
  "pcfDiamHost": "pcf-smsservice. oracle .com",
  "pcfDiamRealm": "oracle .com"
}
```

Sample logs when **Enable UE Identifier Information in Stale Session Deletion Logic** is enabled:

```
{
  "bindingId": "b75b4c43-0636-485f-b9e4-26fb6eea97ca",
  "supi": "imsi-311411000000002",
  "gpsi": "msisdn-5084948009",
  "dnn": "dnn2",
}
```

```
"ipv4Addr": "10.10.10.2",  
"pcfFqdn": "pcf-smsservice.oracle.com",  
"snssai": {  
  "sst": 11,  
  "sd": "abc123"  
},  
"pcfDiamHost": "pcf-smsservice.oracle.com",  
"pcfDiamRealm": "oracle.com"  
}
```

For more information, see [General Settings](#).

Using REST API

To configure enhanced logging for BSF stale session deletions, configure enhancedLogging section in General Configurations API.

The General Configurations API can be accessed using {apiRoot}/oc-bsf-configuration/v1/general/export.

For more information, see *General Configurations* section in *Oracle Communications Core Binding Support Function, REST Specification Guide*.


3.35 Support Multiple Cluster Deployment at CNC Console

The CNC Console supports both single and multiple cluster deployments.

In a single cluster deployment, the CNC Console can manage NFs and Oracle Communications Cloud Native Core, Cloud Native Environment (CNE) common services deployed in the local Kubernetes clusters.


In a multiple instances deployment, the CNC Console can manage multiple BSF instances and CNE common services deployed within a Kubernetes cluster. For more information about single and multiple cluster deployments, see *Oracle Communications Cloud Native Core, Cloud Native Configuration Console Installation, Upgrade, and Fault Recovery Guide*.

The following image represents a Kubernetes cluster with one instance of CNC Console and two instances of BSF. The single instance of the CNC Console is configuring two instances of BSF with different namespaces.

Figure 3-43 Support for Multiple Instance Deployment

With the support of multicluster deployment, BSF deployed in multiple Kubernetes clusters can be accessed using CNC Console. In a multicluster deployment, the CNC Console can manage BSF and OCCNE common services deployed in the remote Kubernetes clusters.

The following image represents multiple Kubernetes clusters with one CNC Console and two BSF deployments. The single instance of CNC Console is configuring two instances of BSF with different namespaces deployed in different clusters.

Figure 3-44 Support for Multicluster Deployment

3.36 Support for 3GPP NF Sets and Binding Headers

Oracle Communications Cloud Native Core Binding Support Function supports the 3GPP NF Sets and Binding Headers in Model-B (Direct communication) and Model-C (Indirect communication). Using this feature, BSF can construct and send a binding header in the response messages to PCF for successful call processing.

Note

Since BSF is a producer NF, it does not send any notifications to consumer NFs. BSF sends an HTTP request message only to NRF for management purpose.

NF Set: NF set is a group of interchangeable NF instances supporting similar services and network slices. In an NF set, the NF instances can be geographically distributed, but have access to the same context data. The NF instances can be deployed in such a pattern so that several instances are present within an NF set to provide distribution, redundancy, and scalability as a set. The NF instances of an NF set are equivalent and share the same MCC, MNC, NID (for SNPN), NF type, and NF Set ID.

Binding Headers: The Binding headers indicate the suitable target NF producer instance for NF service instance selection, reselection, and routing of subsequent requests associated with a specific NF producer resource or context. It allows the NF producer to indicate that the NF consumer, for a particular context, should be bound to an NF service instance, or NF set depending on local policies. Binding can also be used by the NF consumer to indicate the suitable NF consumer instances for notification target instance reselection and routing of subsequent notification requests, associated with a specific notification subscription.

BSF supports the following binding header as defined in 3GPP:


Table 3-35 Supported Headers

Header Name	Description
3gpp-Sbi-Binding	This header is used to communicate the binding information from an HTTP server for storage and subsequent use by an HTTP client. This header contains a comma-delimited list of Binding Indications from an HTTP server for storage and use of HTTP clients. The absence of this parameter in a Binding Indication in a service request is interpreted as "callback". Note: In the current release, the following are not supported: - Binding levels – nfservice-instance and nfservice-set - Attributes – recoverytime and notif-receiver

Binding Support Function supports the NF Set and Binding Header functionality in all SBI interfaces.

Example

The following diagram depicts an example where communication between PCF and BSF Management service takes place:

Figure 3-45 Example of NF Set and Binding Header in BSF

The above call flow diagram describes a scenario where PCF sends a request to register PcfBindings towards the BSF Management service through Ingress Gateway. The register request does not contain a 3gpp-sbi-binding header.

Once the request is received at BSF, the BSF Management service adds 3gpp-sbi-binding to the response and sends it back to the PCF with HTTP status code 201.

If the user does not want to add a binding header to response messages, then it can be configured through CNC Console or REST APIs for BSF.

Managing NF Sets and Binding Header Support

Enable and Configure

The NF Sets and Binding header support can be enabled and configured for the BSF Management service interface using any of the following two ways:

- **Using CNC Console:** Perform the feature configurations on the **NF Communication Profile** page. **NF Communication Profile** must be attached to **PCF Communication Profile** under the **Audit** section of the **Management Service** page. For more information about the configurations, see [Service Configurations](#).
- **Using REST API:** BSF provides the following REST API for NF Sets and Binding Headers configuration:
API: {apiRoot}/oc-bsf-configuration/v1/services/nfcommprofiles

You can perform the GET and PUT operations to configure this feature. For more information about REST API configuration, see the "Management Service" section in *Oracle Communications Cloud Native Core Binding Support Function REST Specification Guide*.

Observe

BSF uses the Ingress metrics to contain information about the NF bindings used by PCF. The following metrics contains the information about NF bindings used by PCF:

- `ocpm_ingress_request_total` with the new dimensions – `pcf_id` and `pcf_set_id`

Note

The dimensions are populated only for `Nbsf_Management_Register` requests.

- `ocpm_ingress_response_total` with the new dimensions – `binding_level` and `binding_id`

For more information, see the [BSF Metrics](#) section.

Maintain

The BSF logs include the NF binding information sent by BSF. The logs include information about the following headers:

- `location`
- `3gpp-sbi-binding`

The following is a sample log for PCF binding register request:

```
{
  "instant": {
    "epochSecond": 1636550691,
    "nanoOfSecond": 280882458
  },
  "thread": "XNIO-1 task-1",
  "level": "DEBUG",
  "loggerName":
"ocpm.bsf.api.management.controller.BindingSupportManagementServiceAPIControl
ler",
  "message": "PCF binding: PcfBinding [supi=imsi-411411000000011,
gpsi=5084948001, ipv4Addr=10.10.10.16, ipv6Prefix=null, ipDomain=null,
macAddr48=null, dnn=internet, pcfFqdn=pcf-smsservice.oracle.com,
pcfIpEndpoints=null, pcfDiamHost=pcf-smsservice.oracle.com,
pcfDiamRealm=oracle.com, snssai=Snssai [sst=11, sd=abc123],
pcfId=fe7d992b-0541-4c7d-ab84-c6d70b1b0123,
pcfSetId=setxyz.pcfset.5gc.mnc015.mcc345, bindLevel=NF_SET]",
  "endOfBatch": false,
  "loggerFqn": "org.apache.logging.slf4j.Log4jLogger",
  "threadId": 45,
  "threadPriority": 5,
  "messageTimestamp": "2021-11-10T13:24:51.280+0000"
}
```

3.37 Support for User-Agent Header

User-Agent header helps the producer Network Function (NF) to identify the consumer NF that has sent the request. To implement this, 3GPP introduced the use of User-Agent header for consumers to include the same in service requests. Additionally, producers may require to

support the validation of the User-Agent headers to complete the request identification process in the network.

With the integration of this feature, User-Agent header helps the producer Network Function (NF) to identify the consumer NF that has sent the request.

The following format is used to generate User-Agent header:

<NF Type>-<Instance-Id> <FQDN>

where, <NF Type> is the type of the Network Function.

<Instance-Id> is the instance ID of the NF.

<FQDN> is the FQDN of the NF.

Example: BSF-54804518-4191-46b3-955c-ac631f953ed8 bsf1.east.5gc.mnc012.mcc234.3gppnetwork.org

Following validations are made once the feature is enabled:

- If the user-agent header is present, then it is matched with the configured NF types. If a match is found, then validation is successful and request is allowed to pass. If a match is not found, then request is rejected with a configurable error code.
- If the user-agent header is present and has multiple values, then the request is rejected with a configurable error code. Hence, the user-agent header, if present should have a single value.
- If the user-agent header is not present and validationType is relaxed, then validation is not made and request is allowed to pass. If validationType is strict, then request will be rejected with a configurable error code.

Managing Support for User-Agent Header in Ingress Gateway

Enable

You can enable the User-Agent Header feature using REST or Helm configuration.

- Helm: Set the value of the parameter `userAgentHeaderValidationConfigMode` to Helm in the custom-values.yaml file. For more information, see the *Oracle Communications Cloud Native Core Binding Support Function Installation and Upgrade Guide*.
- REST API: Set the value of the parameter `userAgentHeaderValidationConfigMode` to REST in the custom-values.yaml file. REST configuration from the JSON bodies sent to path: `/bsf/nf-common-component/v1/igw/useragentheadervalidation` is stored in a database under the `common_config` table. For more information, see the *Oracle Communications Cloud Native Core Binding Support Function REST Specification Guide*.

Configure

You can configure the User-Agent Header using REST or Helm.

To configure User-Agent header at Ingress Gateway using Helm, you need to perform the configurations:

```
#User-Agent header validator configuration
#Mode of configuration. Can be either HELM or REST
userAgentHeaderValidationConfigMode: HELM
userAgentHeaderValidation:
  enabled: true
```

```
# If User-Agent header is not present or it's value is null in the incoming
request then
validation type can be used to skip or perform validation. If set to strict
then validation will be performed.
# If set to relaxed then validation will be skipped.
validationType: relaxed
# List of consumer NF Types to be matched against the value of User-Agent
header in the request
consumerNfTypes:
  - "SMF"
  - "AMF"
  - "UDR"
  ...
```

To configure the User-Agent header at Ingress Gateway using REST API, see *user-Agent Header in Oracle Communications Cloud Native Core Binding Support Function REST Specification Guide*.

Managing Support for User-Agent Header in Egress Gateway

Enable

You can enable the User-Agent Header feature using REST or Helm configuration.

- Helm: Set the value of the parameter `userAgentHeaderConfigMode` to Helm in the custom-values.yaml file. For more information, see the *Oracle Communications Cloud Native Core Binding Support Function Installation and Upgrade Guide*.
- REST API: Set the value of the parameter `userAgentHeaderConfigMode` to REST in the custom-values.yaml file. REST configuration from the JSON bodies sent to path: `"/bsf/nf-common-component/v1/egw/useragentheader"` is stored in a database under the `common_config` table. For more information, see the *Oracle Communications Cloud Native Core Binding Support Function REST Specification Guide*.

Configure

You can configure the User-Agent Header using REST or Helm.

To configure User-Agent header at Egress Gateway using Helm, you need to perform the configurations:

```
userAgentHeaderConfigMode: HELM
userAgentHeader:
  enabled: false # flag to enable or disable the feature
  nfType: "PCF" # NF type of consumer NF
  nfInstanceId: "2d8e8e68-24ad-11ed-861d-0242ac120002" # NF type of consumer
  NF
  addFqdnToHeader: true # Flag to add fqdn. If enabled then user-agent header
  will be
  generated along with the fqdn configured otherwise fqdn will not be added
  nfFqdn: "oracle1.pcf.pacific.org" #fqdn of NF. This is not the fqdn of
  gateway
  overwriteHeader: true

oauthClient:
  enabled: false
  nfInstanceId: fe7d992b-0541-4c7d-ab84-c6d70b1b01b1
  nfType: PCF
```

To configure the User-Agent header at Egress Gateway using REST API, see user-Agent Header in *Oracle Communications Cloud Native Core Binding Support Function REST Specification Guide*.

Observe

BSF provides the following metric specific to User-Agent Header feature:

- `oc.ingressgateway.http.requests`

For more information, see [User-Agent Header Metrics](#) section.

3.38 Support for Active Sessions Counter


Active sessions are the unique PCF binding sessions in every BSF instance.

Note

The active sessions count also includes stale sessions until BSF removes it from its local database.

For more details on how stale sessions are handled, see [Handling Stale Session in BSF](#).

Figure 3-46 Retrieving Active Sessions Count




The Audit-Service periodically counts the simultaneous unique active sessions, over a configurable period of time. The default value of the time period is set to 15 mins. The time period can be configured to any value between 1 to 60 minutes.

The bsf-mgt-svc requests the Audit-Service to count active sessions for configurable time period. The Audit-Service periodically finds active sessions count and publishes the count as a metric.

To get the instantaneous value of active sessions count, you can query the CM service, which internally calls Audit-Service to fetch the value.

Figure 3-47 Active Sessions Count



Note

The audit registration process is used to register for audit registration as well as to enable active sessions counting.

The bsf-mgt-svc sends an audit registration request to the Audit-Service.

If the registration is successful, the Audit-Service responds with a 200 OK message.

If the registration fails due to any issue with the request, the Audit-Service responds with a Bad Request - 400 message.


If the registration fails due to any other internal reason while processing the request, the Audit-Service responds with an Internal Server Error - 500 message.

Note

Audit registration is not required for CM Service to query the instantaneous active sessions count.


Once the registration is successful, the Audit-Service periodically finds active sessions count and publishes the count as a metric.

Figure 3-48 Get Instantaneous Active Sessions Count - Successful



The CM Service sends a Get maximum active sessions count request to Audit-Service. If the CM Service call to Audit-Service is successful, it responds with a 200 OK message.

Figure 3-49 Get Instantaneous Active Sessions Count - Failure



If the CM Service call to Audit-Service fails, it responds with a 404 message.

ENABLE

You can enable the Active Sessions Counter feature using CNC Console or REST API.

- **Using CNC Console:** Set the value of **Count Active Sessions** parameter to true on **Active Sessions Counting** page. For more details, see *Active Sessions Count* under [Service Configurations](#) section.

- **Using REST API:** BSF provides the following REST API for Active Sessions Counter configuration:

API: {apiRoot}/oc-bsf-configuration/v1/services/management-service

You can perform the GET and PUT operations to configure this feature.

To enable the feature, set the value of `activeSessionCounting.countRecords` parameter to true.

For more details, see *BSF REST Specifications* section in *Oracle Communications Cloud Native Core Binding Support Function REST Specification Guide*.

For more information, see [Active Session Query](#).

Configure

You can configure the Active Sessions Counter feature using CNC Console or REST API.

- **Using CNC Console:** Set the value of **Session Count Interval (in minutes)** parameter on **Active Sessions Counting** page. You can set the time interval to any value between 1 to 60 minutes. For more details, see *Active Sessions Count* under [Service Configurations](#) section.

For instantaneous query, use the **Active Session Query** tab under **Status and Query** in CNC Console.

For more information, see [Active Session Query](#).

- **Using REST API:** BSF provides the following REST API for Active Sessions Counter configuration:

API: {apiRoot}/oc-bsf-configuration/v1/services/management-service

You can perform the GET and PUT operations to configure this feature.

To configure the feature, set the value of `activeSessionCounting.countRecordsInterval` parameter.

For instantaneous query, BSF provides the following REST API:

API: {apiRoot}/oc-bsf-configuration/v1/activeSessionCount/pcfBindings

You can perform the GET operation to retrieve the instantaneous active sessions count.

For more details, see *BSF REST Specifications* section in *Oracle Communications Cloud Native Core Binding Support Function REST Specification Guide*.

Observe

BSF provides the following metric specific to Active Sessions Counter feature:

- `oc_db_active_session_count`
- `inbound_requests_total`

For more information, see [Active Sessions Count Metrics](#).

3.39 Controlled Shutdown of an Instance

CNC BSF supports controlled shutdown feature to provide the partial or complete isolation of the site from the network so that the operator can perform the necessary maintenance and recovery procedures when required. It helps the operator to perform the recovery procedures as per the requirement.

The site isolation is achieved by shutting down the load at gateways (Ingress Gateway, Egress Gateway, and Diameter Gateway) and updating the NF status as SUSPENDED at NRF.

Operational State

The site can be in one of the three possible operational states NORMAL, PARTIAL SHUTDOWN, or COMPLETE SHUTDOWN. The operational state can move to any of the states from the current state, there is no definitive order of state change. Currently, the operational state is stored in the common config server of the Ingress Gateway. It is read by Ingress Gateway, Egress Gateway, Diameter Gateway, and App-info periodically and action is triggered based on the current state.

Note

Since the operational state is stored in config server, the service instances will detect the state change after the config refresh is done. If the config refresh interval is set as 5 seconds, then the pods may recognize the operational state change after 5 seconds.

The operational state can be modified through CNC Console or REST API. Operation state configuration stored in the common config server will be read by the following services:

- Ingress Gateway
- Egress Gateway
- Diameter Gateway
- App-info
- Audit Service

Note

If the Disaster Recovery procedure is performed when the config backup was taken when the system was in PARTIAL or COMPLETE SHUTDOWN state, then manual intervention may be required to change the operational state back to NORMAL state.

Load Control

Gateways enforce load control when the system is in a PARTIAL or COMPLETE shutdown state. The level of load control varies based on the shutdown state. When in a PARTIAL shutdown state, no new session establishments are allowed so session creation messages will be rejected (with configured error code) in this state. When in complete shutdown, no messages are allowed.


Note

When the system is in COMPLETE SHUTDOWN state, audit service triggered notification or diameter messages will be rejected at respective gateways.


Call Flow for Diameter Gateway

NORMAL State: If the Controlled shutdown operational state is NORMAL, then the Diameter Gateway processes the message as normal.

PARTIAL SHUTDOWN: If the controlled shutdown operational state is PARTIAL SHUTDOWN, then the Diameter Gateway accepts only in-session messages and rejects all CCR-I and AAR-I messages.



COMPLETE SHUTDOWN: If the controlled shutdown operational state is COMPLETE SHUTDOWN, then the Diameter Gateway rejects all messages.




Call Flow for Egress Gateway

NORMAL State: If the Controlled shutdown operational state is NORMAL, then the Egress Gateway processes the message as normal.

PARTIAL SHUTDOWN: If the controlled shutdown operational state is PARTIAL SHUTDOWN, then the Egress Gateway processes the message as normal.

COMPLETE SHUTDOWN: If the controlled shutdown operational state is COMPLETE SHUTDOWN, then the Egress Gateway processes the request as follows:


- Forward all requests received from NRF Client.
- Reject all requests received from any other services like UDR Connector, SM Service, AM Service, UE Service, and CHF Connector.




Call Flow for Ingress Gateway

NORMAL State: If the Controlled shutdown operational state is NORMAL, then the Ingress Gateway processes the message as normal.

PARTIAL SHUTDOWN: If the controlled shutdown operational state is PARTIAL SHUTDOWN, then the Ingress Gateway accepts only in-session messages and rejects all SM-Create requests.



COMPLETE SHUTDOWN: If the controlled shutdown operational state is COMPLETE SHUTDOWN, then the Ingress Gateway rejects all incoming requests.



Audit Service

The **Disabling Audit on Controlled Shutdown** feature allows you to gracefully shut down audit and timer services in a controlled manner. When enabled, this feature ensures that during a **COMPLETE_SHUTDOWN**, the Audit Service and Timer Service are temporarily disabled, preventing interruptions and ensuring a smooth shutdown process. This feature is particularly useful for maintenance or upgrade scenarios where you need to pause auditing activities without losing data or disrupting ongoing operations.

Depending on the **operationalState**, the system behaves as follows:

- **NORMAL / PARTIAL_SHUTDOWN:** Audit and Timer Services continue to function as usual.
- **COMPLETE_SHUTDOWN:**
 - **Audit Service:** All audit operations are interrupted, including stale session notifications and record deletions. All entries in `audit_schedule_table` are transitioned to `CTRL_SHUTDOWN_INTRPT` state, except those already in `DEREGISTERED` or `INTERRUPTED` states. The **Audit Now** and **Audit Resume** functionalities are disabled. But, the **Audit Pause** functionality remains enabled. The audit service supports registration, deregistration, and fetch requests.
 - **Timer Service:** Timer expiry is disabled. The timer expiry notifications are paused during this state.

Note

Once the system returns to **NORMAL** or **PARTIAL_SHUTDOWN**, the Audit Service resumes normal operations, and all audit entries which are in `CTRL_SHUTDOWN_INTRPT` state are moved back to the *QUEUED* state.

If the current version has the **enableControlledShutdown** flag enabled and the **operationalState** is set to **COMPLETE_SHUTDOWN**, follow these steps before initiating a rollback:

1. Manually transition all audit records to the **INTERRUPTED** state by clicking the **Audit Pause** button in the CNC Console.
2. After the rollback is completed, manually move these records back to the **QUEUED** state to resume auditing.

Note

Failure to complete the audit record state transition process will leave audit records stuck in the **COMPLETE_SHUTDOWN** state, blocking audits from restarting after a rollback to a version that doesn't support this feature.

Managing Controlled Shutdown of an instance

Enable

You can enable or disable the Controlled Shutdown feature by using the `enableControlledShutdown` parameter in the `custom.yaml` file. This parameter is set as `false` by default. You can enable it by setting its value as `true`. For more information, see *Controlled Shutdown Configurations* section in the *Oracle Communications Cloud Native Binding Support Function Installation and Upgrade Guide*.

Configure

Diameter Gateway and Ingress Gateway can be configured through CNC Console. For more information, see [Controlled Shutdown Configurations](#).

Egress Gateway routes configuration for controlled shutdown is done through Helm. For more information, see the *Controlled Shutdown Configurations* section in the *Oracle Communications Cloud Native Core Binding Support Function Installation and Upgrade Guide*.

Observe

Metrics

CNC BSF provides the following metrics specific to controlled shutdown feature:

- `system_operational_state`
For more information, see [CM Service Metrics](#).
- `diam_controlled_shutdown_message_reject_total`
For more information, see [Diameter Gateway Metrics](#).
- `audit_schedule_current_status`
For more information, see [Audit Service Metrics](#).
- `occnp_audit_schedule_status_total`
For more information, see [Audit Service Metrics](#).

Alerts

CNC BSF provides the following alerts for controlled shutdown feature:

- `SYSTEM_IMPAIRMENT_MAJOR`
- `SYSTEM_IMPAIRMENT_CRITICAL`
- `SYSTEM_OPERATIONAL_STATE_NORMAL`

- SYSTEM_OPERATIONAL_STATE_PARTIAL_SHUTDOWN
- SYSTEM_OPERATIONAL_STATE_COMPLETE_SHUTDOWN

For more information, see [BSF Alerts](#).

3.40 Graceful Termination of Kubernetes Pods

This feature is to support BSF NF's Kubernetes pods to terminate gracefully to reduce traffic loss.

In Kubernetes cluster, pods can get deleted due to various events. Few problems that can arise with abnormal termination:

1. A pod that is currently in the middle of processing a request is removed, leads to incomplete processing.
2. Kubernetes routes traffic to pods that have already been deleted, resulting in stale session at local or peer user.
3. The corruption of data in cache or database.

BSF services handles graceful termination of HTTP2 connections by

- accepting new request, but sends an immediate 503 service unavailable response to the clients.
- wait for ongoing response messages to complete the session.
- terminate the TCP connections and any DB transactions gracefully.

The applications performs all the above tasks before termination of the pod or before the grace period expires. If the grace period expires, and the process hasn't gracefully shutdown, the container runtime will force kill, stopping the pod immediately.

Note

Diameter gateway service handles graceful termination of "Disconnect-Peer-Request" and "Disconnect-Peer-Answer" along with HTTP2 connections.

The grace period is configurable in the custom-values.yaml file, the default value is set at 30 seconds. Here is a sample configuration for graceful shutdown parameters in custom-values.yaml file:

```
bsf-management-service:
  gracefulShutdown:
    gracePeriod: 30s
config-server:
  gracefulShutdown:
    gracePeriod: 30s
```

BSF Services that support Kubernetes graceful shutdown:

- BSF Management Services
- Config Server
- CM service

- Audit service
- Query service
- Diameter Gateway
- AppInfo service
- PerfInfo service
- Ingress Gateway
- Egress Gateway
- NRF client service

For more information on setting the gracePeriod value for different BSF services, see section "Customizing Binding Support Function" in *Oracle Communications Cloud Native Core Binding Support Function Installation and Upgrade Guide*.

3.41 NF Scoring for a Site

The NF Scoring feature calculates the score for a site based on Network Function (NF) specific factors such as metrics, and alerts. The NF Scoring feature helps the operator to determine the health of a site as compared to other sites. Comparing the NF scores within or across the sites helps the customers to choose the site.

One of the use cases is the Controlled Shutdown feature that allows the operator to partially or completely isolate the site. The NF Scoring feature helps the operators to choose which site to partially or completely isolate based on NF scoring.

App-Info service queries and calculates NF score as it has the site information.

App Info Scoring Mechanism:

App Info reads the configurations from the common configuration server to check if NF Scoring functionality is enabled or not. It works in the following ways:

- **Continuous NF Score Calculation:** When the NF Scoring feature is enabled, app info periodically reads the configurations to calculate the score.
- **On-Demand NF Score Calculation:** When the NF Scoring feature is enabled, app info fetches all the factors or criteria to calculate the NF Score. It is real-time fetching of factors and then the NF score is calculated on demand.

Table 3-36 NF Scoring Criteria

Factors	Default Score	Formula to calculate Factor Score	Details
TPS	20	$\min((\text{<Current\_TPS>} / \text{<Max\_TPS>} * \text{<Max\_TPS\_Score>}), \text{<Max\_TPS\_Score>})$	$\text{<Current\_TPS>} = \text{IGW} + \text{EGW} + \text{Diameter Ingress} + \text{Diameter Egress}$ Max_TPS: specifies the maximum TPS. Max_TPS_Score: Specifies the maximum score of the TPS.


Table 3-36 (Cont.) NF Scoring Criteria

Factors	Default Score	Formula to calculate Factor Score	Details
Service	30	$A / N * \text{<Max\_SVC\_Score>}$	A = Number of available services N = Number of configured services <Max_SVC_Score> Specifies the maximum score of the Service Health.
Connection	20	$\min((\text{<Conn\_Current>} / \text{<Conn\_Total>} * \text{<Conn\_Score>}), \text{<Conn\_Score>})$	<Conn_current> specifies the number of connections from network to Policy. <Conn_Total> specifies the total number of connections expected from network to Policy. <Conn_Score> specifies the score for the connection.
Replication_health	30	$\min((\text{<Site\_Current>} / \text{<Site\_Total>} * \text{<Site\_Score>}), \text{<Site\_Score>})$	<Site_Total> specifies the total number of possible replication links. <Site_Current> specifies the available active healthy links. Replication_health score specifies the score for the Replication_health.
Locality_Preference	5	NA	The value of Locality_Preference is added for NF score calculation.
Critical_Alerts	2	$\text{CrN} * \text{Configured\_Score\_Critical\_Alerts}$	CrN is the Number of active critical alarms. Configured_Score_Critical_Alerts specifies the score configured by the user.
Major_Alerts	1	$\text{MaN} * \text{Configured\_Score\_Major\_Alerts}$	MaN is the Number of active Major alarms. Configured_Score_Major_Alerts specifies the score configured by the user.
Minor_Alerts	0	$\text{MiN} * \text{Configured\_Score\_Minor\_Alerts}$	MiN is the Number of active Minor alarms. Configured_Score_Minor_Alerts specifies the score configured by the user.


Formula for NF scoring of a site: Sum of TPS score, Service score, Connection score, Replication_health, and Locality_Preference score subtracted from Alerts scores.

Call Flows

This section describes examples of the call flows for the NF Scoring feature:

Figure 3-50 Call flow to Save Configuration Data

The operator sends a request to save the configuration is sent to the CM service. It saves the configuration data to the common config database. App-Info reads the configuration data and returns the acknowledgment.

Figure 3-51 Call flow to Get the NF Score

The operator sends a request to CM service to get the score. CM service requests it to App-Info. App-Info queries and calculates NF_Score.

Managing Controlled Shutdown of an instance

Enable

You can enable this feature by selecting the **Enable NF Scoring** field in the **Settings** page of **NF Scoring**.

For more information about enabling the feature through CNC Console, see [NF Scoring Configurations](#).

Configure

The NF Scoring feature can be configured through CNC Console. For more information, see [NF Scoring Configurations](#).

Note

You can configure the env variable, `<NF_SCORING_INTERVAL>`, in deployment of app-info. Default value is 30 seconds (changing the env variable would result into restart of app-info pod).

Observe

Metrics

BSF provides the following metrics specific to NF Scoring feature:

- nfscore
- nfScoringFactorActualValue

For more information, see [AppInfo Metrics](#).

3.42 NRF Client Retry and Health Check

With the alternate route retry feature, Policy can attempt service requests to an alternate secondary Network Repository Function (NRF) when the primary NRF throws errors. In addition, the health status check feature actively monitors the health of the NRFs and provides the list of the healthy NRFs for session requests only. The NRF client also provides the health information of NRFs to other services if requested, and notifies any change in the health status.

For a given service request, the NRF client initiates a request towards a healthy and the highest priority NRF. If the NRF client receives a failure response for the request or the request timed-out, it attempts to send the request to the same NRF for `NrfRetryConfig.primaryNrfRetryCount` number of times. If a success response is received before the retry count gets exhausted, NRF client accepts the response and does not send any further service requests. However, if NRF client fails to receive a success response, it attempts to send the service request to an alternate NRF. The alternate NRF is selected based on the assigned priority and health status.

If the NRF Client receives a `retryAfterTime` value in the response header from the NRF, the NRF Client halts any further attempts to the NRF and flags the NRF as unhealthy for the specified time period. The NRF client retries the service request to alternate NRFs until any one of the following conditions are met:

- NRF-client receives a success response.
- `NrfRetryConfig.alternateNRFRetryCount` is exhausted.
- All attempts to available healthy NRFs are exhausted.

Once any of the listed conditions are met, NRF-client accepts the response and proceed.

NRF Client marks NRF as unhealthy under the following conditions:

- If the NRF Client receives a *retryAfterTime* value in the response header from the NRF, then NRF will be unhealthy for a time period as defined in *retryAfterTime*.
- If the status code received is available in the default values for *errorCodeReasonsForFailure*, then NRF will be unhealthy for a period of time as defined in `ConfigMap.data: profile.retryAfterTime`.
- If the status code received is available in the default values for *errorCodeReasonsForFailure* and all the retry attempts are exhausted.
- If NRF Client receives an error from Gateway service and the error is configured in the *gatewayErrorCodes* with all the exhausted retry attempts.

Note

- If NRF Client receives an error from Gateway service and the error is not configured in the *gatewayErrorCodes*, then NRF remains marked as healthy.
- *HealthCheckConfig* and *NRFRetryConfig* must be configured for the NRF Client functionality to work as expected.
- NRF Client considers a response as failure only when it is configured in the *errorReasonsForFailure* parameter in the `custom-values.yaml` file. The primary and non-primary NRFs must be geo-redundant for the NRF Retry mechanism to work.
- For autonomous procedures such as *NfRegistration* and *NfHeartbeat*, NRF-client continues to retry sending service requests till a success response is received. For details on NRF Client configuration parameters, see *NRF Client Configuration* section in *Oracle Communications Cloud Native Core, Converged Binding Support Function Installation, Upgrade and Fault Recovery Guide*.

3.43 BSF Message Feed for Monitoring

In order to enable correlation of the internal and external (request/response) messages for all the transactions initiated by the producer and consumer NFs, BSF supports copying the messages at Ingress and Egress Gateways.

This feature allows NFs using Ingress and Egress Gateways to report every incoming and outgoing message to Oracle Communications Network Analytics Data Director (OCNADD) monitoring system.


That is, OCNADD is a message store to keep a copy of each request and response processed through IGW & EGW.

The insights on these messages enable NFs to integrate with external 5G SBI monitoring system for:

- Call Tracing / Tracking
- Live debugging

Architecture

Figure 3-52 BSF Message Feed Architecture



OCNADD is a Network Data Broker part of the Network Analytics suite of products. OCNADD receives network data traffic information from various sources such as 5G NFs and Non-5G Nodes and sends the data securely to subscribed consumer (3rd Party tools) after applying its powerful and configurable filtering, replication, and aggregation rule corresponding to subscribed consumers. For more information on OCNADD, see *Oracle Communications Network Analytics Data Director User Guide*.

5G NF Kafka Producer is used as the source to send the data stream towards OCNADD. The 5G NFs use integrated Kafka producer services to stream the 5G South Bound Interface (SBI) messages along with metadata added by NFs to OCNADD.

Managing BSF Message Feed

Enable

BSF Message Feed feature can be enabled using Helm parameters either at the time of BSF installation or during the software upgrade.

`ingress-gateway.message-copy.enabled` parameter is used to enable copying messages passing through Ingress Gateway.

`egress-gateway.message-copy.enabled` parameter is used to enable copying messages passing through Egress Gateway.

For more information, see *Configuring Ingress Gateway*, *Configuring Egress Gateway*, and *Configuring Kafka for NF Message Feed* sections in *Oracle Communications Cloud Native Core, Binding Support Function Installation, Upgrade, and Fault Recovery Guide*.

Configure

BSF Message Feed feature can be configured using Helm parameters either at the time of BSF installation or during the software upgrade.

For more information, see *Configuring Ingress Gateway*, *Configuring Egress Gateway*, and *Configuring Kafka for NF Message Feed* sections in *Oracle Communications Cloud Native Core, Binding Support Function Installation, Upgrade, and Fault Recovery Guide*.

SASL_SSL Configuration for BSF Message Copy

As there is no certificate-based client authentication required, a trustStore is created at BSF.

BSF contains placeholders to accept caroot certificates, which are then translated into trustStore using Gateway init-containers.

BSF uses native SSL functionality provided by Gateway services. SSL service block gets activated or used, when `enableIncomingHttps` is set to `true`. The same configuration is used for message copy SSL configuration too.

To configure only BSF-DD SSL communication without native SSL functionality, configure **caBundle** and **trustStorePassword** sections with appropriate secret configurations.

To use both native SSL functionality and BSF-DD SSL communication, add the caRoot certificate of Kafka broker to the existing caRoot certificate by appending Kafka broker ca certificate after the existing certificate.

1. Generate SSL certificates.

Note

Creation process for private keys, certificates and passwords is based on discretion of user or operator.

2. Before copying the certificates to the secret, add the DD Root certificates contents into the CA certificate(**caroot.cer**) generated for NRF.

Note

Make sure to add 8 hyphens "-" between 2 certificates.

```
-----BEGIN CERTIFICATE-----
<existing caroot-certificate content>
-----END CERTIFICATE-----
-----
-----BEGIN CERTIFICATE-----
<DD caroot-certificate content>
-----END CERTIFICATE-----
```

3. Create a secret for authentication with DD.

To create a secret store the password in a text file and use the same file to create a new secret.

```
kubectl create secret generic ocingress-secret --from-
file=ssl_ecdsa_private_key.pem --from-file=rsa_private_key_pkcs1.pem --
from-file=ssl_truststore.txt --from-file=ssl_keystore.txt --from-
file=caroot.cer --from-file=ssl_rsa_certificate.crt --from-
file=ssl_ecdsa_certificate.crt --from-file=sasl.txt -n <namespace>
```

```
kubectl create secret generic ocegress-secret --from-
file=ssl_ecdsa_private_key.pem --from-file=ssl_rsa_private_key.pem --from-
file=ssl_truststore.txt --from-file=ssl_keystore.txt --from-
file=ssl_cabundle.crt --from-file=ssl_rsa_certificate.crt --from-
file=ssl_ecdsa_certificate.crt --from-file=sasl.txt -n <namespace>
```

4. Provide appropriate values for the SSL section.
SSL configuration:

```
service:
  ssl:

    privateKey:
      k8SecretName: ocegress-secret
      k8Namespace: bsf
      rsa:
        fileName: rsa_private_key_pkcs1.pem
      ecdsa:
        fileName: ssl_ecdsa_private_key.pem

    certificate:
      k8SecretName: ocegress-secret
      k8Namespace: bsf
      rsa:
        fileName: tmp.cer
      ecdsa:
        fileName: ssl_ecdsa_certificate.crt

    caBundle:
      k8SecretName: ocegress-secret
      k8Namespace: bsf
      fileName: caroot.cer

    keyStorePassword:
      k8SecretName: ocegress-secret
      k8Namespace: bsf
      fileName: key.txt

    trustStorePassword:
      k8SecretName: ocegress-secret
      k8Namespace: bsf
      fileName: trust.txt

    initialAlgorithm: RS256
```

5. Configure the message copy feature.

```
messageCopy:
  enabled: true
  copyPayload: true
  topicName: BSF
  ackRequired: false
  retryOnFailure: 0
  security:
    enabled: true
    protocol: SASL_SSL
    tlsVersion: TLSv1.2
    saslConfiguration:
      userName: ocnadd
      password:
        k8SecretName: ocegress-secret
        k8Namespace: bsf
        fileName: sasl.txt
```

6. Make sure to configure the correct SASL_SSL port in `kafka.bootstrapAddress` attribute. To get the correct value of this, refer to DD Kafka's Values.yaml file.

Observability

Metrics

The following metrics are used to count the ingress and egress messages at the gateways:

- `oc_ingressgateway_msgcopy_requests_total`
- `oc_ingressgateway_msgcopy_responses_total`
- `oc_egressgateway_msgcopy_requests_total`
- `oc_egressgateway_msgcopy_responses_total`

For more information, see:

- [Ingress Gateway Metrics](#)
- [Egress Gateway Metrics](#)

Alerts

The following alerts are raised when OCNADD is not reachable:

- `INGRESS_GATEWAY_DD_UNREACHABLE_MAJOR`
- `EGRESS_GATEWAY_DD_UNREACHABLE_MAJOR`

For more information, see

- [INGRESS_GATEWAY_DD_UNREACHABLE_MAJOR](#)
- [EGRESS_GATEWAY_DD_UNREACHABLE_MAJOR](#)

3.44 Optimizing Lookup in BSF Management Service Database

BSF supports optimized storage and lookup of PCF binding information in BSF Management Service.

BSF stores binding information in the `pcf_binding_v2` table and maintains dedicated lookup tables for IPv4 address, IPv6 prefix, and MAC address identifiers. During DISCOVERY operations, BSF can use these lookup tables and primary-key reads to identify the corresponding PCF binding record. If a lookup cannot be resolved through the new tables, BSF can fall back to legacy index-based lookup.

The supported lookup identifiers are IPv4 address, IPv6 prefix, and MAC address. These identifiers may be used with additional discovery attributes, such as SUPI, GPSI, DNN, IP domain, S-NSSAI, or site ID, depending on the discovery request.

During installation or upgrade, BSF creates the `pcf_binding_v2` table and the associated IPv4, IPv6, and MAC lookup tables. The `json_content` column is stored as `VARCHAR(12000)` and the `binding_id` column is stored as `VARCHAR(128)` in the `pcf_binding_v2` table.

When migration is enabled, REGISTER operations write new binding records to the `pcf_binding_v2` table. Existing binding records are migrated using a lazy migration approach. Eligible legacy records are migrated during Audit processing.

BSF uses a binding identifier format that contains table information used to determine the target table for lookup operations. New records created in the `pcf_binding_v2` table use a v2-prefixed binding identifier. Existing records remain available throughout the migration process.

BSF retains the existing `pcf_binding` table during migration and supports fallback lookup behavior during the transition period. Legacy and migrated records can coexist until migration is complete.

The `pcf_binding` table must not be deprecated until migration is complete and the table contains no binding records. After all binding records are migrated and lookup is configured to use only the `pcf_binding_v2` table, the original `pcf_binding` table is expected to contain zero records and can be deprecated.

For High Availability (HA) or Geo Redundancy (GR) deployments, feature configuration must be consistent across all sites. If the feature is enabled on some sites but not others, BSF may reject requests during the transition period because different sites may use different binding identifier formats.

This feature updates the `binding_id` schema to `VARCHAR(128)`. It does not change the external BSF APIs used for PCF binding registration, discovery, or deletion. Request and response formats for Nbsf Management Service procedures remain unchanged.

Migration Procedure for Multi-Site Deployments

The following table describes the recommended migration sequence and corresponding rollback procedure for migrating PCF binding records from LONG BLOB to VARCHAR in a two-site deployment.

Table 3-37 Migration Sequence

Step	Action	Rollback Procedure	Side Effects
1	Upgrade site1 to 26.1.201.	Software rollback on site1.	-
2	Upgrade site2 to 26.1.201.	Rollback site1 and site2.	-
3	Update the PCF_BINDING_TABLE_LOOKUP advanced setting to 1 on site1 or site2.	Rollback site1 and site2.	-
4	Update the PCF_BINDING_TABLE_LOOKUP advanced setting to 1 on the other site.	Rollback site1 and site2.	-

Table 3-37 (Cont.) Migration Sequence

Step	Action	Rollback Procedure	Side Effects
5	Set the ENABLE_PCF_BINDING_TABLE_MIGRATION advanced setting to true on site1 or site2.	Rollback site1 and site2.	Loss of bindings in the pcf_binding_v2 table.
6	Set the ENABLE_PCF_BINDING_TABLE_MIGRATION advanced setting to true on the other site.	Rollback site1 and site2.	Loss of bindings in the pcf_binding_v2 table.
7	Set the PCF_BINDING_TABLE_LOOKUP advanced setting to 2 on site1 or site2.	Rollback site1 and site2.	Loss of bindings in the pcf_binding_v2 table.
8	Set the PCF_BINDING_TABLE_LOOKUP advanced setting to 2 on the other site.	Rollback site1 and site2.	Loss of bindings in the pcf_binding_v2 table.
9	After migration is complete on a site, set the PCF_BINDING_TABLE_LOOKUP advanced setting to 3 on that site.	Rollback site1 and site2.	Loss of bindings in the pcf_binding_v2 table.
10	After migration is complete on the other site, perform the action described in step 9.	-	-

Table 3-38 High Availability and Geo Redundancy Behavior During Binding Table Migration

PCF_BINDING_TABLE_LOOKUP value set in both sites	Site1, upgraded to 26.1.201: CRUD behavior	Site2, upgraded to 26.1.201: CRUD behavior	Comments	Conclusion
1 - PCF_BINDING_FIRST	Reads from the pcf_binding table first, then falls back to the pcf_binding_v2 table. Creates or updates records in the pcf_binding_v2 table when migration is enabled.	Reads from the pcf_binding table first, then falls back to the pcf_binding_v2 table. Creates or updates records in the pcf_binding_v2 table when migration is enabled.	Initial migration state. Most records are still available in the pcf_binding table.	Recommended when 0% to 50% of data is migrated.
1 - PCF_BINDING_FIRST	Reads old records from the pcf_binding table first. Migrated or new records are still available through fallback to the pcf_binding_v2 table.	Reads old records from the pcf_binding table first. Migrated or new records are still available through fallback to the pcf_binding_v2 table.	Safe while migration is progressing and both old and new records exist.	Supported during migration. Move to value 2 after more than 50% of data is migrated.

Table 3-38 (Cont.) High Availability and Geo Redundancy Behavior During Binding Table Migration

PCF_BINDING_TABLE_LOOKUP value set in both sites	Site1, upgraded to 26.1.201: CRUD behavior	Site2, upgraded to 26.1.201: CRUD behavior	Comments	Conclusion
1 - PCF_BINDING_FIRST	Continues checking the pcf_binding table first, even when most or all records may already be available in the pcf_binding_v2 table.	Continues checking the pcf_binding table first, even when most or all records may already be available in the pcf_binding_v2 table.	Functional, but less optimal because lookup checks the old table first.	Not recommended after more than 50% of data is migrated. Use value 2 instead.
At this point, the feature will be enabled (ENABLE_PCF_BINDING_TABLE_MIGRATION= true) at both sites.				
2 - PCF_BINDING_V2_FIRST	Reads from the pcf_binding_v2 table first, then falls back to the pcf_binding table. Creates or updates records in the pcf_binding_v2 table when migration is enabled.	Reads from the pcf_binding_v2 table first, then falls back to the pcf_binding table. Creates or updates records in the pcf_binding_v2 table when migration is enabled.	Optimized for later migration stages while still allowing fallback to records that remain in the old table.	Recommended when more than 50% but less than 100% of data is migrated.
3 - PCF_BINDING_V2_ONLY	Uses only the pcf_binding_v2 table and lookup tables. No fallback to the pcf_binding table.	Uses only the pcf_binding_v2 table and lookup tables. No fallback to the pcf_binding table.	Use only after migration is complete on both HA/GR sites.	Recommended only when 100% of data is migrated.
0 - PCF_BINDING_ONLY	Uses only the pcf_binding table.	Uses only the pcf_binding table.	Optimized lookup and migration are not used.	Use only before migration starts or if optimized lookup is disabled.

Note

Both HA/GR sites must use the same PCF_BINDING_TABLE_LOOKUP value. Use value 1 while migration is starting or still below 50%. Use value 2 after more than 50% of data is migrated. Use value 3 only after 100% of data is migrated on both sites.

Recommendation to Speed Up Migration

During a maintenance window, the operator can use the audit service to speed up migration of records from the pcf_binding table to the pcf_binding_v2 table and the lookup tables.

1. Set Query to PCF to false.

2. Reduce the maxTTL value to enable faster audit of records.
3. Increase the notification rate so that a larger number of records are sent for audit.
4. At the end of the maintenance window, restore the original maxTTL and notification rate values.

Note

If Query to PCF is not set to false, BSF may send many unnecessary BSF-to-PCF audit requests that are required only for migration.

Migration Complete Handling


When the minor alert BSF_PCF_BINDING_TABLE_MIGRATED_PERCENTAGE is triggered, the operator must update the PCF_BINDING_TABLE_LOOKUP value to 3 in the advancedSettings configuration.

After the migration is complete, the operator can optionally delete the pcf_binding table from the database. Before deleting the table, perform the following steps:

1. Verify that the pcf_binding table is empty.
2. If audit is enabled, disable audit and then re-enable it using the GUI or REST API.
3. Manually drop the pcf_binding table from the database.

Call Flow Models

Register



1. The PCF sends a Binding register request to the BSF Ingress Gateway.
2. The BSF Ingress Gateway forwards the Binding register request to BSF Management.

3. BSF Management evaluates the configured binding table migration settings.
4. If `ENABLE_PCF_BINDING_TABLE_MIGRATION` is set to true and `PCF_BINDING_TABLE_LOOKUP` is set to a value greater than `PCF_BINDING_ONLY(0)`, BSF Management inserts the binding record into the `pcf_binding_v2` table.
5. The database returns a success response to BSF Management after the binding record is stored.
6. For each enabled lookup table, BSF Management inserts the corresponding lookup record or records into the lookup table. These records are populated using the identifier information available in the registered PCF binding.\
7. BSF Management pegs the `ocbsf_db_latency` metric with `db_operation` set to insert and status set to successful.
8. If `ENABLE_PCF_BINDING_TABLE_MIGRATION` is set to false, or if `PCF_BINDING_TABLE_LOOKUP` is set to `PCF_BINDING_ONLY(0)`, BSF Management inserts the binding record into the `pcf_binding` table instead.
9. The database returns a success response to BSF Management after the binding record is stored.
10. After the binding registration operation is complete, BSF Management returns a 201 CREATED response to the BSF Ingress Gateway.
11. The BSF Ingress Gateway forwards the 201 CREATED response to the PCF.


Note

Insertion into alternate lookup tables can fail. In such cases, the audit process handles reinsertion of any missing lookup records.

Deregister**Note**


The following diagrams are part of a single call flow which has been divided into separate sections.

pcf_binding_v2 Only




1. The PCF sends a Binding deregister request to the BSF Ingress Gateway.
2. The BSF Ingress Gateway forwards the Binding deregister request to BSF Management.
3. BSF Management evaluates the deregister request against the following condition: `("v2-" prefix present in binding_id) OR ("v2-" prefix NOT present in binding_id and PCF_BINDING_TABLE_LOOKUP == PCF_BINDING_V2_ONLY(3))`
4. When this condition is met, the binding record is expected to exist only in the `pcf_binding_v2` table. BSF Management does not fall back to the `pcf_binding` table.
5. BSF Management searches the `pcf_binding_v2` table for a binding record that matches the given `binding_id`.
6. The database returns the matching PCF binding record to BSF Management.
7. BSF Management deletes the binding record from the `pcf_binding_v2` table using the `binding_id`.
8. The database returns a success response after the binding record is deleted.
9. For each enabled lookup table, BSF Management deletes the corresponding lookup records using the Binding ID and UE address.
10. BSF Management pegs the `ocbsf_db_latency` metric with `db_operation` set to `delete_by_binding_id_and_ue_addr` and `status` set to `successful`.
11. BSF Management returns 204 NO CONTENT to the BSF Ingress Gateway.
12. The BSF Ingress Gateway forwards the 204 NO CONTENT response to the PCF.

pcf_binding_v2 First



1. The PCF sends a Binding deregister request to the BSF Ingress Gateway.
2. The BSF Ingress Gateway forwards the Binding deregister request to BSF Management.
3. BSF Management evaluates the deregister request against the following condition:
PCF_BINDING_TABLE_LOOKUP == PCF_BINDING_V2_FIRST(2)
4. When this condition is met, BSF Management first attempts to delete the binding record from the pcf_binding_v2 table.
5. BSF Management searches the pcf_binding_v2 table for a binding record that matches the given binding_id.
6. If the binding record is found in the pcf_binding_v2 table, the database returns the matching PCF binding to BSF Management.
7. BSF Management deletes the binding record from the pcf_binding_v2 table using the binding_id.
8. The database returns a success response after the binding record is deleted.
9. For each enabled lookup table, BSF Management deletes the corresponding lookup records using the Binding ID and UE address.
10. BSF Management pegs the ocbsf_db_latency metric with db_operation set to delete_by_binding_id_and_ue_addr and status set to successful.
11. If the binding record is not found in the pcf_binding_v2 table, BSF Management falls back to the pcf_binding table.
12. BSF Management deletes the binding record from the pcf_binding table using the binding_id.
13. The database returns a success response after the binding record is deleted from the pcf_binding table.
14. BSF Management returns 204 NO CONTENT to the BSF Ingress Gateway.
15. The BSF Ingress Gateway forwards the 204 NO CONTENT response to the PCF.

pcf_binding First



1. The PCF sends a Binding deregister request to the BSF Ingress Gateway.
2. The BSF Ingress Gateway forwards the Binding deregister request to BSF Management.
3. BSF Management evaluates the deregister request against the following condition:
`PCF_BINDING_TABLE_LOOKUP == PCF_BINDING_FIRST(1)`
4. When this condition is met, BSF Management first attempts to delete the binding record from the `pcf_binding` table.
5. If the binding record is found in the `pcf_binding` table, BSF Management deletes the binding record from the `pcf_binding` table using the `binding_id`.
6. The database returns a success response after the binding record is deleted.
7. If the binding record is not found in the `pcf_binding` table, BSF Management falls back to the `pcf_binding_v2` table.
8. BSF Management searches the `pcf_binding_v2` table for a binding record that matches the given `binding_id`.
9. The database returns the matching PCF binding record to BSF Management.
10. BSF Management deletes the binding record from the `pcf_binding_v2` table using the `binding_id`.
11. The database returns a success response after the binding record is deleted from the `pcf_binding_v2` table.
12. For each enabled lookup table, BSF Management deletes the corresponding lookup records using the Binding ID and UE address.
13. BSF Management pegs the `ocbsf_db_latency` metric with `db_operation` set to `delete_by_binding_id_and_ue_addr` and `status` set to `successful`.


14. For the default pcf_binding-only behavior, the record is expected to exist only in the pcf_binding table. BSF Management deletes the binding record from the pcf_binding table and does not fall back to pcf_binding_v2.
15. After the delete operation is complete, BSF Management returns 204 NO CONTENT to the BSF Ingress Gateway.
16. The BSF Ingress Gateway forwards the 204 NO CONTENT response to the PCF.

Discovery

Note

Discovery flows and Rx queries do not perform migration. Migration is handled by the audit process.

PCF_BINDING_ONLY (PCF_BINDING_TABLE_LOOKUP == 0)




1. The AF sends an Rx AAR-I request to the BSF Diameter Gateway.
2. The BSF Diameter Gateway sends a Discover Request to BSF Management.
3. BSF Management defines the search criteria for the discovery request.
4. BSF Management queries the pcf_binding table using index-based lookup.
5. If a matching binding record is found, the database returns the PCF binding details to BSF Management.
6. BSF Management returns the PCF binding record with a 200 OK response to the BSF Diameter Gateway.

7. The BSF Diameter Gateway sends an Rx AAR-I request to the PCF.
8. If a matching binding record is not found, BSF Management returns 204 NO CONTENT to the BSF Diameter Gateway.
9. The BSF Diameter Gateway returns an Rx AAA-I response to the AF.


PCF_BINDING_FIRST (PCF_BINDING_TABLE_LOOKUP == 1)

Note

The following diagrams are part of a single call flow which has been divided into separate sections.



1. The AF sends an Rx AAR-I request to the BSF Diameter Gateway.
2. The BSF Diameter Gateway sends a Discover Request to BSF Management.
3. BSF Management defines the search criteria for the discovery request.
4. BSF Management queries the pcf_binding table using index-based lookup.
5. If a matching binding record is found in the pcf_binding table, the database returns the PCF binding record to BSF Management.
6. BSF Management returns the PCF binding record with 200 OK to the BSF Diameter Gateway.
7. The BSF Diameter Gateway forwards the Rx AAR-I request to the PCF.




1. If no binding record is found in the `pcf_binding` table, BSF Management extracts the UE address identifier from the search criteria.
2. If the lookup table corresponding to the extracted UE address identifier is enabled, BSF Management queries the corresponding lookup table, such as the IPv6 prefix lookup table.
3. If lookup records are found, BSF Management pegs the `ocbsf_db_latency` metric with `db_operation` set to `select_by_ue_addr` and status set to `successful`. It also pegs the `ocbsf_alt_key_lookup` metric with `success` status.
4. BSF Management filters the lookup records using the search criteria.
5. If one or more lookup records match the search criteria, BSF Management queries the `pcf_binding_v2` table using the Binding IDs from the matching lookup records.
6. If matching PCF binding records are found in the `pcf_binding_v2` table, BSF Management sorts the records by the Created Date Time timestamp and retrieves the most recent PCF binding.
7. BSF Management returns the PCF binding record with 200 OK to the BSF Diameter Gateway.
8. The BSF Diameter Gateway forwards the Rx AAR-I request to the PCF.
9. If lookup tables are disabled, no lookup records are found, no lookup records match the search criteria, or no binding record is found using the Binding IDs from the lookup records, BSF Management falls back to index-based lookup in the `pcf_binding_v2` table.
10. If the fallback lookup finds matching PCF binding records, BSF Management pegs the `ocbsf_index_based_lookup_fallback` metric with `found` status.
11. BSF Management sorts the returned PCF binding records by the Created Date Time timestamp and retrieves the most recent PCF binding.
12. BSF Management returns the PCF binding record with 200 OK to the BSF Diameter Gateway.
13. The BSF Diameter Gateway forwards the Rx AAR-I request to the PCF.
14. If the fallback lookup does not find a binding record in `pcf_binding_v2`, BSF Management pegs the `ocbsf_index_based_lookup_fallback` metric with `not_found` status.
15. BSF Management returns 204 NO CONTENT to the BSF Diameter Gateway.
16. The BSF Diameter Gateway returns an Rx AAA-I response to the AF.

PCF_BINDING_V2_FIRST (PCF_BINDING_TABLE_LOOKUP == 2)


Note

The following diagrams are part of a single call flow which has been divided into separate sections.




1. The AF sends an Rx AAR-I request to the BSF Diameter Gateway.
2. The BSF Diameter Gateway sends a Discover Request to BSF Management.
3. BSF Management defines the search criteria for the discovery request.
4. BSF Management extracts the UE address identifier from the search criteria.
5. If the lookup table corresponding to the extracted UE address identifier is enabled, BSF Management queries the corresponding lookup table. For example, if the extracted UE address identifier is an IPv6 prefix, BSF Management queries the IPv6 prefix lookup table.
6. If lookup records are found in the lookup table, the database returns the lookup records to BSF Management.
7. BSF Management pegs the ocbf_db_latency metric with db_operation set to select_by_ue_addr and status set to successful.
8. BSF Management pegs the ocbf_alt_key_lookup metric with success status.
9. BSF Management checks whether the returned lookup records match the discovery search criteria.
10. If the lookup records match the search criteria, the flow continues with pcf_binding_v2 lookup using the Binding ID from the matching lookup record.
11. If no lookup record matches the search criteria, BSF Management pegs the ocbf_db_latency metric with db_operation set to select_by_ue_addr and status set to failure.
12. BSF Management pegs the ocbf_alt_key_lookup metric with failure status and not_found cause.
13. The flow continues with index-based lookup in the pcf_binding_v2 table.

14. If no lookup records are found in the lookup table, BSF Management pegs the same failure metrics and continues with index-based lookup in the pcf_binding_v2 table.
15. If lookup tables are disabled, BSF Management skips lookup-table processing and continues with index-based lookup in the pcf_binding_v2 table.



1. If a lookup record matches the search criteria, BSF Management queries the pcf_binding_v2 table using the Binding ID from the matching lookup record.
2. If a binding is found in the pcf_binding_v2 table, the database returns the PCF binding records to BSF Management.
3. BSF Management sorts the PCF binding records using the Created Date Time timestamp.
4. BSF Management retrieves the most recent PCF binding.
5. BSF Management returns the PCF binding record with 200 OK to the BSF Diameter Gateway.
6. The BSF Diameter Gateway forwards the Rx AAR-I request to the PCF.
7. If no binding is found in the pcf_binding_v2 table using the Binding ID from the matching lookup record, BSF Management queries the pcf_binding_v2 table using index-based lookup.
8. If the index-based lookup finds a binding in the pcf_binding_v2 table, the database returns the PCF binding records to BSF Management.
9. BSF Management sorts the PCF binding records using the Created Date Time timestamp.
10. BSF Management retrieves the most recent PCF binding.
11. BSF Management pegs the ocbsf_index_based_lookup_fallback metric with found status.
12. BSF Management returns the PCF binding record with 200 OK to the BSF Diameter Gateway.
13. The BSF Diameter Gateway forwards the Rx AAR-I request to the PCF.
14. If no binding is found in the pcf_binding_v2 table through index-based lookup, BSF Management pegs the ocbsf_index_based_lookup_fallback metric with not_found status.
15. The flow continues with fallback lookup in the pcf_binding table.




1. If the `pcf_binding_v2` lookup does not return a binding, BSF Management queries the `pcf_binding` table using index-based lookup.
2. If a binding is found in the `pcf_binding` table, the database returns the PCF binding to BSF Management.
3. BSF Management sorts the PCF binding records using the Created Date Time timestamp, when multiple records are returned.
4. BSF Management retrieves the most recent PCF binding.
5. BSF Management returns the PCF binding record with 200 OK to the BSF Diameter Gateway.
6. The BSF Diameter Gateway forwards the Rx AAR-I request to the PCF.
7. If no binding is found in the `pcf_binding` table, BSF Management returns 204 NO CONTENT to the BSF Diameter Gateway.
8. The BSF Diameter Gateway returns an Rx AAA-I response to the AF.

PCF_BINDING_V2_ONLY (PCF_BINDING_TABLE_LOOKUP == 3)

Note

The following diagrams are part of a single call flow which has been divided into separate sections.



1. The AF sends an Rx AAR-I request to the BSF Diameter Gateway.

2. The BSF Diameter Gateway sends a Discover Request to BSF Management.
3. BSF Management defines the search criteria for the discovery request.
4. BSF Management extracts the UE address identifier from the search criteria.
5. If the lookup table corresponding to the extracted UE address identifier is enabled, BSF Management queries the corresponding lookup table, such as the `ipv6_prefix_table`.
6. If lookup records are found in the lookup table, the database returns the lookup records to BSF Management.
7. BSF Management pegs the `ocbsf_alt_key_lookup` metric with success status.
8. BSF Management sorts the lookup records from most recent to oldest.
9. BSF Management filters the lookup records using the search criteria.
10. If one or more lookup records match the search criteria, the flow continues with a `pcf_binding_v2` lookup using the Binding ID from the matching lookup record.
11. If no lookup record matches the search criteria, the flow continues with index-based lookup in the `pcf_binding_v2` table.
12. If no lookup records are found in the lookup table, BSF Management pegs the `ocbsf_alt_key_lookup` metric with failure status and `not_found` cause.
13. The flow continues with index-based lookup in the `pcf_binding_v2` table.
14. If lookup tables are disabled, BSF Management skips lookup-table processing and continues with index-based lookup in the `pcf_binding_v2` table.



1. If a lookup record matches the search criteria, BSF Management queries the pcf_binding_v2 table using the Binding ID from the matching lookup record.

2. If a binding is found in the `pcf_binding_v2` table, the database returns the PCF binding record to BSF Management.
3. BSF Management returns the PCF binding record with 200 OK to the BSF Diameter Gateway.
4. The BSF Diameter Gateway forwards the Rx AAR-I request to the PCF.
5. If no binding is found in `pcf_binding_v2` using the Binding ID from the lookup record, BSF Management queries `pcf_binding_v2` using index-based lookup.
6. If index-based lookup finds a binding in `pcf_binding_v2`, the database returns the PCF binding record to BSF Management.
7. BSF Management pegs the `ocbsf_index_based_lookup_fallback` metric with found status.
8. BSF Management returns the PCF binding record with 200 OK to the BSF Diameter Gateway.
9. The BSF Diameter Gateway forwards the Rx AAR-I request to the PCF.
10. If index-based lookup does not find a binding in `pcf_binding_v2`, BSF Management pegs the `ocbsf_index_based_lookup_fallback` metric with not_found status.
11. BSF Management returns 204 NO CONTENT to the BSF Diameter Gateway.
12. The BSF Diameter Gateway returns an Rx AAA-I response to the AF.

Upgrade and Rollback Impact

Upgrade:

- Upon upgrade from an earlier BSF release, the `pcf_binding_v2` table and the three supported lookup tables are created, and BSF can use the configurations associated with this feature.
Note: Data migration does not occur during upgrade. Binding records are migrated through lazy migration using the Audit mechanism.
- When adding a new site, users performing a bulk import to configure the system must verify that the following BSF service Advanced Settings are configured before running traffic:
 - `ENABLE_PCF_BINDING_TABLE_MIGRATION=true`
 - `PCF_BINDING_TABLE_LOOKUP=3`

Rollback:

Upon rollback from the current BSF release, the new tables are removed and BSF can no longer use the configurations associated with this feature.

Note: If the rollback target BSF release does not support alternate lookup tables, entries in those tables must be cleaned up manually. Binding records stored only in the updated binding table are not retained after rollback to an earlier release.

Managing the Support for Optimizing Lookup in BSF Management Service Database

Enable

Support for Optimizing Lookup in BSF Management Service Database can be enabled using CNC Console.

Enable Using CNC Console

Enable the **Enable** field the **Alternate Lookup Table Audit Settings** section on the **Management Service** page under **Service Configurations** on CNC Console. For more information, see [Management Service](#).

Configure

Support for Optimizing Lookup in BSF Management Service Database can be configured using CNC Console.

Configure Using CNC Console

Configure the following fields under **Alternate Lookup Table Audit Settings** section on the **Management Service** page for Service Configurations on CNC console:

- Minimum Time to Live (in minutes)
- Maximum Time to Live (in minutes)
- Notification Rate (per second)
- Minimum Audit Passes Interval (in minutes)

For more information, see [Management Service](#).

Configure using Custom Values and Deployment YAML

Percentiles of Latency metrics can be configured using Custom Values and Deployment YAML. For more information, see Configuring Database Indexing in *Binding Support Function Installation, Upgrade, and Fault Recovery Guide*.

The following BSF Management Service metrics are used for this feature:

- ocbsf_feature_incompatibility
- ocbsf_binding_record_lookup_latency
- ocbsf_binding_table_lookup_total
- ocbsf_binding_query_processing_latency
- ocbsf_alt_key_lookup_total
- ocbsf_index_based_lookup_fallback_total
- ocbsf_lookup_table_reinsertions_total
- ocbsf_lookup_table_stale_deletions_total
- ocbsf_db_latency_seconds_count

DB Latency Metrics are expected in the following scenarios:

DB Operation	Situation
insert	Insert operations for Alt Lookup Tables are only expected to happen when: • Registering a PCF Binding V2 table (Register Call-flow) (Alt Lookup Tables must be enabled) • Re-inserting a missing Lookup Record (Audit Notification for V2 Record - 2xx / 5xx / Timeout scenario) (Alt Lookup Tables must be enabled)

DB Operation	Situation
delete_by_binding_id_and_ue_addr	Delete by Binding ID and UE Addr operations for Alt Lookup Tables are only expected to happen when: - Deregistering a PCF Binding V2 record (Register Call-flow) (Alt Lookup Tables must be enabled) - Deregistering a Duplicate PCF Binding V2 record (Collision Detection / Binding Revalidation Call-flow) (Alt Lookup Tables must be enabled) - Deleting a Stale PCF Binding V2 record (Audit Notification for V2 Record - 4xx / MaxTTL scenario) (Alt Lookup Tables can be enabled or disabled) - Deleting a Stale Lookup Record (Audit Notification for Alt Lookup Record - Table disabled / Missing V2 / MaxTTL scenario) (Alt Lookup Tables can be enabled or disabled)
delete_by_binding_id	Delete by Binding ID for Alt Lookup Tables are only expected to happen when: Session Viewer API to delete a session is called
update_audit_time_by_binding_id_and_ue_addr	Update Audit Time (last_audited_date_time and last_access_timestamp) by Binding ID and UE Addr operations for Alt Lookup Tables are only expected to happen when: - Updating Audit Time of V2 Record (Audit Notification for V2 Record - 2xx / 5xx / Timeout scenario) (Alt Lookup Tables must be enabled). - Updating Audit Time of Alt Lookup Record (Audit Notification for Alt Lookup Record - MinTTL scenario) (Alt Lookup Tables must be enabled).
select_by_ue_addr	Select By UE Addr for Alt Lookup Tables are only expected to happen when: - In Discovery requests, when trying to search for a V2 record first via Alt Lookup Tables (Alt Lookup Tables must be enabled). - In Register requests with Collision Detection enabled and Alternate Lookup Tables enabled, to identify the duplicated binding IDs and obtain the most recent matching V2 records which could be used according to the scenario of the collision detected.
select_by_binding_id_and_ue_addr	Select By Binding ID and UE Addr for Alt Lookup Tables are only expected to happen when: Checking if there is a need to re-insert a missing Lookup Record (Audit Notification for V2 Record - 2xx / 5xx / Timeout scenario) (Alt Lookup Tables must be enabled).

For more information, see [BSF Management Service](#).

Logging

The following are the log entries when a record is migrated.

Note

- BSF Management must be in the appropriate log level for these logs to appear
- For the sake of simplicity, any object representation will be written as {...} in the below logs.

- DISCOVERY

- A Lookup records matches the given Search Criteria

```
..., "level": "DEBUG", ..., "message": "Lookup record found for given
Search Criteria. lookupRecord: {...}, searchCriteria: {...}", ...
```

- No Lookup record matches the Search Criteria

```
..., "level": "INFO", ..., "message": "No Lookup records of Lookup
Table {...} match the given Search Criteria. Falling back to Index-
Based lookup. searchCriteria: {...}", ...
```

- No Lookup records were found on the Lookup Table

```
..., "level": "INFO", ..., "message": "No Lookup records were found for
the Lookup Table {...}. Falling back to Index-Based lookup.
searchCriteria: {...}", ...
```

- REGISTER (including Collision Detection):

- A lookup record has been inserted in a Lookup Table

```
..., "level": "INFO", ..., "message": "Successfully registered lookup
record in Lookup Table {...}. lookupRecord: {...}", ...
```

- A lookup record could not be inserted in a Lookup Table

```
..., "level": "WARN", ..., "message": "An error occurred while trying
to register lookup record in Lookup Table {...}. lookupRecord: {...},
error: {...}", ...
```

- DEREGISTER:

- A lookup record has been deleted in a Lookup Table

```
..., "level": "INFO", ..., "message": "Successfully deregistered lookup
record in Lookup Table {...}. lookupRecord: {...}", ...
```

- A lookup record could not be deleted in a Lookup Table

```
..., "level": "WARN", ..., "message": "An error occurred while trying
to deregister lookup record in Lookup Table {...}. lookupRecord: {...},
error: {...}", ...
```

- AUDIT:

- Audit Notification has been received for a lookup record


```
..., "level": "DEBUG", ..., "message": "Audit Notification received for
lookup record. staleRecordInfo: {...}", ...
```
- Lookup record is already present in lookup table for an audited PCF Binding session


```
..., "level": "DEBUG", ..., "message": "Lookup record is already
present in Lookup Table {...}, no need to register it again.
lookupRecord: {...}", ...
```
- Lookup record is not present in lookup table for an audited PCF Binding session


```
..., "level": "INFO", ..., "message": "No lookup record was found for
Binding ID {...} in Lookup Table {...}, will attempt to register
it", ...
```
- Lookup record is still an active


```
..., "level": "DEBUG", ..., "message": "Found an active PCF Binding
record for audited lookup record. lookupRecord: {...}, pcfBinding:
{...}", ...
```
- Lookup record is stale


```
..., "level": "INFO", ..., "message": "No active PCF Binding record was
found for audited lookup record, proceeding with deletion of stale
record. lookupRecord: {...}", ...
```

3.45 Diameter Gateway Connection Repaving

In large Diameter deployments, connection distribution across Diameter Gateway (DGW) pods can become uneven over time due to node restarts, fail overs, scaling operations, or network events. This imbalance can result in some DGW pods handling a disproportionately high number of Diameter connections while others remain underutilized.

Diameter Gateway Connection Repaving feature provides a controlled mechanism to re-balance Diameter connections across the cluster without restarting services or causing a complete outage. It allows to force the disconnection of the existing external Diameter connections so that external peers reconnect through the Layer-4 load balancer and are re-distributed across available Diameter Gateway pods.

The feature supports:

- Controlled and staggered disconnection of Diameter connections
- Continuous acceptance of new incoming connections during repaving
- Historical auditing of repaving operations
- Configurable throttling and jitter controls
- Enable/disable control through configuration settings

The repaving process uses two levels of staggering:

- Pod-Level Jitter: Prevents all pods from disconnecting simultaneously.

- Connection-Level Throttle: Prevents flooding peers with mass disconnects

The repaving process can be triggered through either:

- CNC Console
- Diameter Repaving REST API

When the repaving process is triggered:

1. The operator, automation client, or CNC Console invokes the Diameter Repaving REST API on CM Service. CM Service validates the request, generates `invocationId` and `requestTimestamp`, updates the Config Server trigger topic, and records the trigger in `historical_data`. Diameter Gateway pods detect the Config Server topic update (`common.public.connectionrepaving.diameter-gateway.trigger.disconnect`) and executes the repaving operation when the feature is enabled.
2. Writes an audit entry into the `historical_data` table.
3. Diameter Gateway pods monitor this topic.
4. Upon detecting the update, if **Enable Diameter Gateway Repaving** field is enabled on **Diameter Connection Repaving** page under **Diameter Configuration** on CNC Console:
 - Each pod calculates a random stagger delay. This prevents simultaneous disconnections across pods.
 - After the jitter delay expires, existing external Diameter connections established before the trigger timestamp are disconnected. The connections are disconnected sequentially. A configurable delay is applied between disconnect operations.
 - New connection requests are continued to be accepted throughout the repaving process, ensuring service continuity.
5. The pod performs the throttled disconnection of connections established before the trigger timestamp. Once disconnections are complete, it marks processing for the trigger as finished and sets `isRepavingBusy` to `false`. The pod checks its internal queue and returns to idle monitoring when the queue is empty.
6. External Peers, having their old connections terminated, will automatically attempt to re-establish their connections (if they haven't already done so with new connections accepted during the process).
7. These new connection attempts are accepted by the Diameter Gateway pods.
8. Diameter Gateway processing completes after the pod has handled all eligible external Diameter connections established before the trigger timestamp and completed any applicable queued trigger processing. External peers subsequently reconnect, and the Layer-4 load balancer redistributes the new connections. The resulting connection distribution must be verified using external monitoring.

Note

The Diameter Gateway side repaving completes after eligible connections are disconnected. The final balancing depends on peer re-connection behavior and load balancer distribution.

Call Flows

Figure 3-53 Diameter Gateway Connection Termination





Figure 3-54 Overlapping Repaving Triggers with Queue



Scenario 1: Successful Connection Repaving Execution

1. Initiation (Operator Action & CM Service Update)

A request is sent to the CM Service's dedicated repaving API endpoint (such as `{apiRoot}/oc-bsf-configuration/v1/diameterRepaving/triggerDisconnection`).

The CM Service validates the request, generates a unique `InvocationId` and `requestTimestamp`.

The CM Service updates the configuration topic (`common.public.connectionrepaving.diameter-gateway`) in the Config Server by

replacing the trigger information with the latest one, holding the latest invocationId and requestTimestamp.

The CM Service logs a single entry to the `historical_data` table with `type="diameterRepavingTrigger"`, including the invocationId and requestTimestamp in the value field.

2. Configuration Change Detection (Diameter Gateway Pods)

Each Diameter Gateway pod's configuration client, which is monitoring the `common.public.connectionrepaving.diameter-gateway.trigger.disconnect` topic, detects the update (the newly added entry).

The config client invokes the `onChange` (or equivalent) callback method within the Diameter Gateway service, passing the updated list of invocation records.

3. Two-Tiered Disconnection Execution (Diameter Gateway Pods)

a. Trigger Validation: Inside the `onChange` handler, each pod:

Retrieves the latest invocation record from the received data.

Compares its invocationId with the ID of the last trigger it processed (stored locally).

If the ID is new: Checks if the feature is enabled

(`common.public.connectionrepaving.diameter-gateway.config.enabled` is `true`)

and if the minimum re-invocation delay

(`DGW_CONNECTION_REPAVING.REINVOCATION_DELAY_MIN`) has passed.

b. Tier 1 (Pod-Level Jitter): If validation passes, the pod:

Stores the new invocationId as the last processed ID.

Independently calculates a random stagger delay based on

`DGW_CONNECTION_REPAVING.MAX_STAGGER_JITTER_MS`.

`Wait Time = random(0, MAX_STAGGER_JITTER_MS)`

Pauses execution for this unique, random duration (for example, Pod-A might wait 17 ms, Pod-B might wait 142 ms). Logs that it is waiting.

c. Service Continuity: During this entire process (both the initial wait and the disconnection phase), all Diameter Gateway pods remain active and accepts new incoming connections from peers.

d. Tier 2 (Connection-Level Throttle): After its random pod-level delay expires, the pod begins its disconnection logic.

It iterates through all its existing external Diameter connections that were established before the requestTimestamp associated with the current invocationId.

The order of disconnection is not based on load or complex logic.

To prevent flooding the peer, the pod pauses between terminating each connection for the interval specified in

`DGW_CONNECTION_REPAVING.CONNECTION_DISCONNECT_INTERVAL_MS`.

4. Sequential Trigger Processing with Internal Queue: To enhance robustness in scenarios where repaving might be triggered very frequently (potentially faster than a pod's disconnection cycle + `DGW_CONNECTION_REPAVING.REINVOCATION_DELAY_MIN`), a queue and the `isRepavingBusy` flag are implemented In this model:

The flag `isRepavingBusy` determines if a given repaving trigger is completed or it is still in progress, changing the status (`true/false`) when it starts and finishes.

Each Diameter Gateway pod maintains an internal queue for incoming repaving triggers (represented by their invocationId and requestTimestamp).

When a new trigger is detected using `onChange`, is added to the queue if the pod is currently busy processing a previous trigger.

Once a pod completes a repaving cycle the `isRepavingBusy` flag is set to `false`. It checks its internal queue. If the queue is not empty, it dequeues the next trigger and begins processing it. After validating the `DGW_CONNECTION_REPAVING.REINVOCAATION_DELAY_MIN` against the dequeued trigger's timestamp relative to the previously completed one.

This ensures that the rapid triggers are not lost or rejected but are processed sequentially, preserving the order of operations while still preventing simultaneous execution within a single pod. This approach offers greater resilience against very high trigger rates.

5. Completion and Recovery:

External Peers, having their old connections terminated, automatically attempts to re-establish their connections (if they are not already done so with new connections accepted during the process).

These new connection attempts are accepted by the Diameter Gateway pods.

The new connections are naturally re-balanced across the cluster by the Layer-4 load balancer, completing the repaving process.

Scenario 2: New Connection Attempt During Force Disconnection

This scenario describes what happens when an external system tries to connect while a pod is actively repaving.

1. Upon receiving the configuration change event, the Diameter Gateway pod starts the disconnection of all external connections established prior to the timestamp included in the event.
2. An external peer sends a new connection request (such as a Diameter CER message) to the Diameter Gateway pod.
3. Diameter Gateway pod continues accepting external connections during force disconnections.
4. The external peer likely retry the connection after a back off period, by which time the Diameter Gateway pod may have completed its process.

Scenario 3: Diameter Repaving is Triggered While this Feature is Disabled

This scenario covers the safety mechanism provided by the feature flag.

1. An API request to the CM Service to trigger Diameter Repaving.
2. Configuration topic is updated in config server database.
3. The Diameter Gateway configuration watcher detects the configuration change.
4. It checks if the feature is enabled based on the value in `common.public.connectionrepaving.diameter-gateway.config` topic and finds that it is set to `false`.
5. The client ignores the state change and does not execute the diameter repaving.
6. The process stops immediately. The Diameter Gateway pods are never notified and continue to operate normally without any interruption.

Scenario 4: Overlapping Repaving Triggers with Queuing

This scenario describes how a Diameter Gateway pod handles a new repaving trigger that arrives while it is already processing a previous one, utilizing an internal queue.

1. Initial State: Diameter Gateway Pod-A is idle (not currently repaving) and is monitoring the `common.public.connectionrepaving.diameter-gateway` topic. Its last processed `invocationId` is null or some older ID.
2. Trigger 1 Initiation: A repaving request is initiated through the API. The CM Service updates the config topic with `{ "invocationId": "trigger-UUID-1", "requestTimestamp": "T1" }`. CM Service logs the trigger in `historical_data`.
3. Diameter Gateway Pod-A Detects & Starts Trigger 1:
Pod-A's config client detects the change. The `onChange` handler runs.
It compares "trigger-UUID-1" (latest) with its last processed ID (old/null).
Feature flag and re-invocation delay check is passed.
Pod-A stores "trigger-UUID-1" as its last started ID.
Pod-A calculates its random jitter delay (such as Jitter-1).
Pod-A enters a "busy" state (for example, sets an internal `isRepavingBusy` flag to true) and begins the Jitter-1 wait period.
4. Trigger 2 Initiation (During Pod-A's Busy State): Before Pod-A completes its disconnection cycle for Trigger 1 (either during the Jitter-1 wait or the subsequent throttled disconnection phase), the operator triggers repaving again. The CM Service updates the config topic, appending `{ "invocationId": "trigger-UUID-2", "requestTimestamp": "T2" }`. CM Service logs this second trigger.
5. Diameter Gateway Pod-A detects trigger 2:
Pod-A's config client detects the second change (the list now contains both triggers). The `onChange` handler runs again.
It compares "trigger-UUID-2" (latest) with its last started ID ("trigger-UUID-1").
It checks its internal state and finds that `isRepavingBusy` flag is set to true.
Instead of immediately processing it, Pod-A adds Trigger 2 ("trigger-UUID-2", T2) to its internal queue. It does not calculate jitter or start waiting for Trigger 2 yet.
6. Diameter Gateway Pod-A Completes Trigger 1:
Pod-A finishes the Jitter-1 wait.
It performs the throttled disconnection of connections established before T1.
Once the disconnections for Trigger 1 are complete, it marks processing for "trigger-UUID-1" as finished and sets `isRepavingBusy` flag back to false.
7. Diameter Gateway Pod-A Dequeues & Starts Trigger 2:
Pod-A checks its internal queue and finds Trigger 2. It dequeues it.
It performs the validation checks for Trigger 2:
 - Feature flag check.
 - Re-invocation delay check: Compares T2 against the completion time of Trigger 1 using `REINVOCATION_DELAY_MIN`. The check is passed.Pod-A stores "trigger-UUID-2" as its last started ID.
Pod-A calculates its random jitter delay (Jitter-2).
Pod-A sets `isRepavingBusy` flag to true and begins the Jitter-2 wait period.
8. Diameter Gateway Pod-A Completes Trigger 2:
Pod-A finishes the Jitter-2 wait.

It performs the throttled disconnection of connections established before T2 (but after T1, if any were established and accepted in between).

Once disconnections for Trigger 2 are complete, it marks processing for "trigger-UUID-2" as finished, sets `isRepavingBusy` flag to false.

Pod-A checks its internal queue again; it's empty. The pod returns to idle monitoring.

Managing Diameter Gateway Connection Repaving

Enable

This feature can be enabled using CNC Console and REST API.

Enable Using CNC Console

To enable this feature using CNC Console, enable **Enable Diameter Gateway Repaving** field on **Diameter Connection Repaving** page under **Diameter Configuration** on CNC Console.

For more information, see [Diameter Repaving](#).

Enable Using REST API

To enable this feature using REST API, use Diameter Repaving API.

The URI to access the Diameter Repaving API is `{apiRoot}/oc-bsf-configuration/v1/diameterRepaving/config`.

Configure

This feature can be configured using CNC Console and REST API.

Configure Using CNC Console

Configure the following advanced settings keys under **Settings** page for **Diameter Configurations** on CNC Console:

- `DGW_CONNECTION_REPAVING.MAX_STAGGER_JITTER_MS`
- `DGW_CONNECTION_REPAVING.REINVOCATION_DELAY_MIN`
- `DGW_CONNECTION_REPAVING.CONNECTION_DISCONNECT_INTERVAL_MS`
- `DGW_CONNECTION_REPAVING.DPR_DELAY_MS`
- `DGW_CONNECTION_REPAVING.DPR_DISCONNECT_CAUSE`

For more information, see [Settings](#).

Configure Using REST API

This feature can be enabled and configured using the Diameter Repaving API.

Use the `{apiRoot}/oc-bsf-configuration/v1/diameterRepaving/config` to configure the feature.

Use the `{apiRoot}/oc-bsf-configuration/v1/diameterRepaving/triggerDisconnection` to disconnect the existing external Diameter connections established before the trigger timestamp.

Also, configure the following advanced settings keys on Diameter Settings API using the following URI: `{apiRoot}/oc-bsf-configuration/v1/diameter/settings`:

- `DGW_CONNECTION_REPAVING.MAX_STAGGER_JITTER_MS`
- `DGW_CONNECTION_REPAVING.REINVOCATION_DELAY_MIN`

- `DGW_CONNECTION_REPAVING.CONNECTION_DISCONNECT_INTERVAL_MS`
- `DGW_CONNECTION_REPAVING.DPR_DELAY_MS`
- `DGW_CONNECTION_REPAVING.DPR_DISCONNECT_CAUSE`

Note

The default repaving values in these advanced settings keys are likely implicit unless they are overridden.

Observability

Metrics

The `diam_repaving_disconnections_total` metric is to count the total number of connections terminated by the "force disconnect" mechanism. For more information, see [Diameter Gateway Metrics](#).

Note

- The success of the re-balancing relies on the Layer-4 Load Balancer distributing new connection attempts effectively across the available pods after the disconnections occur. Common load balancing methods like "round robin" or "least connections" are expected to facilitate this re-balancing, assuming external clients promptly attempt to reconnect.
- It is assumed that the external Diameter clients automatically attempt to reconnect following the forced disconnection. The success of the connection repaving relies on this client behavior, even if not all connections are immediately re-established due to varying client load or specific retry logic.
- External monitoring tools, such as KPIs must be used to detect connection imbalance.
- This feature is not autonomous and requires explicit operator intervention to initiate.
- The solution's effectiveness requires the Diameter Gateway pods to be up and running. While it works effectively with direct load balancer-to-pod routing and is expected to work with service meshes like Istio with ASM sidecars (acting as TCP proxies), certain Kubernetes network configurations using kube-proxy might limit the overall re-balancing effectiveness. This is because each kube-proxy instance on different nodes can independently load balance connections across the Diameter Gateway pods, potentially leading to less predictable or less globally optimized connection distribution after the repaving process completes.
- The disconnection action triggers the termination of all existing external inbound connections on each pod, rather than calculating a minimal, targeted number of connections to drop.

3.46 Security Context

The Security Context feature improves the security posture of deployed workloads in Oracle Communications Cloud Native Core, Binding Support Function (BSF), by applying Kubernetes pod and container security context settings.

A security context defines privilege and access-control settings for a pod or container. These settings control how containers run, which user and group IDs are used, whether privilege escalation is allowed, which Linux capabilities are available, and whether the root file system is writable.

This feature applies the SA-approved deployment security-context recommendations to BSF services, Common Services consumed by BSF, and one-click BSF deployment values. The implementation supports the following deployment models:

- OC-CNE-based deployments.
- Non-OC-CNE deployments where the customer has no sidecar injections, or sidecars do not require special or elevated privileges.
- Non-OC-CNE deployments where the customer requires sidecar injections with different or elevated privileges.

The Security Context feature is supported for Converged Policy services and related services, including:

- binding-service
- bulwark
- cmservice
- config-server
- diam-connector
- diam-gateway
- ldap-gateway
- notifier
- nwdaf-agent
- pcf-amservice
- pcf-smsservice
- pcf-ueservice
- pcf-userservice
- pcrf-core
- perf-info
- policy-runtime
- policyds
- queryservice
- session-state-audit
- soapconnector
- usage-mon
- app_info

The Security Context feature supports the following default configuration:

```
enableContainerSecurityContext: true
enablePodSecurityContext: false
```

```
podSecurityContext:
  runAsNonRoot: true
  runAsGroup: 25
  runAsUser: 25

containerSecurityContext:
  runAsNonRoot: true
  readOnlyRootFilesystem: true
  allowPrivilegeEscalation: false
  privileged: false
  capabilities:
    drop:
      - ALL
```

The container security context is enabled by default. The pod security context is available but disabled by default to avoid compatibility issues with deployments that use sidecar containers or other pod-level requirements.

Managing Security Context Feature

Enable

You can enable the Security Context Feature using Helm Configuration:

Enable Container Security Context:

You can enable or disable the feature by setting the `enableContainerSecurityContext` parameter to `true` or `false` in the `custom-values.yaml` file. The container security context is enabled by default.

Enable Pod Security Context:

You can enable or disable the feature by setting the `enablePodSecurityContext` parameter to `true` or `false` in the `custom-values.yaml` file. The pod security context is disabled by default.

Service-Specific Configuration:

Use service-specific configuration when a service requires a different user ID, group ID, or security context value. For example:

```
config-server::
  podSecurityContext:
    runAsNonRoot: true
    runAsUser: 10001
    runAsGroup: 10001
  containerSecurityContext:
    runAsNonRoot: true
    allowPrivilegeEscalation: false
    privileged: false
    readOnlyRootFilesystem: true
    capabilities:
      drop:
        - "ALL"
```

Deployment Template:

The deployment template applies the pod and container security context settings when the corresponding values are enabled.

```
apiVersion: apps/v1
kind: Deployment
metadata:
  ...
spec:
  ...
  template:
    spec:
      ...
      {{- if .Values.global.setHostUsers }}
      hostUsers: false
      {{- end }}
      {{- if .Values.enablePodSecurityContext }}
      securityContext:
        {{- toYaml .Values.podSecurityContext | nindent 8 }}
      {{- end }}
      containers:
        - name: {{ .Chart.Name }}
          ...
          {{- if .Values.enableContainerSecurityContext }}
          securityContext:
            {{- toYaml .Values.containerSecurityContext | nindent 12 }}
          {{- end }}
```

Observability

There are no new Metrics, KPIs or alerts available for this feature.

4

Configuring BSF Using CNC Console

This chapter describes how to configure different global and service parameters in Oracle Communications Cloud Native Core Binding Service Function (BSF) using Oracle Communications Cloud Native Configuration Console (CNC Console).

Oracle Communications Cloud Native Configuration Console (CNC Console)

This section provides an overview of the CNC Console, which includes an interface to help in creating global and service parameters in BSF.

You can use BSF integration with CNC Console only after logging successfully in to the CNC Console application. To log in to the CNC Console, make the following updates to the hosts file available at the C:\Windows\System32\drivers\etc location.

1. In Windows system, open the **hosts** file in a notepad as an Administrator and append the following set of lines at the end:

```
<IP Address> cncc-iam-ingress-gateway.cncc.svc.cluster.local
```

```
<IP Address> cncc-core-ingress-gateway.cncc.svc.cluster.local
```

where:

<IP Address> is the host address of the deployment cluster. It depends on the deployment cluster.

Example:

```
10.75.225.189 cncc-iam-ingress-gateway.cncc.svc.cluster.local
```

```
10.75.225.189 cncc-core-ingress-gateway.cncc.svc.cluster.local
```

Note

The IP Address can change when deployment cluster changes.

2. Save and close the hosts file.

Note

Before logging into CNC Console, create a CNC user and password. Using these user details, you can log in to the CNC Console application. For more information about creating a CNC Console user and password, see *Oracle Communications Cloud Native Configuration Console (CNC Console) Installation, Upgrade, and Fault Recovery Guide*.

To log in to CNC Console:

1. Open a web browser and enter the URL: `http://cncc-core-ingress-gateway.cncc.svc.cluster.local:port number/` and press Enter.


Note

port number is cncc-iam-ingress-port number.

The login page opens.

2. Enter the **Username** and **Password**.
3. Click **Log In**.
4. On the Welcome page, select the required NF instance from the **Please Select Instance** drop-down field.
This opens the CNC Console home page for the selected NF instance:

Figure 4-1 CNC Console for BSF



5. To use BSF services integrated with CNC Console, click **BSF** in the left navigation pane.

4.1 General Configurations

This section describes the general configurations for Oracle Communications Cloud Native Core Binding Support Function (BSF).

To access the General Configurations page from the CNC Console home page, click **BSF**, and then click **General Configurations**.

It consists of the following two sections:

- [General Settings](#)
- [SBI Error Codes](#)

4.1.1 General Settings

This section describes the general settings, which can be configured using Cloud Native Core Console for Binding Support Function (BSF).

To access the General Settings page from the CNC Console home page, click **BSF**, select **General Configurations**, and click **General Settings**.

To edit the existing general settings, perform the following steps:

1. Click  **Edit** .
This opens the **Edit General Settings** page.
2. Enter the values for the following fields:

Table 4-1 General Settings

Field Name	Description
Enable Tracing	Specifies whether to enable or disable tracing for the BSF deployment. By default, this configuration is enabled.
Enable Metrics	Specifies whether to enable or disable system metrics for the BSF deployment. By default, this configuration is enabled.
Enable Collision Detection	Specifies whether to enable or disable collision detection for the BSF deployment. By default, this configuration is disabled.
Binding Discovery Request Timeout	Specifies the request timeout value for the discovery request sent by the BSF Diameter Gateway towards the BSF Management Service.
Enable Subscriber Activity Logging	Specifies whether to enable or disable subscriber activity logging. The default value is <code>false</code> .
Enable SBI Correlation	Specifies whether to enable or disable correlation-info header in BSF. The default value is <code>false</code> .

Table 4-2 General Settings - Enhanced Logging Configuration

Field Name	Description
Enable Enhanced Logging	Specifies whether to enable or disable enhanced logging for the BSF deployment. Default value: <code>false</code>

Table 4-2 (Cont.) General Settings - Enhanced Logging Configuration

Field Name	Description
Log Level for Stale Session Deletions	Specifies the log level for stale session deletions. The available log levels for stale session deletions are: • TRACE • DEBUG • INFO • WARN • ERROR The logs for BSF stale binding deletion are generated at the specified level. Note: The value of this field does not modify in any way the log level that BSF Management service has. The root log level of BSF Management service must be either the same log level defined in the selected or lower for these deletion logs to appear. For example, if the stale session deletions log level is at INFO, then BSF Management service log level must be either TRACE, DEBUG, or INFO. Default value: INFO.
Enable UE Identifier Information	Specifies whether to enable or disable UE Identifier information for the BSF deployment. Default value: false
Enable UE Identifier Information in Stale Session Deletion Logic	Specifies whether to enable or disable UE Identifier information in stale session deletion logging. When this field is disabled, the UE Identifier information in the logs are masked. That is, the identifier information appears like 'xxxx'. Default value: false

3. Click **Save** to save the chosen general settings.

4.1.2 SBI Error Codes Configurations

This section describes how to customize the SBI error codes according to the network requirements using the SBI Error Codes page.

The **SBI Error Codes** page on CNC Console allows users to view and edit conditions defined by default for the Binding Support Function (BSF) network function. This page also provides the options to import and export SBI error codes.

The following table describes the errors supported by BSF:

Table 4-3 Error Codes and Responses

Condition ID and Name	Error Description	HTTP Status Code	Application Error Code
ERROR_CODE_RESOURCE_URI_STRUCTURE_NOT_FOUND Binding resource URI is invalid	This error is returned when bindingId is missing in the request URI or when the request URI is invalid.	404	RESOURCE_URI_STRUCTURE_NOT_FOUND
ERROR_CODE_BINDING_NOT_FOUND Binding key not found	This error is returned when binding data is not found in the database.	404	NOT_FOUND
ERROR_CODE_INVALID_QUERY_PARAM Parameters in binding key are not acceptable	This error is returned when parameters such as multiple UE addresses are sent in the request.	400	INVALID_QUERY_PARAM
ERROR_CODE_MANDATORY_QUERY_PARAM_INCORRECT Required parameter in binding key is not acceptable	This error is returned when the mandatory query parameter is invalid.	400	MANDATORY_QUERY_PARAM_INCORRECT
ERROR_CODE_OPTIONAL_QUERY_PARAM_INCORRECT Optional parameter in binding key is invalid	This error is returned when the optional query parameter is invalid.	400	OPTIONAL_QUERY_PARAM_INCORRECT
ERROR_CODE_MANDATORY_QUERY_PARAM_MISSING Required parameter in binding key is missing	This error is returned when the mandatory query parameter is missing in the URI request.	400	MANDATORY_QUERY_PARAM_MISSING
ERROR_CODE_INVALID_MSG_FORMAT Parameter(s) in binding data is not supported	This error is returned in case of invalid binding data is sent as payload. For instance, when an IP domain is present without an IPv4 address, this error is reported.	400	INVALID_MSG_FORMAT
ERROR_CODE_MANDATORY_IE_INCORRECT Required parameter in binding data is invalid	This error is returned when the mandatory parameter such as SNSSAI is assigned a semantically incorrect value.	400	MANDATORY_IE_INCORRECT

Table 4-3 (Cont.) Error Codes and Responses

Condition ID and Name	Error Description	HTTP Status Code	Application Error Code
ERROR_CODE_MANDATORY_IE_MISSING Required parameter in binding data is missing	This error is returned when the mandatory parameter such as SNSSAI is missing.	400	MANDATORY_IE_MISSING
ERROR_CODE_OPTIONAL_IE_INCORRECT Optional parameter in binding data is invalid	This error is returned when the Information Element (IE) in the request is invalid.	400	OPTIONAL_IE_INCORRECT
ERROR_CODE_INTERNAL_SERVER_ERROR Unexpected server error	This error is returned in the following situations: - Issue in saving data successfully to the database. - URI pattern for nbsf-management/v1/pcfBindings throws IllegalArgumentException.	500	INTERNAL_SERVER_ERROR
ERROR_CODE_METHOD_NOT_ALLOWED Request method not supported	This error is returned when the requested method is not supported. For instance, if the user sends a PUT request instead of a POST request to register PCFBinding, BSF returns error 405 in response to the request.	405	METHOD_NOT_ALLOWED

Editing SBI Error Codes

To edit any of the defined conditions, perform the following steps:

- From the navigation menu, click **BSF**, then select **General Configurations**, and click **SBI Error Codes**.
This opens the SBI Error Codes page that lists the condition names along with their **HTTP Status code** and **Application Error Code**.
- Click Edit against the condition that you need to customize.
This opens the Edit SBI Error Codes page.
- Update the required values for the fields as described in the following table:

Table 4-4 Parameters for Edit SBI Error Codes

Parameter	Description
Error Description	Specifies the description for a defined condition. It is recommended to use descriptions that clearly explain the condition.
HTTP Status Code	Specifies the HTTP Status code for a defined condition. This is a mandatory field and cannot be left blank. Note: Currently, the value for HTTP status code can be selected only from the supported set of values. To view the values, see Table 4-3 .
Application Error Code	Specifies the application error code for a defined condition. Users can customize application error codes as per their requirements. This is a mandatory field and cannot be left blank.

- Click **Save**.

Supported HTTP Status Codes

The following table describes the HTTP status codes that can be used as an input for the **HTTP Status Code** field while configuring the SBI Error codes feature:

Table 4-5 1xx Informational Series

Status Code	Description
100	CONTINUE
101	SWITCHING_PROTOCOLS
102	PROCESSING
103	CHECKPOINT

Table 4-6 2xx Success Series

Status Code	Description
200	OK
201	CREATED
202	ACCEPTED
203	NON_AUTHORITATIVE_INFORMATION
204	NO_CONTENT
205	RESET_CONTENT
206	PARTIAL_CONTENT
207	MULTI_STATUS
208	ALREADY_REPORTED
226	IM_USED

Table 4-7 3xx Redirection Series

Status Code	Description
300	MULTIPLE_CHOICES
301	MOVED_PERMANENTLY

Table 4-7 (Cont.) 3xx Redirection Series

Status Code	Description
302	FOUND
302	MOVED_TEMPORARILY
303	SEE_OTHER
304	NOT_MODIFIED
305	USE_PROXY
307	TEMPORARY_REDIRECT
308	PERMANENT_REDIRECT

Table 4-8 4xx Client Error Series

Status Code	Description
400	BAD_REQUEST
401	UNAUTHORIZED
402	PAYMENT_REQUIRED
403	FORBIDDEN
404	NOT_FOUND
405	METHOD_NOT_ALLOWED
406	NOT_ACCEPTABLE
407	PROXY_AUTHENTICATION_REQUIRED
408	REQUEST_TIMEOUT
409	CONFLICT
410	GONE
411	LENGTH_REQUIRED
412	PRECONDITION_FAILED
413	PAYLOAD_TOO_LARGE
413	REQUEST_ENTITY_TOO_LARGE
414	URI_TOO_LONG
414	REQUEST_URI_TOO_LONG
415	UNSUPPORTED_MEDIA_TYPE
416	REQUESTED_RANGE_NOT_SATISFIABLE
417	EXPECTATION_FAILED
418	I_AM_A_TEAPOT
419	INSUFFICIENT_SPACE_ON_RESOURCE
420	METHOD_FAILURE
421	DESTINATION_LOCKED
422	UNPROCESSABLE_ENTITY
423	LOCKED
424	FAILED_DEPENDENCY
425	TOO_EARLY
426	UPGRADE_REQUIRED
428	PRECONDITION_REQUIRED
429	TOO_MANY_REQUESTS
431	REQUEST_HEADER_FIELDS_TOO_LARGE

Table 4-8 (Cont.) 4xx Client Error Series

Status Code	Description
451	UNAVAILABLE_FOR_LEGAL_REASONS

Table 4-9 5xx Server Error Series

Status Code	Description
500	INTERNAL_SERVER_ERROR
501	NOT_IMPLEMENTED
502	BAD_GATEWAY
503	SERVICE_UNAVAILABLE
504	GATEWAY_TIMEOUT
505	HTTP_VERSION_NOT_SUPPORTED
506	VARIANT_ALSO_NEGOTIATES
507	INSUFFICIENT_STORAGE
508	LOOP_DETECTED
509	BANDWIDTH_LIMIT_EXCEEDED
510	NOT_EXTENDED
511	NETWORK_AUTHENTICATION_REQUIRED

4.1.3 SBI Ingress Error Code Profiles Collection

This procedure provides information about how to use the SBI Ingress Error Code Profiles Collection page to create and manage SBI Ingress error code profiles collection in General Configurations.

To configure Error Code profiles, perform the following steps:

1. From the navigation menu, under **BSF**, click **General Configurations**, and then select **SBI Ingress Error Code Profiles Collection**.
This opens the SBI Ingress Error Code Profiles Collection page.
2. Click Edit.
This opens the Edit SBI Ingress Error Code Profiles Collection page.
3. Click  **Add** .
This opens the Add SBI Ingress Error Code Profiles Collection page.
4. Enter values for the available input fields as described in the following table:

Table 4-10 Error Code Profiles Collection Configurations

Field Name	Description
Name	Specifies a unique name to identify the error profile.
Error Code	Specifies the HTTP Code that is populated in the error response when a message request is rejected due to overload control.

Table 4-10 (Cont.) Error Code Profiles Collection Configurations

Field Name	Description
Error Cause	Specifies the error cause that is populated in the error response when a message request is rejected due to overload control.
Error Title	Specifies the error title that is populated in the error response when a message request is rejected due to overload control.
Error Description	Specifies the error description that is populated in the error response when a message request is rejected due to overload control.

- Click **Save** to save the error code profile.
To discard the changes, click **Cancel**

The value gets listed on the SBI Ingress Error Code Profiles Collection page. Use  or  available under the **Actions** column to update or delete the profile.

4.2 Error Handling

This section describes how to manage and view the error configurations in BSF, using the **Error Handling** Configurations page

4.2.1 Error Configurations

The error handling framework allows the users to configure an error state and an action for it. The action contains two parts, an error rule and an error context. On the Console UI, the operator configures the error state specific to the BSF services and the list of actions for it.

The **Error Configurations** page displays the error configurations related to different BSF services.

To add error handler template for a BSF service:

Error Configurations for Diameter Gateway Service

- From the navigation menu, under **BSF**, click **Error Handling**, and select **Error Configurations** page. This opens the **Error Handling Configurations** page.
- From the Select Service Name drop-down list select the value **diam-gateway**. This page displays diameter message retry configurations for AAR, ASR, RAR and STR messages. On the page **Error Handler Templates of diam-gateway** and **Error Configurations of diam-gateway** subsections are displayed.
- The **Error Handler Templates of diam-gateway** provides two options:
 - Error Code Configuration**
 - Timeout Error Configuration**
- The **Error Configurations of diam-gateway** provides default error handling configurations to retry on all error codes (except diameter result code 2xxx) and timeout for Rx AAR, Rx ASR, Rx RAR and Rx STR failed diameter messages.
- To configure the **Error Code Configuration** in **Error Handler Templates of diam-gateway**, Click  **Edit**. This opens the **Error Handler Template** editing page.

6. Enter values for the available input fields. The following table describes the fields:

Table 4-11 Create Error Code Configuration - Edit

Field Name	Description
On Rx	Specifies the list of diameter interfaces. The values are: • AAR • RAR • ASR • STR
Status	Specifies the Diameter error status pattern to be matched for the selected Rx message type The value must follow these rules: • Must be exactly four characters. • Each character can be a number (0–9) or the wildcard character (*). • The wildcard character (*) matches any single digit. • 2xxx status patterns are not supported because 2xxx Diameter result codes represent success responses. Valid examples: 3002, 3***, 4***, 5***, 30**, 5012. Invalid examples: 2001, 2***, 20**, 300, 30020, 3xx*, 30A*. During upgrade, legacy 2xxx rules configured in earlier releases are retained. However, if such configurations are exported and then imported, the import operation completes with PARTIAL_SUCCESS because 2xxx status patterns are not supported during import.

Table 4-12 Error Cause Configure

Field Name	Description
Error Cause Field	Specifies to search for which error causing filed in the diameter answer message. Default value: ALL
Match Operator	Specifies the match operator to search error section in the diameter answer message. Default value: ANY
Error Response Originator	The peer from which the error origination occurs. User can choose from the following options: • ANY • INTERMEDIATE_PEER • DESTINATION PEER
Message	Specifies the error message to search in the error section of diameter answer message. Default value: ANY

Table 4-12 (Cont.) Error Cause Configure

Field Name	Description
Status	Specifies the Diameter error status pattern to search in the error section of the Diameter answer message. The value must follow these rules: • Must be exactly four characters. • Each character can be a number (0–9) or the wildcard character (*). • The wildcard character (*) matches any single digit. • 2xxx status patterns are not supported because 2xxx Diameter result codes represent success responses. Valid examples: 3002, 3***, 4***, 5***, 30**, 5012. Invalid examples: 2001, 2***, 20**, 300, 30020, 3xx*, 30A*. During upgrade, legacy 2xxx rules configured in earlier releases are retained. However, if such configurations are exported and then imported, the import operation completes with PARTIAL_SUCCESS because 2xxx status patterns are not supported during import. Default value: ANY
Cause	Specifies the field that matches the cause during error ends with 'not found'. User can choose from the following options: • ANY • RESPONSE_TIMEOUT

Table 4-13 Action

Field Name	Description
Action	The action to be performed in the event of failed diameter message on Rx interface. User can choose from the following options: • RETRY TO ALTERNATE PEER • ONE RETRY TO ALTERNATE PEER

- Click **Save** to save the changes.

 Note

Click **Cancel** to discard the changes.

- To configure **Timeout Error Configuration** in **Error Handler Templates of diam-gateway**, Click  **Edit**. This opens the **Error Handler Template** editing page.
- Enter values for the available input fields. The following table describes the fields:

Table 4-14 Create Timeout Error Configuration

Field	Description
On Rx	Specifies the list of diameter interfaces. The values are: • AAR • RAR • ASR • STR Default value: AAR
Status	Specifies the error status to be provided by the user. Default value: ANY

Table 4-15 Error Cause Configure

Field	Description
Error Cause Field	Specifies to search for which error causing filed in the diameter answer message. Default value: MESSAGE
Match Operator	Specifies the match operator to search error section in the diameter answer message. Default value: EQUALS
Message	Specifies the error message to search in the error section of diameter answer message. Default value: TIMEOUT_EXCEPTION

Table 4-16 Action

Field	Description
Action	The action to be performed in the event of response timeout on Rx interface. User can choose from the following options: • RETRY TO ALTERNATE PEER • ONE RETRY TO ALTERNATE PEER

10. Click **Save** to save the changes or click **Cancel** to discard the changes. Using the edit error handling configuration option, you can enable or disable the configurations. The priority for each error handling configuration can be set using the



and



arrows.

Error Configurations for BSF Management Service

1. From the navigation menu, under **BSF**, click **Error Handling**, and select **Error Configurations** page. This opens the **Error Handling Configurations** page.
2. From the Select Service Name drop-down list select the value BSF Management Service. This page displays BSF Management Service.
3. Enable the error handler configurations using the **Enable Error Handler Configurations** toggle button.
4. The error handler template provides *Error Enhancement Configurations*.
5. To configure the *Error Enhancement Configurations* in Error Handler Templates of the required service, Click  **Edit** . This opens the **Error Handler Template** editing page.
6. Enter values for the available input fields. The following table describes the fields:

Table 4-17 Error Handler Template

Field	Description
On	Specifies the Application Error. Default value: Application Error

Table 4-18 Action

Field	Description
Action	Specifies the action to be performed in the event of failed message. Default value: Reject with Enhanced Detail
Exclude from error message	Specifies exclusion of the provided components from detail error message. By default, "Error State and "Problem Cause" are excluded.

7. Click **Save** to save the changes.

 Note

Click **Cancel** to discard the changes.

Importing Error Handling Configurations:

1. Click the **Import** icon.
The **File Upload** dialog box opens.
2. Upload the file in JSON format by using the **Drag and Drop** button.
3. Click **Import**.

4.3 Logging Configurations

This section describes how to customize the log level and subscriber logging activity settings in BSF using the Logging Configurations pages.

4.3.1 Logging Level

This procedure describes how to configure the log level for different services through CNC Console.

Note

The default log level for each service is **Warn**.

The Logging Level page displays the log level configured for different services. The page allows you to edit the log level configurations.

To configure the log level:

1. From the navigation menu, under **BSF**, click **Logging Level**.
This opens the **Logging Level Configuration** page. You can add or edit the log level and package log level for each service type from this page.
2. Click  **Edit**.
This opens the **Edit Log Level** page.
3. From the **Service Type** drop-down list, select the service for which you need to view, edit, or delete the logs.
4. From the **Application Log Level** drop-down list, select the root log level of the application for the selected service type. Possible values are:
 - **TRACE**
 - **DEBUG**
 - **INFO**
 - **WARN**
 - **ERROR**

Note

The value for the **Application Log Level** field is the mandatory value, and the **Package Log Level** is the optional value.

5. Expand the **Package Log Level** group to enter the package log level information:

Note

This step is only applicable when Oracle Engineering is trying to isolate an issue and requests one or more package names be added and logs collected after the reproduction of an issue.

- a. Click  **Add**.
The page opens the **Add Package log Level** dialog box.

- b. Enter the value in the **Package** field. The value of **Package** field is dependent on the package's name in each application. Before you set the value of **Package** field, you need to know what package is existed in that application.
- c. From the **Log Level** drop-down menu, select the log level for the package. Possible values are:
 - **TRACE**
 - **DEBUG**
 - **INFO**
 - **WARN**
 - **ERROR**
- d. Click **Save**.
The Package log level information for the selected service is saved.

 Note

Use  **Edit** or  **Delete** available in the next column to update or delete the package log level information.

6. Click **Save**.
The log level information for the selected service type is saved.

4.3.2 Subscriber Activity Logging

Subscriber Activity Logging allows you to define a list of the subscribers (identifier) that you may require to troubleshoot the NFs and trace all the logs related to the subscribers separately to view. This functionality can be used to troubleshoot problematic subscribers without enabling logs or traces that can impact all subscribers. You can capture and monitor subscriber logs for Binding Register and Deregister call flow between Ingress Gateway and BSF Management Service and Binding Discovery call flow between Diameter Gateway and BSF Management Service.

To enable the subscriber activity logging functionality, set value of the **Enable Subscriber Activity Logging** parameter to **true** on the **General Configurations** page. By default, this functionality remains disabled. For more information about enabling the functionality, see [General Settings](#)

This procedure provides information about how to configure and manage subscriber logging.

The **Subscriber Activity Logging** page allows you to create new and manage existing subscribers. The page displays the list of defined subscribers and provides the options to import, export, or add lists.

You can configure the list of subscribers using the **Subscriber Activity Logging** page.


To configure a list of subscribers for logging:

To configure Subscriber Activity Logging:

1. From the navigation menu under **BSF**, navigate to **General Configurations**, click **Logging Configurations**, and then select **Subscriber Activity Logging**.
This opens the **Subscriber Activity Logging** page. The page lists the existing configurations. You can add or import new subscriber activity logging configurations using this page.

Note

Click  **Export** to download the available listings in the JSON file format on your system.

Figure 4-2 Adding Subscriber Activity Logging

2. Click  **Add** .
This opens the **Create Subscriber Activity Logging** page.
3. On the **Create Subscriber Activity Logging** page, enter the following information:

Table 4-19 Create Subscriber Activity Logging Field Description

Field Name	Description
Identifier Type	Select the subscriber identifier type. Supported subscriber identifier type are: • GPSI • SUPI • IPV4 • IPV6
Identifier Value	The identifier value for the selected identifier type.
Enable	Use this switch to enable or disable the subscriber logging functionality for the selected subscriber.

Figure 4-3 Creation of Subscriber for Logging with Identifier Console

4. Click **Save**.
The configuration gets listed on the **Subscriber Activity Logging** page. The page defines the Subscriber Activity Logging configuration in the BSF database.

Note

Use  or  available under the **Actions** column to update or delete the configuration.

Figure 4-4 Added List of Subscribers

Subscriber Activity Logging			
Type to Filter 		 Refresh	 Import
		 Export	 Add
Identifier Type 	Identifier Value 	Enable 	Actions
SUPI	imsi-1234567	true	 

To import Subscriber Activity Logging configuration:

1. Click  **Import** .
The page opens the **File Upload** dialog box.
2. Upload the file in JSON format by using the **Drag and Drop** button.
3. Click **Import**.

Subscriber Identifiers**SUPI**

In the 5G system, a globally unique Subscription Permanent Identifier (SUPI), known as IMSI (International Mobile Subscriber Identity) till 4G, is assigned for each subscription. The SUPIs are assigned in such a manner that it helps in identifying subscriptions and is independent of the user equipment.

For 4G systems, the value of IMSI is structured as:

imsi: <value>

For 5G systems, the value of SUPI is structured as:

supi: imsi-<value>

GPSI

General Public Subscription Identifier (GPSI), known as MSISDN (Mobile Station International Subscriber Directory Number) till 4G, is a 3GPP defined subscriber public identifier that can be used both inside and outside of the 3GPP system. The association between GPSI and its related SUPI are stored in the subscription data in a 5G system.

For 4G systems, the value of MSISDN is structured as:

msisdn/e164:<value>

For 5G systems, the value of GPSI is structured as:

gpsi: msisdn-<value>

4.4 Service Configurations

This section describes how to customize the BSF Management Service and Audit Service according to the network requirements using the Service Configuration page.

4.4.1 Management Service

You can view and edit configurations for BSF Management Service on the **Management Service** page using the CNC Console.

To access this screen from the Home screen of CNC Console, under **BSF**, click **Service Configurations** and then **Management Service**.

To edit management service configurations, perform the following steps:

1. From the navigation menu, under **BSF**, click **Service Configurations**, and select **Management Service**.

This opens the **Management Service** page.

2. Click  **Edit**.

This opens the **Edit Management Service** page.

3. Update the required values for the parameters as described in the following table:

Table 4-20 Parameters for Edit Management Service Configurations

Parameter	Description
Server Root URL	Specifies the callback URI for notifications to be received by the user.

Table 4-21 Audit

Parameter	Description
Enabled	Use this flag to enable or disable stale session handling feature. By default, the feature is disabled.
Notification Rate (per second)	Specifies the number of notifications that Audit service sends to the BSF Management service in one second. The recommended value is 50. Default and Recommended Value: 50 Note: To configure higher number than the recommended value, contact My Oracle Support (https://support.oracle.com)
Binding Age (in minutes)	Specifies the binding age for binding records. Once the binding age for a record exceeds the configured value, audit service marks the record as suspected stale. The recommended value is 3600.
Maximum Binding Age (in minutes)	Specifies the maximum binding age for binding records. Once the binding age for a record exceeds the configured value, audit service marks the record as stale and BSF removes the record from its local database. The recommended value is 7200.

Table 4-21 (Cont.) Audit

Parameter	Description
Minimum Audit Attempts	Specifies the minimum number of consecutive failed audit attempts until maxTTL / forceTTL is reached. If maxTTL is reached and audit_attempts + 1 >= Minimum Audit Attempts for maxTTL, Audit service sends notification to Management Service with maxTTL flag set to <i>true</i>. Management Service deletes the record. Range: 0-255 Default Value: 0 Note: If maxTTL is not reached and if audit attempts are reached, the number of audit attempts are incremented until maxTTL is reached.
Minimum Audit Passes Interval (in minutes)	Specifies the minimum interval between two consecutive audits. The recommended value is 10.
Answer with Result Code Configuration	Choose the value that BSF compares with the result code in the AAA-I answer. If both the values match, BSF Management service initiates stale record notification. To add the result code, perform the following steps: - Click Add. The Add Answer with Result Code Configuration dialog box opens. - For the Answer with Result Code, select any of the following valid values from the drop-down list: - DIAMETER_UNABLE_TO_COMPLY - DIAMETER_UNABLE_TO_DELIVER - EXPERIMENTAL_RESULT_CODE Note: If you select EXPERIMENTAL_RESULT_CODE, enter the required values for Result Code and Vendor ID. - Click Save on the dialog box.
Query to PCF	Indicates whether BSF management service queries PCF to confirm the status of a PcfBinding record, which is suspected as stale by the audit service. Default value: false Note: When Query to PCF parameter is set to false, the value of "Minimum Audit Attempts" parameter in Service Configurations of Management Service, and "Forced Deletion - Minimum Audit Attempts" parameter in Service configurations of Audit Service should be set to 0.
Vendor ID	Specifies the vendor ID that BSF retrieves from the Vendor Specific Attribute to send query requests towards PCF. The vendor ID should be 6-digit long. Note: PCF sends the Vendor Specific Attribute in the request body at the time of binding registration.
Minimum required 404 responses	This configuration is used by the BSF Management service to determine when a PCF binding should be considered stale. A binding is marked stale after the configured number of HTTP 404 responses is received from PCF instances. Minimum: 1 Maximum: Must be less than or equal to the Alternate Routing Retry (ARR) retry count configured in the Retry Profile.

Table 4-21 (Cont.) Audit

Parameter	Description
Expect 404 from different PCF	When this flag is disabled, all 404 responses are counted toward the required minimum, whether they come from the same PCF instance or different ones. When this flag is enabled, only 404 responses from different PCF instances are counted. Uniqueness is determined using the Server header in the 404 response to identify distinct PCF instances. Default: False
PCF Communication Profile	The NF communication profile to be used for BSF to PCF custom audit notification API. It is used to configure Binding Header, Server Header and SBI Correlation.

Table 4-22 Active Bindings Counting

Parameter	Description
Count Active Binding	Enables or disables the active sessions counting. By default, the active sessions counting is disabled. To enable the feature, set the value of this parameter to <code>true</code> .
Bindings Count Interval (in minutes)	Specifies the time interval (in minutes) for which maximum active sessions are reported as a metric. Default value is 15 minutes. You can set the time interval to any value between 1 to 60 minutes.
Root Log Level	Specifies the log level of BSF Management service. The available values for this field are as follows: - Trace - Debug - Information - Warn (Default) - Error - Always
Log Levels	To add the log levels, perform the following steps: - You can add log levels using the Log Levels group on this page. - To add log level: - Click  Add . This opens the Add Log Levels dialog box. - On the dialog box, enter the values shown in Table 4-23. - Click Save on the dialog box. The log level information for the selected service type is saved.

Table 4-23 Parameters for Add Log Levels Configurations

Parameter	Description
Logger Name	Specifies the name of the logger.

Table 4-23 (Cont.) Parameters for Add Log Levels Configurations

Parameter	Description
Level	Specifies the log level of BSF Management service. Select any of the following valid values: • Trace • Debug • Information • Warn (Default) • Error • Always

Table 4-24 Message Priority Profile

Field	Description
Message Priority Profile	Specifies the message priority profile to be applied to message priorities. The support for SBI Message Priority header gets enabled when a message priority profile is selected for Message Priority Profile field, and 3gpp-Sbi-Message-Priority header is added to the outgoing and incoming HTTP requests and responses to BSF depending on the configuration of the Message Priority Profile selected for BSF Management Service.

Table 4-25 Alternate Lookup Table Settings

Field	Description
Enable IPv4 Address Table	Enables/disables the ipv4_addr_lookup table. Default Value: true
Enable IPv6 Prefix Table	Enables /disables the ipv6_prefix_lookup table. Default Value: true
Enable MAC Address 48 Table	Enables /disables the mac_addr_48_lookup table. Default Value: true

Table 4-26 Alternate Lookup Table Audit Settings

Field	Description
Enabled	Enables/disables audit for all the alternative lookup tables (IPv4, IPv6, MAC). Set this to true when all lookup tables are enabled in Alternate Lookup Table Settings. Default Value: false
Minimum Time to Live (in minutes)	Minimum Time to Live (minTTL) in minutes for any record in all lookup tables. Default Value: 7200

Table 4-26 (Cont.) Alternate Lookup Table Audit Settings

Field	Description
Maximum Time to Live (in minutes)	Maximum Time to Live (maxTTL) in minutes for any record in the lookup tables. Default Value: 9800
Notification Rate (per second)	Notification rate per second at which Audit will notify BSF to check if a record in the lookup tables is stale. Default Value: 50
Minimum Audit Passes Interval (in minutes)	Frequency in minutes at which Audit will send notifications whenever no successful response (2xx or 404) has been received for a lookup record. Default Value: 10

4. Click **Save** on the **Edit Management Service** page to save your changes.
5. Perform the following steps to configure **Advanced Settings**:
 - a. Click the  **Add**.
The page opens the **Add Advanced Settings** dialog box.
 - b. In the dialog box, enter the following **keys** and respective **values**:

The following table describes the keys and values:

Table 4-27 Add Advanced Settings Configurations

Key	Value
BSF_MGMT_GET_PCF_BINDING_REQUEST_PRIORITY	Get Pcf Bindings: /nbsf-management/v1/pcfBindings Allowed Values: 0-31 Default Value: 15
BSF_MGMT_POST_PCF_BINDING_REQUEST_PRIORITY	Post Pcf Bindings: /nbsf-management/v1/pcfBindings Allowed Values: 0-31 Default Value: 20
BSF_MGMT_DELETE_PCF_BINDING_REQUEST_PRIORITY	Delete Pcf Bindings: /nbsf-management/v1/pcfBindings/{bindingId} Allowed Values: 0-31 Default Value: 16
BSF_MGMT_POST_AUDIT_NOTIFY_REQUEST_PRIORITY	Post Audit Notification: /audit/notify Allowed Values: 0-31 Default Value: 25
BSF_MGMT_DELETE_AUDIT_NOTIFY_REQUEST_PRIORITY	Delete Audit Notification: /audit/notify Allowed Values: 0-31 Default Value: 25
BSF_MGMT_GET_SESSION_VIEWER_PCF_BINDING_REQUEST_PRIORITY	Get Session Viewer Pcf Bindings: /oc-bsf-query/v1/pcfBindings Allowed Values: 0-31 Default Value: 28

Table 4-27 (Cont.) Add Advanced Settings Configurations

Key	Value
BSF_MGMT_POST_SESSION_VIEWER_DELETE_PCF_BINDING_REQUEST_PRIORITY	Post Session Viewer Pcf Bindings Cleanup: /oc-bsf-query/v1/pcfBindings/cleanup/{bindingId} Allowed Values: 0-31 Default Value: 18
BSF_MGMT_DELETE_SESSION_VIEWER_DELETEALL_PCF_BINDING_REQUEST_PRIORITY	Delete Session Viewer Pcf Bindings Cleanup All: /oc-bsf-query/v1/pcfBindings/cleanupAllBsfPcfBindings Allowed Values: 0-31 Default Value: 18
BSF_METRICS_SEGREGATED_STATUS_CODES	Enable the different HTTP status codes to be segregated in the status codes for BSF Metrics. For example, 2XX, 4XX, and 5XX Error codes. Allowed values: 100-599 Default value: 404
ENABLE_PCF_COOKIE	Indicates whether to enable the support for optimizing N7 session lookup for AAR messages in Rx call flows. Possible values: • true (to enable the feature) • false (to disable the feature) If this key is enabled, BSF stores the Rx session identifier details received from Policy in ocnf-service-cookie header during binding registration request. Later when BSF receives a AAR-I request, it retrieves the Rx session identifier details from BSF Management Service database and sends the session identifier details as OCNF-SERVICE-COOKIE Custom AVP in the AAR-I request to Policy. Default value: false
ENABLE_PCF_BINDING_TABLE_MIGRATION	Enables/disables this feature. PCF_BINDING_TABLE_LOOKUP must be at least 1 to set this to true. An alert will be set and default settings will be used in case the conditions are not met. Default value: false

Table 4-27 (Cont.) Add Advanced Settings Configurations

Key	Value												
PCF_BINDING_TABLE_LOOKUP	Signifies the table lookup order.												
Table 4-28 PCF_BINDING_LOOKUP_TABLE													
<table><tr><th>Value</th><th>Name</th><th>Description</th></tr><tr><td>0</td><td>PCF_BINDING_ONLY</td><td>Only looks up pcf_binding table (backward compatibility).</td></tr><tr><td>1</td><td>PCF_BINDING_FIRST</td><td>Looks up both tables - pcf_binding table first, followed by pcf_binding_v2 table. It allows BSF to update or delete a binding record in the table where the record currently exists. It does not migrate the record between tables. This supports replication across GR sites. Customers must set this value to 1 on all GR sites before enabling ENABLE_PCF_BINDING_TABLE MIGRATION on any GR site. Ideally, binding table migration should be enabled on all GR sites while this configuration remains enabled.</td></tr><tr><td>2</td><td>PCF_BINDING_V2_FIRST</td><td>Looks up both tables - pcf_binding_v2 table first, followed by pcf_binding table. This value must be set to 2 when most of the records (>50%) are migrated. This setting does not affect functionality. However, it may</td></tr></table>		Value	Name	Description	0	PCF_BINDING_ONLY	Only looks up pcf_binding table (backward compatibility).	1	PCF_BINDING_FIRST	Looks up both tables - pcf_binding table first, followed by pcf_binding_v2 table. It allows BSF to update or delete a binding record in the table where the record currently exists. It does not migrate the record between tables. This supports replication across GR sites. Customers must set this value to 1 on all GR sites before enabling ENABLE_PCF_BINDING_TABLE MIGRATION on any GR site. Ideally, binding table migration should be enabled on all GR sites while this configuration remains enabled.	2	PCF_BINDING_V2_FIRST	Looks up both tables - pcf_binding_v2 table first, followed by pcf_binding table. This value must be set to 2 when most of the records (>50%) are migrated. This setting does not affect functionality. However, it may
Value	Name	Description											
0	PCF_BINDING_ONLY	Only looks up pcf_binding table (backward compatibility).											
1	PCF_BINDING_FIRST	Looks up both tables - pcf_binding table first, followed by pcf_binding_v2 table. It allows BSF to update or delete a binding record in the table where the record currently exists. It does not migrate the record between tables. This supports replication across GR sites. Customers must set this value to 1 on all GR sites before enabling ENABLE_PCF_BINDING_TABLE MIGRATION on any GR site. Ideally, binding table migration should be enabled on all GR sites while this configuration remains enabled.											
2	PCF_BINDING_V2_FIRST	Looks up both tables - pcf_binding_v2 table first, followed by pcf_binding table. This value must be set to 2 when most of the records (>50%) are migrated. This setting does not affect functionality. However, it may											

Table 4-27 (Cont.) Add Advanced Settings Configurations

Key	Value									
Table 4-28 PCF_BINDING_LOOKUP_TABLE <table><thead><tr><th>Value</th><th>Name</th><th>Description</th></tr></thead><tbody><tr><td></td><td></td><td>have a performance impact.</td></tr><tr><td>3</td><td>PCF_BINDING_V2_ONLY</td><td>Looks up binding records only in the pcf_binding_v2 table, considering all binding records are migrated. At this point, the original pcf_binding table is expected to contain no records and can be deprecated.</td></tr></tbody></table>		Value	Name	Description			have a performance impact.	3	PCF_BINDING_V2_ONLY	Looks up binding records only in the pcf_binding_v2 table, considering all binding records are migrated. At this point, the original pcf_binding table is expected to contain no records and can be deprecated.
Value	Name	Description								
		have a performance impact.								
3	PCF_BINDING_V2_ONLY	Looks up binding records only in the pcf_binding_v2 table, considering all binding records are migrated. At this point, the original pcf_binding table is expected to contain no records and can be deprecated.								
AUDIT.ALT_LOOKUP_TABLE.IPV4_ADDR_TABLE.ENABLED	Enables Audit for the ipv4_addr_lookup table. Default Value: false									
AUDIT.ALT_LOOKUP_TABLE.IPV4_ADDR_TABLE.MIN_TTL	Minimum Time to Live (minTTL) in minutes for any record in ipv4_addr_lookup table. Default Value: 7200									
AUDIT.ALT_LOOKUP_TABLE.IPV4_ADDR_TABLE.MAX_TTL	Maximum Time to Live (maxTTL) in minutes for any record in ipv4_addr_lookup table. Default Value: 9800									
AUDIT.ALT_LOOKUP_TABLE.IPV4_ADDR_TABLE.NOTIFICATION_RATE	Notification rate per second at which Audit will notify BSF to check if a record in the ipv4_addr_lookup table is stale. Default Value: 50									
AUDIT.ALT_LOOKUP_TABLE.IPV4_ADDR_TABLE.FREQUENCY	Default Value: 10									
AUDIT.ALT_LOOKUP_TABLE.IPV6_PREFIX_TABLE.ENABLED	Enables Audit for the ipv6_prefix_lookup table. Default Value: false									
AUDIT.ALT_LOOKUP_TABLE.IPV6_PREFIX_TABLE.MIN_TTL	Minimum Time to Live (minTTL) in minutes for any record in ipv6_prefix_lookup table. Default Value: 7200									
AUDIT.ALT_LOOKUP_TABLE.IPV6_PREFIX_TABLE.MAX_TTL	Maximum Time to Live (maxTTL) in minutes for any record in ipv6_prefix_lookup table. Default Value: 9800									
AUDIT.ALT_LOOKUP_TABLE.IPV6_PREFIX_TABLE.NOTIFICATION_RATE	Notification rate per second at which Audit will notify BSF to check if a record in the ipv6_prefix_lookup table is stale. Default Value: 50									

Table 4-27 Add Advanced Settings Configurations

Key	Value
AUDIT.ALT_LOOKUP_TABLE.IPV6_PREFIX_TABLE.FREQUENCY	Frequency in minutes at which Audit will send notifications whenever no successful response has been received (2xx or 404) for a lookup record in the ipv6_prefix_lookup table. Default Value: 10
AUDIT.ALT_LOOKUP_TABLE.MAC_ADDR_48_TABLE.ENABLED	Enables Audit for the mac_addr_48_lookup table. Default Value: false
AUDIT.ALT_LOOKUP_TABLE.MAC_ADDR_48_TABLE.MIN_TTL	Minimum Time to Live (minTTL) in minutes for any record in mac_addr_48_lookup table. Default Value: 7200
AUDIT.ALT_LOOKUP_TABLE.MAC_ADDR_48_TABLE.MAX_TTL	Maximum Time to Live (maxTTL) in minutes for any record in mac_addr_48_lookup table. Default Value: 9800
AUDIT.ALT_LOOKUP_TABLE.MAC_ADDR_48_TABLE.NOTIFICATION_RATE	Notification rate per second at which Audit will notify BSF to check if a record in the mac_addr_48_lookup table is stale. Default Value: 50
AUDIT.ALT_LOOKUP_TABLE.MAC_ADDR_48_TABLE.FREQUENCY	Frequency in minutes at which Audit will send notifications whenever no successful response has been received (2xx or 404) for a lookup record in the mac_addr_48_lookup table. Default Value: 10
ENABLE_DISCOVERYPARAM_SERVICE_NAMES	Enables inclusion of the service name in the NF On-Demand Discovery request. Default Value: false
SBI_CALLBACK_HEADER_VALUE	Specifies the callback header value for audit notification requests, overriding the default value defined by the feature.

Note

- Enable audit for each alternate lookup table only when the corresponding lookup table is enabled in the advanced settings. For example:
 - Enable `AUDIT.ALT_LOOKUP_TABLE.IPV4_ADDR_TABLE.ENABLED` only when Enable IPv4 Address Table is enabled.
 - Enable `AUDIT.ALT_LOOKUP_TABLE.IPV6_PREFIX_TABLE.ENABLED` only when Enable IPv6 Prefix Table is enabled.
 - Enable `AUDIT.ALT_LOOKUP_TABLE.MAC_ADDR_48_TABLE.ENABLED` only when Enable MAC Address 48 Table is enabled.
- In the NF On-Demand Discovery (ODD) request, the service name in the discovery parameters is set either from the binding header or by using the default value. However, since this is an optional parameter, it can be disabled by setting the advanced configuration key `ENABLE_DISCOVERYPARAM_SERVICE_NAMES=false`.
- By default, the callback header value is sent as `Nbsf_Management_myNotification` when it is enabled in the NF Communication Profile. However, this value can be customized using the advanced configuration key `SBI_CALLBACK_HEADER_VALUE` by setting it to the required string value.
- Fresh installations automatically set `ENABLE_PCF_BINDING_TABLE_MIGRATION=true`, `PCF_BINDING_TABLE_LOOKUP=3`, `Enable IPv4 Address Table=true`, `Enable IPv6 Prefix Table=true`, and `Enable MAC Address 48 Table=true`. These configurations should never be modified.

4.4.2 Audit Service

This section provides information about configuring the Audit Service.

4.4.2.1 Audit

The **Audit Service** page displays the Audit Service configurations. The page allows you to edit the configurations.

To configure Audit service:

1. From the navigation menu under **BSF**, navigate to **Service Configurations**, and select **Audit Service** and browse to **Audit** page.
This opens the **Audit** page. The page displays the existing configurations.
2. Click  **Edit**.
This opens the **Edit Audit** page.
3. Make sure that the value of the **Audit Enabled** switch is enabled.
This field determines if auditing is enabled for BSF Management Service. By default, this switch is enabled.
4. Expand the **Forced Deletion** group and configure the **Minimum Audit Attempts** parameter.

Minimum Audit Attempts specifies the minimum number of audit attempts until ForceTTL is reached.

If ForceTTL is reached and `audit_attempts + 1 >= Minimum Audit Attempts of ForceTTL`, then Audit service deletes the identified stale records from its respective database.

The default value of this parameter is 0 and the value can range between 0 to 255.

5. Click **Save**.
The page saves the Audit service configurations.

① Note

Important considerations during Audit configurations:

In situations where the number of stale records (records that are eligible to be audited) at any given point of time in the system (for a given microservice database) is expected to be high (for example, > 1M), then the **Session Age** and/or the **Notification Rate** parameters should be set appropriately such that at least 2 audit cycles can be finished before the records that were assessed by the first cycle fall stale again. The Audit procedure having a **Session Age** less than this recommendation, may not be able to assess all stale records as the already assessed ones will be stale again too soon. It is recommended to keep a minimum of 24 hours for the **Session Age** and a minimum of 48 hours for **Max Binding Age**.

The time taken to complete an Audit Cycle and begin the next one can be calculated as below:

Audit Cycle Time = $S / (N * 60) + I$ minutes, where,

S = expected number of stale sessions at any given time,

N = notification rate (per second),

I = minimum Audit Interval

4.4.2.2 Audit Schedule Data

Audit service supports multiple pod, using the Audit Schedule and Audit TaskScheduler. Audit schedule has list of registered services that needs audit service with all the scheduled details. Audit TaskScheduler polls for all those audit jobs in the QUEUED status.

Figure 4-5 Audit Scheduled Service List Data

Service Name	Table Name	Next Schedule	Status	Detail View
bsf-management-service	pcf_binding	2023-10-05T12:29:41.000+00:00	QUEUED	


Table 4-29 Audit Schedule Data Fields

Field Name	Description
Service Name	Name of the service that had registered for audit service.

Table 4-29 (Cont.) Audit Schedule Data Fields

Field Name	Description
Table Name	Name of the database table that has requested for the audit service
Next Schedule	Next schedule availability time
Status	Scheduler job status
Start Time	Audit start time of pod.
End Time	Audit end time of pod.
Last Polled Time	Pod and associated scheduler details.
Scheduler Details	Last time at which pod updated this table.
Is Service Dependent	If this is true, notification will be triggered for one table in the service, for all table's in the service, notifications will be sent parallelly based on the notificationRate set for each table.

Audit Schedule task can be in one of the state described below:

Figure 4-6 Audit Schedule Job Work Flow**Table 4-30 Audit Schedule Job Status**

Status	Description
QUEUED	When a registered request is received it is added to the schedule table, the job status is set to Queued.
INITIALISING	When polling task is complete and is initializing for AuditTaskManager to audit, the job status is set to Initializing.
RUNNING	When Audit task manager starts the task, the job status is set to Running.
COMPLETE	When Audit TaskManager completes the Audit, the status is set to Complete.

Table 4-30 (Cont.) Audit Schedule Job Status

Status	Description
INTERRUPTED	When the audit process is paused from GUI, then the job status is set to Interrupted. When resumed from GUI, job status is set to Queued When deregistered, the job data will be removed from the table. when a service is register, it is set to Queued.
DEREGISTERED	When deregistered, the job status is set to Deregistered for Initializing or Running and the entry will be deleted for any other status(except Initializing or Running). In case the a service register request is received before AuditTaskManager deletes the table entry for deregistered job statuses, the status would be moved to Queued.

For more information on Audit Service and Audit Schedule REST API details, see the section *Audit Service in Oracle Communications Cloud Native Core, Binding Support Function REST Specification Guide*.

4.4.3 IGW

1. From the navigation menu under **BSF**, navigate to **Service Configurations**, and then select **IGW**. This opens the **Error Code Series List** page.
2. Click Edit on the Error Code Series List.
This opens **Edit Error Code Series List** page.
3. Click **Add** to add the error code series.
The **Add Error Code Series** page is displayed.
4. Configure the fields in the **Add Error Code Series** page as mentioned in .
5. Click **Add** in the **Error Code Series** section.
The **Add Error Code Series** page is displayed.
6. Configure the fields in the **Add Error Code Series** page as mentioned in .
7. Click **Save** on the **Add Error Code Series** page to save the details. Click **Cancel** to discard your progress and go back to **Create Error Code Series** page.
8. Click **Save** on the **Create Error Code Series**page to save the details. Click **Cancel** to discard your progress and go back to **Error Code Series** page.

Note

Use the **Edit** icon or **Delete** icon available in the next column of the specific entry to update or delete the error code series information.

Table 4-31 Error Code Series

Field Name	Description
Profile	Specifies the name for error code series list profile.
Error Code Series	Lists the error codes for a specific service.

Table 4-32 Error Code Series Configuration

Field Name	Description
Error Set	Possible values for "errorSet" attribute: 5xx, 4xx, 3xx
Error Codes	Possible values include all error codes in the respective HttpSeries value assigned for Error Set . Note: Use "-1" if all error codes in that HttpSeries are to be considered.

4.4.4 Common

This section describes how to perform common configurations for BSF Management Service.

4.4.4.1 Message Profiles

BSF allows to configure creation, update, and addition of message profiles for Management service.

Perform the following procedures to create and configure message profiles:

1. Navigate to **Message Profiles** page under **Common Data** for **Service Configurations** in CNC Console for BSF.
2. Type a key word in **Type to Filter** text box to filter the required message profile name from the available list of message profiles.
3. To refresh the list of message profiles, click the **Refresh** button ( **Refresh**).
4. To delete an existing message profile, click **Delete** button ().
5. To create a message profile, and perform the following configurations.:
 - a. Click **Add** button ( **Add**).
 - b. In the **Create Message Profiles** page perform the following configurations:
 - i. Type the name of the message profile in **Message Profile Name** text box.
 - ii. From the **Service Type** list box, select the name of the service for which message profile is applied.

Note

Currently, BSF supports only BSF Management Service.

- iii. Configure the rules defined for each message profile as follows:

- i. Click **Add** button ( **Add**) in the **Message Profile Attributes** table.
- ii. In the **Add Message Profile Attributes** window configure the message profile attributes as follows:

Table 4-33 Add Message Profile Attributes

Field	Description
Message ID	Specifies the unique identifier that defines a message profile rule.
Rule Priority	Specifies the priority of the rule. Rule priority is an integer value between 1 to 65535. Note: Each message priority profile includes a unique rule priority for all the rules created in it and not for each message type. Unique rule priority is for per message profile , not for per message type.
Interface Type	Specifies the interface type such as Nbsf_Management or Npcf_Binding_Audit. for which the rule is defined. Default value: Nbsf_Management
Message Type	Specifies the message type for an interface for which the rule is defined. Possible values are: • For Nbsf_Management: – Register – Deregister – Discovery – Any Note: When message type is Any, it is applicable to all the message types for that interface. • For Npcf_Binding_Audit: – Audit_Notify – Any Default value: For Nbsf_Management interface type, default Message Type should be Register. For Npcf_Binding_Audit interface type, default Message Type should be Audit_Notify.

Table 4-33 (Cont.) Add Message Profile Attributes

Field	Description
Message Profile Filter Conditions	Configure the message profile filter conditions to filter out rules that matches these attributes as explained in Table 4-34 and click Save. Note: Message Profile Filter Conditions is not a mandatory attribute for rules. Note: When Request Priority Override is selected as IfNotPresent, the Message Profile Filter Conditions is not valid. If Message Profile Filter Conditions is not configured the rule is applicable for all priority values.
Request Priority Override	Specifies the rule to override the request priority value mentioned in the header, consider the request priority as configured in Request Priority Value field and assign a response priority as configured in Response Priority Assign field. Possible values are: • Never: Indicates never to consider this value and whatever is present in request header will be propagated. • Always: Indicates to always consider this value irrespective of header request value. • IfNotPresent: Indicates to consider this value if respective header request value is not present. Default value: Never
Request Priority Value	This field appears only when the value of Request Priority Override field is Always or IfNotPresent. It allows to mention the custom request priority value. Range: 0-31
Response Priority Assign	It specifies the rule to specify response priority value. Possible values are: • Never: Indicates not to assign any response priority. • Custom: Allows to use specify the desired response priority in the Response Priority Assign field. • EvaluatedRequestPriority: Indicates that whatever is the final request priority, the same gets assigned to the response. Default value: Never
Response Priority Value	This field appears only when the value of Response Priority Assign field is Custom. It allows to mention the custom response priority value. Range: 0-31

Table 4-34 Add Message Profile Filter Conditions

Field	Description
Attribute Name	Name of the attribute to be filtered out such as 3gpp-sbi-message-priority or oc-message-priority. Default value: 3gpp-sbi-message-priority
Attribute Value	This is an optional field to mention the value of the attribute. That is, it is not mandatory for each rule, and it can be left blank. This field accepts either a single value or a range of values between 0 to 31. It does not accept any value outside the range 0-31.


- c. Click **Save** to save the message profile attribute configuration.

6. To edit an existing message profile, click the **Edit** button () corresponding to the message profile name.

In the **Edit Message Profiles** window, perform the following:

- a. Edit the **Service Type**.
- b. Edit the message profile attributes:

- i. Click the **View** button () corresponding to the required **Message ID** to view the details of the message profile attributes.

- ii. Click the **Add** button () to add a new set of message profile attributes.

- iii. Click the **Delete** button () corresponding to the required **Message ID** to delete an existing set of message profile attributes.

- iv. Click the **Edit** button () corresponding to the required **Message ID** to edit an existing set of message profile attributes.

- v. Click **Save** to save the changes.

4.4.4.2 Retry Profiles

This procedure provides information about how to create and manage retry profiles.


The Retry Profile page allows you to create new and manage existing retry profiles. The page displays the list of defined retry profiles with the option to add profiles.

To configure retry profile:

1. From the navigation menu under **BSF**, navigate to **Service Configurations**, click **Common Data**, and then select **Retry Profiles**.

This opens the **Retry Profiles** page. The page lists the existing retry profiles. You can add new retry profiles using this page.

2. Click



. This opens the **Create Retry Profile** page.

3. On the **Create Retry Profile** page, enter the following information:

Table 4-35 Create Retry Profile

Field Name	Description
Name	The unique name for the retry profile. This name is used to refer to the retry profile in other service configuration screen, such as BSF Management Service. The name can only contain the characters A–Z, a–z, 0–9, period (.), hyphen (-), and underline (_). The maximum length is 255 characters.
Retry on Internal Send Failure	Specifies whether to enable or disable retry for failed messages between core microservices of BSF. On enabling this switch, ensure that you configure the parameters under Retry Settings group. However, if the switch is disabled, the configurations in the Retry Settings group are not taken into consideration. When enabled, retries are attempted when BSF encounters internal send failures for egress messages. A failure to send egress messages internally indicates an exception in the core service, such as BSF Management service. The exception can occur while connecting to the next microservice or a failure on the microservice itself. For example, egress message did not reach the egress gateway. This also includes the connection errors and timeouts on the core microservice. Note: Since the failure is internal, the external NF and SCP (if applicable) remain the same during retry attempt.
Enable Alternate Routing	Specifies whether to enable or disable alternate routing. On enabling this switch, ensure that you configure the parameters under the Alternate Routing Settings group. When enabled, alternate routing is attempted when the BSF messages are not delivered successfully to an external NF destination. It may happen because of an exception on the Egress Gateway, such as connection error, timeout, or a failure returned by the external NF or an intermediate router.
Server Header Support	Specifies whether to enable or disable server header. The value configured in this parameter is used when BSF acts as a consumer and receives server header in error response messages. Depending on the requirements, you can select any of the following values from the drop-down list: - Disabled (default) - Select this option to disable server header. On disabling server header, BSF as a consumer searches for alternate producer without considering the value of server header. - Single Instance - Select this option to accept only one value in the server header. When the value is configured as Single Instance and server header contains multiple values, BSF rejects the header.

Table 4-35 (Cont.) Create Retry Profile

Field Name	Description
Pattern To Ignore Server Header Value	Specifies the pattern to ignore server header values. This field appears only when single Instance or Multi Instance is selected from the drop-down of the Server Header Support field.
Retry Settings	Provide Retry Settings details:
Maximum Number Of Retries	Specifies the number of retries that BSF attempts on encountering failure while sending egress messages from BSF management service. Note: The value configured for this parameter does not include the initial attempt. For instance, if you set the value as 3, BSF performs the retry cycle three times after the initial attempt.
Alternate Routing Settings	Provide Alternate Routing Settings details:
Maximum Number Of Alternate Routing Attempts	Specifies the number of retries that BSF attempts on receiving error response from Egress Gateway. The number of retry attempts can be in the range of 1 to 10. Note: The value configured for this parameter does not include the initial attempt. For instance, if you set the value as 3, BSF performs the retry cycle three times after the initial attempt. If the retry attempts are exhausted, CNC BSF fails the transaction with an exception. In certain cases, the number of retry attempts can be greater than the available alternate destinations to be tried. If such a situation arises, BSF fails the transaction when no alternate destination is available even if retry attempts are still available.
Attempt Alternate Route for Following Error Codes	Allows to configure the HTTP error codes for which BSF reattempts the requests after the ARR request failures. To add a value, click Add and enter values for the following fields: - Error Code - Specifies the HTTP error code. For example, 504 - Error Causes - Specifies the error cause. For example, GATEWAY_TIMEOUT, NF_TIMEOUT. You can configure up to 32 error codes with each error code supporting up to five causes. Note: The Error Causes and Error Code are unique. In case you import them from earlier version containing duplicate values, it will pick only one value.
Use Alternate SCP for Alternate Routing	Specifies whether to enable or disable choosing alternate SCP to reach the alternate destination. If you set the value for this parameter as true, Egress Gateway uses an alternate SCP for routing to an alternate NF destination. If there are fewer SCPs than alternate NF destinations, Egress Gateway tries sending the request through one of the previously used SCPs. Note: Additional SCP configuration is required for egress gateway. For more information, see Oracle Communications Cloud Native Core BSF Installation and Upgrade Guide.
Use Binding Information from Binding Header	Specifies whether to enable or disable the use of information received in 3gpp-sbi-binding header. When BSF receives the unique NF Instance ID in the error response, it retrieves the 3gpp-sbi-binding header for that particular NF instance. It extracts the SetId to choose the alternate producer. This parameter is considered only for subsequent retry requests. By default, this switch is disabled.

Table 4-35 (Cont.) Create Retry Profile

Field Name	Description
NF Set Resolution	Specifies whether to enable the on demand discovery of subsequent messages based on the NFSet ID, in case of failure response. The possible values are: NRF Discovery: If this value is selected, then BSF sends a request to NRF to get the list of NFs based on the SetId value from the 3gpp-Sbi-Binding header of the NF with a successful Initial Request. By default NRF Discovery remains selected. Note: NRF discovery is applicable for subsequent messages functionality only.
Locality Preference	Specifies the locality preference when choosing the alternate producer for a failed request. BSF excludes the producer (received in the failed request) and selects an alternate producer based on the value you choose from the drop-down menu: - Same: Choose the next producer from the same locality as the previous request. - Different: Choose the next producer from a different locality. - None (default): Locality filter does not apply.
Retry again at end of alternate NF options	When you set the value for this parameter to true, BSF initiates the retry cycle and sends request to the selected NF when the following conditions are true: - Retry attempts made to all available NF instances - Maximum Number of Alternate Routing Attempt is not exhausted

4. Click Save.

The retry profile gets listed on the **Retry Profile** page. The page defines the retry profile in the BSF database and it is available to be used in BSF.

Note

Use



or



available under the Actions column to update or delete the profiles.


Note

- The option to configure a Retry Profile for the Initial Message is not available.
- In the PCF configuration, the option to configure a Retry Profile is available on the Service Configuration page in addition to the NF Communication Profile. However, in the BSF configuration, the Retry Profile can only be configured through the NF Communication Profile.
- The Enable DNS Resolution and its related subsequent configurations will not be present in the GUI.

4.4.4.3 NF Communication Profiles

The NF Communication Profiles page allows you to create new and manage existing NF communication profiles. The page displays the list of defined profiles with the option to add profiles. To configure NF communications profile, perform the following steps:

1. From the navigation menu under **BSF**, navigate to **Service Configurations**, click **Common Data**, and then select **NF Communication Profiles**. This opens the **NF Communication Profiles** page. The page lists the existing profiles. You can add new profiles using this page.
2. Click



. This opens the **Create NF Communication Profile** page.

3. On the **Create NF Communication Profile** page, enter the following information:

Table 4-36 Create NF Communication Profile

Field Name	Description
NF Communication Profile Name	The unique name for the NF communication profile. This name is used to refer to the profile in other service configuration screen, such as BSF Management Service. The name can only contain the characters A–Z, a–z, 0–9, period (.), hyphen (-), and underline (_). The maximum length is 255 characters.
BSF NF Communication Model	Indicates the communication model as defined in 3GPP Technical Specification 23.501 Annex E, in which the BSF NF is deployed. Select any of the following values from the drop down menu: • Custom: Select this option to create a customized communication profile. • Default Value: Custom
NF Discovery Settings	Provide NF Discovery Settings details:

Table 4-36 (Cont.) Create NF Communication Profile

Field Name	Description
Send Discovery Header in Subsequent Message	This switch indicates if BSF must include the following 3gpp-sbi-discovery-<discovery-parameter> parameters in SBI messages: - Parameters to update a subscription, such as PUT messages when unable to fill in a 3gpp-Sbi-Routing-Binding header. - Parameters to delete a subscription, such as DELETE messages when unable to fill in a 3gpp-Sbi-Routing-Binding header. By default the switch remains disabled. (Not Supported)
Send Target API Root Header in Initial Messages	Indicates if BSF must include 3gpp-sbi-target-apiroot in initial request messages.
On Demand Discovery Caching	Provide On Demand Discovery Caching details:
Force Discovery	Indicates whether to cache the NF profiles in NRF client or to skip caching and receive the response directly from NRF. Value of this parameter can be 0 or 1. - When the value of OC-Force-Rediscovery parameter is set to 0, global cache configuration is taken into account. - When the value of OC-Force-Rediscovery parameter is set to 1, global cache configuration is ignored and NF profiles get retrieved directly from NRF while responding to the query from the backend services. Default value: 0
Retention Period (in ms)	Indicates the time a record is allowed to stay in database after expiry. This parameter accepts an integer value to indicate the retention period in milliseconds. When retention period is not configured or saved as null, backend will not send this header to NRF connector. Default value: 0
NF Bindings Settings	Provide NF Bindings Settings details:
Binding Level	Indicates the binding level that must to be included in the 3gpp-sbi-binding header when BSF adds this header in a message to another NF. Select any of the following values from the drop down menu: - NF Set - NF Instance Default Value: NF Set
Send Binding Header	Indicates if BSF includes the 3gpp-sbi-binding header in SBI messages for the notification requests and responses, as applicable. By default the switch remains disabled.

Table 4-36 (Cont.) Create NF Communication Profile

Field Name	Description
Send BSF Service Name in Binding Header	Indicates whether to include the service name in the binding header or not. Possible values: - true: Includes svcname in the binding header sent towards NFs such as PCF, which contains the 3GPP defined BSF service name. - false: Does not include svcname in the binding header. Default value : false (Not Supported)
Send Routing Binding Header	Indicates if BSF includes the 3gpp-sbi-routing-binding header in SBI messages for the notification requests, as applicable. By default the switch remains disabled.
Send Callback Header	Indicates if BSF includes the 3gpp-sbi-callback header in SBI messages for the notification requests, as applicable. By default the switch remains disabled. If enabled, then operator has advanced settings SBI_CALLBACK_HEADER_VALUE in BSF Management Service to modify the callback header default value Nbsf_Management_myNotification being sent in audit notification towards PCF.
NF Server Settings	Provide NF Server Settings details:
Send Server Header	Indicates if services BSF microservices include server header while sending an error response.
Server Header Error Codes	Specifies the error codes for which service header is generated. The error codes can be from 100 to 999. Note: If no error is specified, then BSF sends server header for all error responses.
Retry & Alternate Routing Settings	Provide Retry & Alternate Routing Settings details:
Retry Profile for Subsequent Messages	Specifies the retry or alternate routing options when a message is not delivered successfully.
NF Correlation Settings	Provide NF Correlation Settings details:
Send Correlation-info header	This specifies if the BSF should send Subscriber identifier as sbi-header-info or not. By default the switch remains disabled.
Allowed Correlation-info Header Generation Type(s)	Specifies that if correlation header is not received from Consumer NFs, BSF should generate the header. The Correlation-Type supported - SUPI - GPSI (Select either both or none)

4. Click **Save**.

The NF communication profile gets listed on the NF Communication Profiles page.

Note

- The NF Communication Profile must be deleted in the Service Configurations before deleting from this page as cyclic references are not handled currently.
- The NF Communication Profile screen will provide only a single communication model option, Custom. Users must configure Model-C communication using the Custom communication type.
- The options for Initial Discovery Parameters and Send Discovery Header in Initial Message are not available.
- The Send Target API Root Header in Initial Message flag may be enabled or disabled; however, regardless of this setting, BSF Management sends the 3gpp-Sbi-Target-apiRoot header in all audit notifications to PCF.

Note

Use



or




available under the Actions column to update or delete the profile.

4.4.5 NRF Agent Service

This page describes how to configure the NRF Agent service.

Perform the following procedure to configure the NRF Agent service:

1. Navigate to **NRF Agent Service** page under **Service Configurations** in CNC Console.
2. To refresh the existing NRF Agent service configuration, click the **Refresh** button



3. To edit the existing NRF Agent service configuration, click the **Edit** button ().
4. In the **Edit NRF Agent Service** page, perform the following configurations:

Table 4-37 Edit NRF Agent

Field	Description
NRF Agent Communication Profile	Specifies the NRF Communication Profile to handle correlation info header toward NRF. (Not Supported)
On Demand Discovery Caching	Provide details about On Demand Discovery Caching
Enable Caching	Indicates NRF Client to cache the profile data and use it to process the future discovery on demand requests. Default value: false
Query Parameters	Accepts the list of parameters that NRF Client uses to fetch the profile data from cache. Note: The query parameters in the request from BSF Management service must match and must be in the same order as the data received from NRF. Otherwise, NRF Client can not process the request. Supported values: • data-set • target-nf-type • requester-nf-type • target-nf-set-id • service-names • amf-region-id • GUAMI • target-nf-instance-id • group-id-list
Retention Period (in ms)	Indicates the time a record is allowed to stay in database after expiry and before any clean process is executed. Default value: 864000000
Refresh Before Expiry Time (in ms)	Indicates the time in milliseconds before expiration when a rediscovery must be performed for the requested discovery parameters. Default value: 4500
Cached entry limit	Parameter to set the upper bound of cache registries allowed to be retained. Default value: 2000
Auditor Scheduler Interval (in ms)	Parameter to set time in milliseconds to activate the Discovery Cache auditor service to clean stale records. Default value: 1800000
Auditor Purge Percentage	Parameter to set the percentage of the cached discovery requests that will be pruned once the threshold has been passed. Default value: 20
Traffic Prioritization	Provide details about Traffic Prioritization
Enable Feature	Specifies whether to enable or disable the traffic prioritization feature at NRF Agent service. Default value: Disabled

Table 4-37 (Cont.) Edit NRF Agent

Field	Description
Incoming Priority Header	Specifies the name of the incoming header. Default value: 3gpp-sbi-message-priority
Outgoing Priority Header	Specifies the name of the outgoing header. Default value: 3gpp-sbi-message-priority
Message Types	Click the Edit button () corresponding to the required Message Type to edit an existing message type as explained in Table - Message Type. Click the Delete button () corresponding to the required Message Type to delete an existing message type. Click the Add button ( Add) to create a message type as explained in Table - Message Type.

Table 4-38 Message Types

Field	Description
Message Type	This field lists all the BSF interactions with NRF message types with priority. Possible values: - Register(AutonomousOnDemandNFRegistration) - Deregister(AutonomousNfDeRegistration) - Update(AutonomousNfPatch) - Heartbeat(NfHeartBeat) - Health_Check(AutonomousHealthCheck)
Priority	Specifies the priority for the selected message type. Default value: Default value of the priority is set to 1 for the following message types: - Register (AutonomousOnDemandNFRegistration) - Heartbeat (NfHeartBeat) - Update (AutonomousNfPatch) - Deregister (AutonomousNfDeRegistration) - Health_Check (AutonomousHealthCheck)

4.5 Diameter Configurations

This section describes how to manage and view the Diameter Configurations in BSF using the Diameter Configurations pages.

The Diameter configuration includes:

- [Settings](#)
- [Peer Nodes](#)
- [Diameter Routing Table](#)

4.5.1 Settings

The **Settings** page displays the general configurations related to the Diameter Gateway. The page allows you to edit the configurations.

To edit settings:

1. From the navigation menu, under **BSF**, click **Diameter Configurations** and select **Settings**.
This opens the **Settings** page. The page displays the existing configurations.
2. Click  **Edit**.
This opens the **Edit Settings** page.
3. Enter the following information under the respective groups:

Table 4-39 Timer

Field Name	Description
Reconnect Delay (sec)	Enter the time frame to delay before attempting to reconnect after a connection failure in seconds. The default value is 3 seconds.
Response Timeout (sec)	The amount of time Diameter Gateway waits for the answer to come from the sent request. It is a global value applicable for all the interfaces messages. Enter the response timeout interval in seconds. Note: To enable Application or Command code Response timeout value, see Enhanced Timer Configuration. The default value is 5 seconds.
Connection Timeout (sec)	Enter the connection timeout interval in seconds. The default value is 3 seconds.
WatchDog Interval (sec)	Enter the watchdog interval in seconds. The default value is 6 seconds.

Table 4-40 Transport

Field Name	Description
Protocol	The protocol supported is TCP.

Table 4-41 Congestion Control

Field Name	Description
Load Shedding Profile	Select any one of the configured load shedding profiles for congestion control on Diameter interface from the drop-down list.
Message Priority Profile	Select any one of the configured message priority profiles for congestion control on Diameter interface from the drop-down list.

Table 4-42 Overload Control

Field Name	Description
Load Shedding Profile	Select any one of the configured load shedding profiles for overload control on Diameter interface from the drop-down list.
Message Priority Profile	Select any one of the configured message priority profiles for overload control on Diameter interface from the drop-down list. Note: The following message priority data that was exported prior to BSF 23.2.0 cannot be imported as the data may be corrupt: • message containing Sd as interface • Sy-SLR as condition message The data with Sd interface or Sy-SLR condition messages that are exported only with BSF 23.2.0 or later versions can be imported.

Table 4-43 Enhanced Timer Configuration

Field Name	Description
Application Name	Request Timer configuration for applications name like Rx, Gx, Sy, Sd.
Application Response Timeout (milliseconds)	Enter the application response timeout in milliseconds. The range of this value is between 3 seconds to 2147483647.

Table 4-44 Command Code Response TimeOut

Field Name	Description
AAR (milliseconds)	The command code response timeout value for AAR. The allowed value ranges from 3 to 2147483647. Default Value: 5000
STR (milliseconds)	The command code response timeout value for STR. The allowed value ranges from 3 to 2147483647. Default Value: 5000

Table 4-44 (Cont.) Command Code Response TimeOut

Field Name	Description
RAR (milliseconds)	The command code response timeout value for RAR. The allowed value ranges from 3 to 2147483647. Default Value: 5000
ASR (milliseconds)	The command code response timeout value for ASR. The allowed value ranges from 3 to 2147483647. Default Value: 5000

The order of precedence (from highest to lowest) of response timeout configurations is:

- a. Command Code Response Timeout (ms) - Message level configurations i.e at AAR, STR etc.
 - b. Application Response Timeout (ms) - Interface level configuration i.e at Gx, Rx etc.
 - c. Response Timeout (ms) - General level configuration
4. Perform the following steps to configure **Advanced Settings**:
- a. Under the Settings page of Diameter Configurations, click the  **Add**. The page opens the **Add Advanced Settings** dialog box.

- b. In the dialog box, enter the following **keys** and respective **values**:

The following table describes the keys and values:

Table 4-45 Add Advanced Settings Configurations

Key	Value
DIAMETER.Enable.Validate.Realm	Used to validate the destination-realm received in the AAR-I message against the destination-realm in BSF's Diameter Gateway host realm. Default value: false
DIAMETER.BSF.Enable.Validate.Binding.Realm	Used to validate the destination-realm of AAR-I against the discovered pcfBinding's realm. Default value: false
DIAMETER.BSF.Enable.Overwrite.Realm	Used to configure whether to overwrite the pcfBinding's realm and identity information in AAR-I destination-realm AVP. Default value: false

Table 4-45 (Cont.) Add Advanced Settings Configurations

Key	Value
CONGESTION_DIAMETER_RESPONSE_CODE	This is to configure the diameter result code for the congestion control discards and data type of value will be numerical. If value is not configured, then default congestion result code will be <i>3004 (DIAMETER_TOO_BUSY)</i> Default Value: 3004 Examples: CONGESTION_DIAMETER_RESPONSE_CODE=5065 CONGESTION_DIAMETER_RESPONSE_CODE=3004
CONGESTION_DIAMETER_VENDOR_ID	This is to configure the diameter vendor id for the congestion control discards and data type of value will be numerical. If value is non-zero, then the <i>CONGESTION_DIAMETER_RESPONSE_CODE</i> considered as the experimental result code. If value is zero / not configured, then <i>CONGESTION_DIAMETER_RESPONSE_CODE</i> considered as the normal result code. Examples: CONGESTION_DIAMETER_VENDOR_ID=10415 CONGESTION_DIAMETER_VENDOR_ID=21274
DGW_CONNECTION_REPAVING.MAX_STAGGER_JITTER_MS	Specifies the maximum duration in seconds for stagger window. All pod disconnections are randomly distributed within this time frame. This approach does not depend on the number of pods. Default value: 300 Range: 300–14400
DGW_CONNECTION_REPAVING.REINVOCATION_DELAY_MIN	Specifies the minimum time between repaving operations. This is the wait between invocations, to prevent the force disconnection from being triggered multiple times in a short period. Default value: 10 Range: 5–480
DGW_CONNECTION_REPAVING.CONNECTION_DISCONNECT_INTERVAL_MS	The delay in milliseconds that the pod pauses between terminating each target connection. This throttles the process to prevent flooding a peer. Default value: 5000 Range: 1000–120000
DGW_CONNECTION_REPAVING.DPR_DELAY_MS	Specifies the delay between DPR transmission and forced disconnect. Default value: 100 Range: 100–200

Table 4-45 (Cont.) Add Advanced Settings Configurations

Key	Value
DGW_CONNECTION_REPAVING.DPR_DISCONNECT_CAUSE	Defines the cause included in DPR. Possible values: - REBOOT - BUSY - DO_NOT_WANT_TO_TALK_TO_YOU
DIAMETER.ErrorHandler.MaxRetryCount.Rx.AAR	It is used to set the maximum retries that can be performed for failed AAR messages. Default Value: 1
DIAMETER.ErrorHandler.MaxRetryCount.Rx.STR	It is used to set the maximum retries that can be performed for failed STR messages. Default Value: 1
DIAMETER.ErrorHandler.MaxRetryCount.Rx.RAR	It is used to set the maximum retries that can be performed for failed RAR messages. Default Value: 1
DIAMETER.ErrorHandler.MaxRetryCount.Rx.ASR	It is used to set the maximum retries that can be performed for failed ASR messages. Default Value: 1
DIAMETER.ErrorHandler.CycleBackRetry.Rx.AAR	It is used to set whether peers can be cycled back for retries for AAR messages. Default Value: false
DIAMETER.ErrorHandler.CycleBackRetry.Rx.STR	It is used to set whether peers can be cycled back for retries for STR messages. Default Value: false
DIAMETER.ErrorHandler.CycleBackRetry.Rx.RAR	It is used to set whether peers can be cycled back for retries for RAR messages. Default Value: false
DIAMETER.ErrorHandler.CycleBackRetry.Rx.ASR	It is used to set whether peers can be cycled back for retries for ASR messages. Default Value: false
DIAMETER.ErrorHandler.Enable.UpdateDestinationHost	By default, the BSF Diameter Gateway uses the same Destination-Host AVP in the retry message as the one received in the original request message. To update the Destination-Host AVP in the retry message with the destination peer selected by the error handling configuration, set <code>DIAMETER.ErrorHandler.Enable.UpdateDestinationHost</code> to true. If this parameter is set to false, the retry message retains the original Destination-Host AVP. Default value: false

5. Click **Save** to save the settings.

4.5.2 Peer Nodes

This procedure provides information about how to define and manage Peer Nodes in Diameter Configurations.

The **Peer Nodes** page allows you to create new peer nodes and manage existing peer nodes. The page displays the list of defined configurations and provides the options to import, export, or add data.

To configure Peer Nodes:

1. From the navigation menu, under **BSF**, click **Diameter Configurations**, and select **Peer Nodes**.
This opens the **Peer Nodes** page. The page lists the existing Peer Nodes. You can add or import new nodes using this page.

 Note

Click  **Export** to download the available listings in the JSON file format on your system.

2. Click  **Add** .
This opens the **Create Peer Node** page.
3. On the **Create Peer Node** page, enter values for the available input fields.
The following table describes the various field names:

Table 4-46 Create Peer Node Configurations

Field Name	Description
Name	Unique name of the peer node. Example value: ocs
Type	Defines which type of diameter service must be selected. The values can be • PCF • Application function (AF) • backend • diameter routing agent (dra) • online charging system (ocs) • tdf • udr
Reconnect Limit (sec)	Defines the reconnect limit. Configure this value as the Diameter peer configuration. Example value: 10
Initiate Connection	Set to true to initiate the connection with peer node.
Transport	Defines the type of transport ways for configuring a peer. The values can be: • TCP • TLSv1.2 • TLSv1.3 • TLSv1.2_OR_TLSv1.3
Port	Enter the port number. Enter a number from 0 to 65535. Example value: 8007
Host	Enter the host name. Enter an FQDN, ipv4, or ipv6 address available for establishing diameter transport connections to the peer node.

Table 4-46 (Cont.) Create Peer Node Configurations

Field Name	Description
Realm	Enter the realm name, that is, FQDNs to all of the computers that transact diameter traffic. For example, to add the realm detail of the OCS peer, enter XXX.COM.
Identity	Enter an identity to define a node in a realm. For example, to add the identity detail of the OCS peer, provide value enter OCS.

- Click **Save** to save the changes.

Note

Click **Cancel** to discard the changes.

The value gets listed on the **Peer Node** page. Use  or  available in the next column to update or delete the listing.

Importing Peer Nodes

To import peer node:

- Click  **Import**.
The page opens the **File Upload** dialog box.
- Upload the file in JSON format by using the **Drag and Drop** button.
- Click **Import**.

4.5.3 Diameter Routing Table

Configuration allows routing Diameter request messages to next hop peer based on Diameter application-id, Destination-Realm, and Destination-Host.

When using routing table, there are two ways of configure next hop route:

- Host-Based Routing:** The destination-host of incoming message is checked in the routing table, and then the message is routed to the top priority matching route's peer.
- Realm-Based Routing:** The destination-realm of incoming message is checked in the routing table, and then the message is routed to the top priority matching route's peer.

Routing decision at Diameter-Gateway

Diameter gateway follows below steps in order:

- If the incoming request message has destination-host and the specified peer is directly connected with gateway pod, then the message is routed to the peer specified in destination-host.
- If the incoming request message has destination-host and is not directly connected via any other diameter-gateway pods in the cluster, then the message will be inter-pod routed.
- The routing table is scanned for a matching route by:
 - If the host is reachable, message is sent.

- If the host is not reachable directly, find if it can be reached by another diameter gateway pod, message is sent using inter-pod route
- If the host is not reachable directly or indirectly, lookup the routing table again for the next matching route.

The **Diameter Routing Table Configurations** page displays the Diameter routing configurations. This page allows you to edit the configurations.

To configure the Diameter routing table:

1. From the navigation menu, under **BSF**, click **Diameter Configurations**, and select **Routing Table**.
This opens the **Diameter Routing Table Configurations** page. The page displays the existing configurations.
2. Click  **Edit**.
This opens the **Edit Diameter Routing Table Configurations** page.
3. Expand the **Diameter Route Table** Table group.
The expanded group allows you to add route table entries.
4. To add routing table:
 - a. Click  **Add**.
The page opens the **Add Diameter Route Table** dialog box.
 - b. Enter the values for the following input fields:

Table 4-47 Add Diameter Route Table Configuration

Field Name	Description
Priority	Defines the order of use when one or more routes have overlapping criteria. It can be a number in the range of 0 to 65535. The lowest priority value indicates the highest priority. Note: If there are more than one routing table entry with same priority, it will consider only first row from multiple rows with same priority.
Name	Specifies the unique name of the diameter routing table.
Type	Specifies whether the diameter route table is Host or Realm based.
Realms/Hosts	Specifies the value of the Realms or Hosts depending on the Type selected by the user. For Realms, you can add multiple FQDNs to this field.
Application ID	Specifies the type of application or interface. The available values are • Rx • Gx • Sh • Sy • All Users can select multiple values for this parameter.

Table 4-47 (Cont.) Add Diameter Route Table Configuration

Field Name	Description
Server Identifier	Specifies the server to which the message is to be routed. This identity must also be present in the Identity field of the peer node. Note: If multiple server identifiers are configured one after the other separated by a (,) comma, they will be considered in a sequential manner.

An example of adding the diameter routing details is shown below:

Figure 4-7 Adding first Diameter Routing Element

The screenshot shows a dialog box titled "Edit Diameter Route Table" with a close button (X) in the top right corner. The dialog contains the following fields and values:

- Priority:** A text input field containing "1". Below the field is a hint: "enter a range in 0-65535 number".
- Name :** A text input field containing "one".
- Type:** A dropdown menu with "Host" selected.
- Hosts:** A text input field containing "ocpcf.oracle.com" with a close button (X) on the right.
- Application ID:** A text input field containing "Rx" with a close button (X) on the right.
- Server Identifier:** A text input field containing "oc-diam-gatewav" with a close button (X) on the right.

At the bottom right of the dialog are two buttons: "Save" and "Cancel".

Figure 4-8 Adding Second Diameter Routing Element

Edit Diameter Route Table

Priority: enter a range in 0-65535 number

Name :

Type:

Hosts: X

Application ID: X

Server Identifier: X

- c. Click **Save** on the **Add Diameter Routing Table** dialog box.
5. On the **Edit Diameter Routing Table Configurations** page, expand the **Default Route** group.
6. Enter a value for the **Server Identifier** drop-down list.
The server identifier drop-down list shows the list of the configured peer nodes on the **Peer Nodes** configuration page. For more information on configuring Peer Nodes, see [Peer Nodes](#).

On selecting any of the values, make sure that the name is the same as the value of server identifier.

Note

* (asterisk) wildcard character is allowed in **Hosts**, **Realms**, and **Server Identifier** fields.

Figure 4-9 Diameter Routing Table Configurations

Edit Diameter Routing Table Configurations

Diameter Route Table:

Type ↑↓	Realms ↑↓	Hosts ↑↓	Application ID ↑↓	Server Identifier ↑↓	⊕ Add
host		ocpcf.oracle.com	Rx	oc-diam-gateway	 Edit  Delete
host		ocpcf	Rx	oc-diam-gateway	 Edit  Delete

▼ **Default Route**

Server Identifier: Server Identifier
oc-diam-gateway ▼

7. Click **Save**.
The configuration gets listed on the **Diameter Routing Table Configurations** page.

 Note

Use  or  available under the **Actions** column to update or delete the Diameter Routing Table configurations.

4.5.4 Diameter Error Codes Configurations

This section describes how to customize the Diameter error codes according to the network requirements using the Diameter Error Codes page.

The **Diameter Error Codes** page on CNC Console allows users to view and edit conditions, which are defined by default for the Binding Support Function (BSF) network function. This page also provides the options to import and export diameter error codes.

The following table describes the diameter errors supported by BSF:

Table 4-48 Error Codes and Responses

Condition Name	Error Message	Diameter Result Code
Unable to Route Diameter Message	BSF cannot route a diameter message to the selected destination.	3002 DIAMETER_UNABLE_TO_DELIVER
Diameter Request Message Timeout	BSF sent a diameter request message to the selected destination but did not receive a response in the specified time.	3002 DIAMETER_UNABLE_TO_DELIVER
Unsupported Diameter Interface Received	A diameter message was received for a diameter interface that is not supported by BSF.	3007 DIAMETER_APPLICATION_UNSUPPORTED
AVP Value Invalid	A diameter message was received containing an AVP with a value that is invalid (indicated in FailedAVP)	5004 DIAMETER_INVALID_AVP_VALUE
Required AVP Not Present	A diameter message was received that did not have a required AVP (indicated in FailedAVP)	5005 DIAMETER_MISSING_AVP
Binding Not Found	A binding record was not found for the subscriber key(s) present in the Diameter AAR Initial message	5065 IP-CAN_SESSION_NOT_AVAILABLE
Internal Error	An internal failure occurred	5012 DIAMETER_UNABLE_TO_COMPLY

Table 4-49 Error Codes and Responses

Condition Name	Error Message	Diameter Result Code
Unable to Route Diameter Message	BSF cannot route a diameter message to the selected destination.	3002 DIAMETER_UNABLE_TO_DELIVER
Diameter Request Message Timeout	BSF sent a diameter request message to the selected destination but did not receive a response in the specified time.	3002 DIAMETER_UNABLE_TO_DELIVER
Unsupported Diameter Interface Received	A diameter message was received for a diameter interface that is not supported by BSF.	3007 DIAMETER_APPLICATION_UNSUPPORTED
AVP Value Invalid	A diameter message was received containing an AVP with a value that is invalid (indicated in FailedAVP)	5004 DIAMETER_INVALID_AVP_VALUE
Required AVP Not Present	A diameter message was received that did not have a required AVP (indicated in FailedAVP)	5005 DIAMETER_MISSING_AVP
Binding Not Found	A binding record was not found for the subscriber key(s) present in the Diameter AAR Initial message	5065 IP-CAN_SESSION_NOT_AVAILABLE

Table 4-49 (Cont.) Error Codes and Responses

Condition Name	Error Message	Diameter Result Code
Internal Error	An internal failure occurred	5012 DIAMETER_UNABLE_TO_COMPLY

Table 4-50 Error Codes and Responses

Condition Name	Error Message	Diameter Result Code
Unable to Route Diameter Message	BSF cannot route a diameter message to the selected destination.	3002 DIAMETER_UNABLE_TO_DELIVER
Diameter Request Message Timeout	BSF sent a diameter request message to the selected destination but did not receive a response in the specified time.	3002 DIAMETER_UNABLE_TO_DELIVER
Unsupported Diameter Interface Received	A diameter message was received for a diameter interface that is not supported by BSF.	3007 DIAMETER_APPLICATION_UNSUPPORTED
AVP Value Invalid	A diameter message was received containing an AVP with a value that is invalid (indicated in FailedAVP)	5004 DIAMETER_INVALID_AVP_VALUE
Required AVP Not Present	A diameter message was received that did not have a required AVP (indicated in FailedAVP)	5005 DIAMETER_MISSING_AVP
Binding Not Found	A binding record was not found for the subscriber key(s) present in the Diameter AAR Initial message	5065 IP-CAN_SESSION_NOT_AVAILABLE
Internal Error	An internal failure occurred	5012 DIAMETER_UNABLE_TO_COMPLY

Editing Diameter Error Codes

To edit any of the defined conditions, perform the following steps:

1. From the navigation menu, click **BSF**, then select **Diameter Configurations**, and click **Diameter Error Codes**.
This opens the Diameter Error Codes page that lists the condition names along with their **Result code**, **Vendor Id**, and **Application Error Code**.
2. Click **Edit** against the condition that you need to customize.
This opens the Edit Diameter Error Codes page.
3. Update the required values for the fields, described in the following table:

Table 4-51 Parameters for Edit Diameter Error Codes

Parameter	Description
Use Experimental Result	Indicates whether to use the Result Code AVP (268) or Experimental Result AVP (297) when an error result is generated by BSF.
Result Code	Specifies the Diameter result code for a defined condition. When Use Experimental Result switch is disabled, this field cannot be left blank. Note: The value must be a standard diameter result code as defined in the RFC 6733.
Experimental Result Code	Specifies the custom Diameter result code for a defined condition. When Use Experimental Result switch is enabled, this field cannot be left blank. Note: The value must be a standard diameter result code, from 3000 to 9999, as defined in the <i>3GPP Technical Specification 29.230</i>
Vendor Id	Specifies the Vendor ID of the operator or governing body that manages the code entered by the user in the Experimental Result Code field. When Use Experimental Result switch is enabled, this field cannot be left blank.
Error Message	A message that explains the nature of the error. This error message is only for user understanding and must not be parsed by network entities.

4. Click **Save**.

Importing Diameter Error Codes

To import Diameter error codes, perform the following steps:

1. Click **Import**.
The **File Upload** dialog box opens.
2. Using the **Drag and Drop** button, upload the file in JSON format.
3. Click **Import**.

Exporting Diameter Error Codes

To export Diameter error codes, click **Export**. A file named `bsf.diameter.errorcodes.json` is saved to your device.

4.5.5 Diameter Repaving

The **Diameter Repaving** page displays the configurations related to Diameter Gateway connection repaving. The page allows you to edit the configurations.

To edit Diameter Repaving:

1. From the navigation menu, under **BSF**, click **Diameter Configurations** and select **Diameter Repaving**.
This opens the **Diameter Repaving** page. The page displays the existing Diameter Gateway connection repaving feature configuration.

2. To enable or disable the feature, configure **Enable Diameter Repaving** flag under **Feature Configuration** section.
3. Click **Trigger Disconnection** under **Diameter Repaving Trigger** section to disconnect the existing external Diameter connections established before the trigger timestamp.
4. The **Disconnection Trigger History** section displays the disconnection history, if any.

4.6 Status and Query

This section describes how to retrieve status of BSF profile registration and query sessions using the Session Viewer page.

4.6.1 Session Viewer

You can use the Session Viewer page to query and view PCF binding information for a UE by using any of the following parameters:

- SUPI
- GPSI
- UE Address

To access this screen from the Home screen of CNC Console, click **BSF** and then **Session Viewer**.

To view PCF bindings of a specific UE, perform the following steps:

1. On the **Session Viewer** page, enter the value of SUPI, GPSI, or UE Address.
Query Parameters Session Viewer

Table 4-52 Address

Parameter	Description
IPv4 Address	Specifies the IP addresses in IPv4 format
IPv6 Prefix	Specifies the IPv6 Address Prefix. Note: When you use IPv6 prefix to query a session, ensure that you provide the full notation value. Example: 2011:db8:3c4d:0:0:0:0:0/48
IP Domain	Specifies the IPv4 address domain identifier
MAC Address	Specifies the MAC address, which is formatted as six groups of two hexadecimal digits separated by colons (:) or hyphens (-). For example, in the format hh:hh:hh:hh:hh:hh.

Table 4-53 User

Parameter	Description
SUPI	Specifies the Subscription Permanent Identifier. For example - imsi-450081100100001.
GPSI	Specifies the Generic Public Subscription Identifier. For example - msisdn-9192503899.

Table 4-54 Slice Information/DNN

Parameter	Description
DNN	Specifies the Data Network Name (DNN).
S-NSSAI_SST	Specifies the Slice or Service type for a given S-NSSAI (Single Network Slice Selection Assistance Information).
S-NSSAI_SD	Specifies the Slice Differentiator (SD) for a given S-NSSAI (Single Network Slice Selection Assistance Information). This optional information is used to difference slice or service type across multiple network slices.

2. Click **Query**.

The page displays all the PCF binding records of the UE. For georedundant BSF deployments, the query results include PCF binding data across all sites.

The following screenshot shows the binding IDs when the user searches with GPSI as 9192503855:

Figure 4-10 Query Results using Session Viewer

If binding data is not available, the page displays **No bindings found** message.

Note

As SUPI and GPSI are optional parameters, PCF may not add these values when sending a query request. In such cases, BSF returns a **No session found** message despite binding data being available in the database.

Delete Bindings in Session Viewer

Upon receiving the search results for a query, user may want to delete one or multiple PCF binding IDs on the Session Viewer page.

To delete PCF binding IDs individually, click delete under **Actions** against the required binding ID.

To delete multiple or all the binding records for a subscriber, select binding IDs, then click **Delete** button, and select **YES** on the dialog box.

For georedundant BSF deployments, user can delete PCF binding data across all sites by selecting the binding records from the query results.

4.6.2 BSF NF Data

This section provides information on NF status.

4.6.2.1 BSF Registration Profile

This page lists the BSF profile registered with NRF.

To make updates to any of the parameters of the BSF registration profile, perform the following steps:

1. Click **Edit** button.
2. Update the values of the required parameters.
3. Click **Save**.

To download the BSF profile, click **Download** button. A file named `bsfRegistrationProfile.json` is saved on your system.

4.6.2.2 BSF NRF Status

This page provides the consolidated status of BSF instances registered with NRF.

On the **BSF NRF Status** page, you can view the status of BSF and the NRFs deployed in the cluster.

To view the page:

1. From the navigation menu under **BSF** navigate to **Status and Query**. Click **BSF NRF Data** and select **BSF NRF Status**. This opens the **BSF NRF Status** page.
2. The page displays the registered BSF instances and current NRF health status.

Figure 4-11 BSF Registration Status at NRF





Figure 4-12 BSF NRF Health Status



3. For BSF

You can view the following details for BSF:

- **BSF status with NRF** - It shows whether the BSF instance is **REGISTERED**, **SUSPENDED**, or **DEREGISTERED** with NRF.
- **FQDN** - It shows the FQDN of the BSF registered with NRF.
- **Instance ID** - It shows the unique Instance ID of BSF was registered with NRF.
- **Registration Time** - It shows the time at which BSF was registered with NRF.

If you want to view more details of the BSF instance such as its registration profile, click **View More Details**. It opens the **NF Registration Profile** page.

4. For NRF

You can view the following details for NRF:

- **Health Status** - This ribbon-styled badge shows the health status of the NRF instance. It could be in either healthy or unhealthy state.
- **Primary NRF** - The circular icon with the label P indicates that the NRF is primary.
- **Active Status** - The pulsating green circular icon shows that the NRF is currently active.
- **FQDN** - It shows the FQDN of the NRF.
- **Priority** - It shows the priority of the NRF instances. An NRF instance with priority 1 is treated as primary NRF.
- **Last Connected Time** - It shows the time when BSF last connected with primary NRF.
- **Last Status Change Time** - It shows the time when the NRF status changed.

When the status of an NRF instance changes from healthy to unhealthy, the error reason is also displayed on the page.

The number of NRF's displayed on the page are dynamic and get updated according to the NRF's configured in the network.

5. Based on when the data refreshes, the **Last Refresh Time** on the page is also updated.

In case of any network error, the page displays **Unable to load NRF data** error.

4.6.3 Active Session Query

This section describes the **Active Session Query** tab under **Status and Query** in CNC Console.

This page allows you to query active sessions count instantly.

To get the active sessions count, click **Query Active Sessions**.

Sample output is as follows:

Count: 7032
DateTime: 07-07-2022 11:42:28

4.6.4 Custom Query

This **Custom Query** page allows users to retrieve the migration completion percentage of PCF bindings from the legacy pcf_binding table to the new pcf_binding_v2 table for a specific site_id.

1. Navigate to **Custom Query** page under **Status and Query** in CNC Console.
2. On the **Custom Query** page, configure the following fields:

Table 4-55 Custom Query Fields

Value	Description
Select the service	Specifies the BSF service against which the custom query is run.
Select the query	Specifies the custom query to run. The PCF Binding Migration Completion query retrieves the migration completion percentage of PCF bindings from the legacy pcf_binding table to the pcf_binding_v2 table for a specified site.
Query Description	Displays information about the selected query, including the type of data returned and the query parameters required to retrieve the results. For the PCF Binding Migration Completion query, the description indicates that the query returns the percentage of PCF bindings migrated to the pcf_binding_v2 table for a specified site.
Query Params and Result	This section is used to configure the query parameters.
siteId	Specifies the site identifier for which the PCF binding migration completion percentage is retrieved. The query uses the selected site identifier to determine the number of PCF bindings present in the pcf_binding and pcf_binding_v2 tables and calculates the migration completion percentage.

3. Click **Query** to get the query result. The query result includes the number of non migrated records and the last run time.

Result:

Returns a single numeric value from 0 to 100 representing the percentage of bindings migrated to v2 for the selected site. 0 means none migrated, or no bindings exist. 100 means fully migrated.

4.7 Administration

This section describes how to perform administration tasks, such as bulk import and bulk export of configurable objects into the BSF system.

4.7.1 Import and Export

This section describes how to perform the bulk export or bulk import of the managed objects (MOs) configured on BSF.

You can perform the bulk export and import of BSF data using the following methods:


- **Using CNC Console for BSF:**

BSF provides the GUI to perform bulk export and import of BSF data.

To access the export and import functionality from the CNC Console home page, expand **BSF**, navigate to **Administration**, and select **Import & Export**.

The page displays the **Export** and **Import** tabs. By default, the **Export** tab remains selected. The following screen capture illustrates an example of the **Import & Export** page:

Figure 4-13 Import & Export



You can perform the following operations using the **Import & Export** page:

- [Exporting BSF Configurations](#)
- [Importing BSF Configurations](#)

Note

- Importing the Overload Threshold Profile will be rejected if the CPU validation mentioned under **Configure Threshold Values** section in [Overload Control Threshold](#) fails for any of the three threshold levels.
- The service names mentioned in the json file used to import the Overload Threshold Profile must be same as mentioned in the exported json file.
- The default Overload Threshold Profile cannot be exported or imported. The default Overload Threshold Profile must be customized with a different profile name before exporting.

- **Using REST API for BSF:**

BSF provides REST APIs to bulk export and import BSF data. For more information about REST API configuration, see [Using REST API for BSF Import & Export](#).

Note

The following message priority data that was exported prior to BSF 23.2.0 cannot be imported as the data may be corrupt:

- message containing Sd as interface
- Sy-SLR as condition message

The data with Sd interface or Sy-SLR condition messages that are exported only with BSF 23.2.0 or later versions can be imported.

4.7.1.1 Exporting BSF Configurations

The export functionality allows you to export BSF configurations with the respective data.

The BSF data export is aligned with the left navigation menu options under BSF on the CNC Console. You can export either all the configurations or the configurations of the selected menu options.


To export BSF configuration data::

1. From **BSF**, navigate to **Administration**, and select **Import & Export**. This opens the **Import & Export** page, displaying the **Export** and **Import** tabs. By default, the **Export** tab remains selected.
2. Click



This opens the **Objects to be Exported** dialog box, displaying the list of BSF configurations in a menu tree structure. The dialog box allows you to select the configurations to be exported. The following screen capture displays an illustration of the Objects to be Exported dialog box:

Figure 4-14 Objects to be Exported dialog box



3. In the dialog box, select the configurations according to the export requirement. Selecting or deselecting a parent folder automatically selects or deselects all the child nodes respectively.

Note

To select or deselect all the configurations, click Selection and perform the required operation.

4. Click **Export**.

Note

Click Cancel to discard the export operation.

This starts the export of all the selected configurations. A row is created in the export status table on the **Import & Export** page, displaying the export status with the following details:

- **Export ResourceId**: A new export resource ID is generated for each export operation. You can use this ID to get the export status.
- **Creation TimeStamp**: The timestamp of generation of **Export ResourceId**.
- **Progress (%)**: Shows the export progress in form of a percentage bar. The page auto refreshes the status until the progress reaches 100 percent.
- **Status**: The status of the export operation. It can be any of the following:

- **INIT**: The validation of policies is in progress, and the export of the configurations has not yet started.
- **IN_PROGRESS**: The export is running.
- **DONE**: The export is complete.
- **Actions**: Provides the buttons to download the following:
 - **Export configuration files in ZIP file format**: The exported configurations in ZIP file format. The ZIP file contains the configuration data in JSON file format. You can download the exported data by clicking



under the **Action** column.

- **Export report in TEXT file format**: The export report provides results for each JSON file present in the exported ZIP file. It also provides the reason for failure in case the export of any of the configurations fails.

You can download the export report by clicking  under the **Action** column.

Note

The buttons remain enabled only for the export operation with **DONE** status.

- **Result**: Provides the result of an export operation. This result is available only for the export operations with **DONE** status. Following are the possible values:
 - **SUCCESS**: The export is successful
 - **FAILED**: The export fails
 - **PARTIAL_SUCCESS**: The export is partially successful

4.7.1.2 Importing BSF Configurations

The import functionality allows you to import BSF data configurations. Using this functionality, you can import the same set of BSF data to different BSF systems.

To import BSF data in JSON or ZIP file format:


1. From **BSF**, navigate to **Administration**, and select **Import & Export**.

This opens the **Import & Export** page, displaying the **Export** and **Import** tabs. By default, the **Export** tab remains selected.

2. Select the **Import** tab and click  **Import** .

This opens the **Import Data** dialog box.

Figure 4-15 Import Data



The dialog box titled "Import Data" has a close button (X) in the top right corner. Inside, there is a dashed box labeled "Drag and Drop" with the text "Select a file or drop one here." Below this is a section labeled "Selected file:" containing a text box that says "No items to display." Underneath is a link "Handle collisions between items to import and existing items:" followed by two radio button options: "Ignore records already present in DB from overwrite" (which is selected) and "Replace all the records". At the bottom right are "Import" and "Cancel" buttons.


3. Upload the file in JSON or ZIP format by using the **Drag and Drop** button.
4. Select any of the following options from **Handle collisions between items to import the existing item**.
 - **Attribute Configurations Screens:**

Note

This option is applicable to all the attribute configuration screens such as **Data Sources** under **Data Source Configurations** or **Subscriber Activity**

Logging under **Goggling Configurations** which have **+ Add** button at the top right corner along with a table. In such screens each row is considered as managed object.

Figure 4-16 Sample Attribute Configurations Screen



The screen shows a table with columns: Name, Type, Host, and Actions. Above the table is a search bar labeled "Type to Filter" with a clear button (X). To the right of the search bar are "Refresh" and "Add" buttons. The table has a header row with dropdown arrows for Name, Type, and Host.

Conflict Resolution Strategy

- **Ignore records already present in DB from overwrite** : Ignores records already present in the database and does not overwrite. For each object in the import file,


if the object already exists in the system, the import does not update the object with the configurations provided in the import file. If an object does not exist, then it is added to the system.

- **Replace all the records** : For each object in the import file, if the object already exists in the system, the import replaces the object with the configuration provided in the import file. If an object does not exist, then it is added to the system. If an object existing in the system and is not present in the imported file, it is retained in the system.
- **Service Configuration Screens:**

Note

This option is applicable to all the service configuration screens such as **Management Service** under **Service Configurations** or **Settings** under **Diameter Configurations**, which have  **Edit** button at the top right corner along with a table.

Figure 4-17 Sample Service Configurations Screen



Conflict Resolution Strategy

- **Ignore records already present in DB from overwrite** : Ignores the records and data which are getting imported. For each object in the import file, the imported record will be ignored.
- **Replace all the records**: For each object in the import file, data to be imported will overwrite the existing data in the database, irrespective of whether the object is already present in the database or not.

5. Click Import.

Note

Click **Cancel** to discard the import operation.

This starts the import of configuration objects and their settings to the database. A row is created in the Import status table, displaying the import status with the following additional details:

- **Import ResourceId**: A new import resource ID is generated for each import operation. You can use this ID to get the import status.
- **Creation TimeStamp**: The timestamp of generation of **Import ResourceId**.
- **Progress (%)**: Shows the import progress in form of the percentage bar. The page auto refreshes the status until the progress reaches 100 percent.

- **Status:** The status of the import operation. It can be any of the following:
 - **IN_PROGRESS:** The import is running.
 - **DONE:** The import is complete.
- **Actions:** Provides a button to download the import report in text format. This button gets enabled once the status is **DONE**. The report provides results for each JSON file present in the imported ZIP file. You can download the import report by clicking  under the **Action** column.
- **Result:** Provides the result of an import operation. This result is available only for operations with **DONE** status. Following are the possible values:
 - **SUCCESS:** The import is successful
 - **FAILED:** The import has failed
 - **PARTIAL_SUCCESS:** The import is partially successful

4.7.1.3 Using REST API for BSF Import & Export

This section describes how to perform the bulk export or import of BSF configurations and BSF Data using REST APIs. BSF provides cURL commands for export and import.

cURL Commands for Bulk Import

Import:

```
curl -X POST "http://<ipAddress>:<port>/oc-bsfconfiguration/v1/administration/import"-H
"accept: */*" -H "Content-Type: multipart/form-data" -F "importFile=@<exported zip file
name>;type=application/x-zip-compressed"
```

where,

<exported zip file name> specifies the name of the zip file to be imported.

<ipAddress>:<port> is the host and port where CNC BSF is running.

Import Report:

```
curl -X GET "http://<ipAddress>:<port>/oc-bsfconfiguration/v1/administration/import/
{importResourceId}/report" -H "accept: application/octet-stream"
```

where, *<importResourceId>* is the resource id generated in response to the POST request for import. The ResourceId is the background task id for the POST operations. This id can be used to track the import requests, and download the data.

Import Status:

```
curl -X GET "http://<ipAddress>:<port>/oc-bsfconfiguration/v1/administration/import/
{importResourceId}/status/<importResourceId>/status" -H "accept: application/json"
```

For more information about Bulk Import REST APIs, see "Bulk Import Export Controller" in *Oracle Communications Cloud Native Core BSF REST Specification Guide*.

cURL Commands for Bulk Export

Export All:

```
curl -X POST "http://<ipAddress>:<port>/oc-bsf-configuration/v1/administration/export" -
H "accept: */*" -d""
```

where, *<ipAddress>:<port>* is the host and port where CNC BSF is running.

Export with Managed Objects:

```
curl -X POST "http://<ipAddress>:<port>/oc-bsf-configuration/v1/administration/export/?managedObjects=PCF%20Session%20Management" -H "accept: */*" -d""
```

Download:

```
curl -X GET "http://<ipAddress>:<port>/oc-bsf-configuration/v1/administration/export/<exportResourceId>/download" -H "accept: application/octet-stream"
```

where, *<exportResourceId>* is the resource id generated in response to the POST request for export. The ResourceId is the background task id for the POST operations. This id can be used to track the export requests, and download the data.

Export Report:

```
curl -X GET "http://<ipAddress>:<port>/oc-bsf-configuration/v1/administration/export/<exportResourceId>/report" -H "accept: application/octet-stream"
```

Export Status:

```
curl -X GET "http://<ipAddress>:<port>/oc-bsf-configuration/v1/administration/export/<exportResourceId>/status" -H "accept: application/json"
```

For more information about Bulk Export REST APIs, see "Bulk Import Export Controller" in *Cloud Native Binding Support Function REST Specification Guide*.

4.8 Controlled Shutdown Configurations


This section describes how to perform Controlled Shutdown configurations for Diameter and Ingress interface.

4.8.1 Operational State

To change Operational State of a site:

- From the navigation menu, under **BSF**, click **Controlled Shutdown**, and then select **Operational State**. This opens the page displaying the two groups, **Switch Operational State** and **Operational State History**:

Figure 4-18 Operational State



- Switch Operational State It displays the following operational states:
 - * **NORMAL:** NF instance is discoverable and services all requests.
 - * **PARTIAL SHUTDOWN:** NF instance is non discoverable and no new session creation requests accepted.
 - * **COMPLETE SHUTDOWN:** NF instance is non discoverable and no new session creation requests accepted.

Note

By default, NORMAL state is assigned to a site. The current state of any site can be identified with a tick mark.

You can switch to a different operational state by clicking the NORMAL, PARTIAL SHUTDOWN, or COMPLETE SHUTDOWN button.

- Operational State History: It displays the history of the operational states along with the timestamp.

Note

It displays maximum of ten records at a time. On scrolling further, another set of ten records is displayed. The maximum number of record maintained is hundred.

4.8.2 Diameter Error Mapping

To configure Diameter Error Mapping, perform the following steps:

1. From the navigation menu under **BSF**, click **Controlled Shutdown** and then select **Diameter Error Mapping**.
This opens the **Diameter Error Mapping** page. The page lists the existing configurations. You can add or import new diameter error mapping configurations using this page.

Note

Click  **Export** to download the available listings in the JSON file format on your system.

2. Click  **Add** .

This opens the Create Diameter Error Mapping page.

3. On the **Create Diameter Error Mapping** page, enter the following information:

Table 4-56 Create Diameter Error Mapping

Field Name	Description
Message Type	Type of the request

Table 4-56 (Cont.) Create Diameter Error Mapping

Field Name	Description
Answer with Result Code	Specifies the result code, returned in the answer response, when request message is rejected as part of congestion control. Users can select any of the following values from the drop-down menu: • DIAMETER_TOO_BUSY • DIAMETER_UNABLE_TO_COMPLY • DIAMETER_UNABLE_TO_DELIVER • CUSTOM_RESULT_CODE Note: When the CUSTOM_RESULT_CODE value is selected, the following two fields are populated on the page: • Result Code: Enter a custom result code. • Use Experimental Result Code: This is disabled by default. You can enable it by clicking the icon against it. When it is enabled, Vendor ID field is populated on the page: – Vendor ID: Enter a valid value to specify vendor ID.

4. Click **Save**.
The configuration gets listed on the **Diameter Error Mapping** page. The page defines the Diameter Error Mapping configuration in the BSF database and it is available to be used in a BSF.

 Note

Use  or  available under the **Actions** column to update or delete the configuration.

Importing Diameter Error Mapping

To import Diameter Error Mapping configuration:

1. Click  **Import**.
The page opens the **File Upload** dialog box.
2. Upload the file in JSON format by using the **Drag and Drop** button.
3. Click **Import**.

4.8.3 SBI Ingress Error Mapping

To configure SBI Ingress Error Mapping, perform the following steps:

1. From the navigation menu, under **BSF**, click **Controlled Shutdown**, and then select **SBI Ingress Error Mapping**.
This opens the SBI Ingress Error Mapping page.
2. Click **Edit**.
This opens the Edit SBI Ingress Error Mapping page.

- Click  **Add** .

This opens the Add SBI Ingress Error Mapping page.

- Enter values for the available input fields as described in the following table:

Table 4-57 Ingress Error Mapping Configurations

Field Name	Description
Id	Specifies the list of IDs available for BSF.
Error Code Profile	Select an error code profile from the dropdown list. It displays the list of error profiles configured using the SBI Ingress Error Code Profiles Collection.

- Click **Save** to save the Ingress error mapping.
To discard the changes, click **Cancel**.

The value gets listed on the SBI Ingress Error Mapping page. Use  or  available under the **Actions** column to update or delete the profile.

4.9 Overload and Congestion Control Configurations

This section describes how to perform overload and congestion control configurations.

To use the Error Code Profiles page to create and manage error code profiles in Overload Control Configurations for Diameter Gateway and SBI interface, see *SBI Ingress Error Code Profiles Collection* section.

Note

When overload control feature is enabled, it should be enabled for both Diameter gateway and SBI interface. The overload control manager needs data from both Ingress Gateway and Diameter Gateway to determine the overall load on BSF management. It is not possible to enable overload control for Diameter Gateway and disable the same in Ingress Gateway.

4.9.1.1 Load Shedding Profiles

This procedure provides information about how to create and manage load shedding profiles in Diameter Configurations.

The Load Shedding Profiles page allows you to create new and manage existing load shedding profiles. The page displays the list of defined profiles and provides the options to import and export data as well.

To configure Load shedding profiles, perform the following steps:

- From the navigation menu, under **BSF**, click **Diameter Configurations**, and select **Load Shedding Profiles**.
This opens the Load Shedding Profiles page.

- Click  **Add** .

This opens the Create Load Shedding Profiles page.

3. Enter values for the available input fields described in the following table:

Table 4-58 Load Shedding Profiles Configurations

Field Name	Description
Name	Unique name of the load shedding profile.
Scheme	Allows to configure the discard policy based on • Priority: to discard messages based on priority range • Priority and Percentage: to discard messages based on priority range and percentage for each range
Type	Defines the type of load shedding profile. You can select any of the following values from the drop-down list: • Congestion Control • Overload Control

To add load shedding rules for the profile type congestion control, perform the following steps:

- a. Under **Load Shedding Rules**, click  **Add** .
This opens the **Add Load Shedding Rules** dialog box.
- b. Enter values for the available input fields, described in the following table:

Table 4-59 Load Shedding Rules Configurations When the Selected Scheme is "Priority"

Field Name	Description
State	This field appears when the Type of load shedding profile is Congestion Control. Specifies the type of state for which the rule is being defined. Select any of the following values using the drop-down: • Danger of Congestion • Congested
Discard Priority	This field appears when the Scheme of load shedding profiles is Priority. Specifies the discard priority for the load shedding rule. The discard priority value can be a number from 0 to 15. Any request message with equal or greater priority is rejected.

Table 4-59 (Cont.) Load Shedding Rules Configurations When the Selected Scheme is "Priority"

Field Name	Description
Answer with Result Code	Specifies the result code, returned in the answer response, when request message is rejected as part of congestion control. Users can select any of the following values from the drop-down menu: • DIAMETER_TOO_BUSY • DIAMETER_UNABLE_TO_COMPLY • DIAMETER_UNABLE_TO_DELIVER • EXPERIMENTAL_RESULT_CODE Note: When the EXPERIMENTAL_RESULT_CODE value is selected, the following two fields are populated on the page: • Result Code: Enter a custom result code. • Vendor ID: Enter a valid value to specify vendor ID.

Table 4-60 Load Shedding Rules Configurations When the Selected Scheme is "Priority and Percentage"

Field Name	Description
State	This field appears when the Type of load shedding profile is Congestion Control. Specifies the type of state for which the rule is being defined. Select any of the following values using the drop-down: • Danger of Congestion • Congested
Discard Priority Percentage	Allows to configure the Discard Priority Percentage as explained in the following step.

- c. To configure discard Priority Percentage, perform the following steps:
- Under **Discard Priority Percentage**, click  **Add** .
This opens the **Add Discard Priority Percentage** dialog box.
Enter values for the available input fields, described in the following table:

Table 4-61 Adding Discard Priority Percentage

Field Name	Description
Priority Range	Specifies the discard priority range for the load shedding rule. The discard priority value can be a number from 0 to 15. Any request message with equal or greater priority will be rejected based on the percentage discard configured.
Discard Percentage	Specifies the discard percentage for the specified priority range.

Table 4-61 (Cont.) Adding Discard Priority Percentage

Field Name	Description
Answer with Result Code	Specifies the result code, returned in the answer response, when request message is rejected as part of congestion control. Users can select any of the following values from the drop-down menu: • DIAMETER_TOO_BUSY • DIAMETER_UNABLE_TO_COMPLY • DIAMETER_UNABLE_TO_DELIVER • EXPERIMENTAL_RESULT_CODE Note: When the EXPERIMENTAL_RESULT_CODE value is selected, the following two fields are populated on the page: • Result Code: Enter a custom result code. • Vendor ID: Enter a valid value to specify vendor ID.

ii. Click **Save** to save the discard priority percentage.

d. Click **Save** to save the load shedding rule.

OR

To add load shedding rules for the profile type overload control, perform the following steps:

- a. Under **Load Shedding Rules**, click  **Add** .
This opens the **Add Load Shedding Rules** dialog box.
- b. Enter values for the available input fields, described in the following table:

Table 4-62 Load Shedding Rules Configurations When the Selected Scheme is "Priority"

Field Name	Description
Level	Specifies the name of the level. The name specified in this parameter must match the level name in Ingress Gateway's ocdiscardpolicies. Select any of the following values using the drop-down list: • L1 • L2 • L3 Note: If the load levels are not configured, then level transitions will not happen and it will stay at the same level. Also, any existing L4 level data will be removed, as L4 is not supported.
Discard Priority	Specifies the discard priority for the load shedding rule. The discard priority value can be a number from 0 to 15. Any request message with equal or greater priority is rejected.

Table 4-62 (Cont.) Load Shedding Rules Configurations When the Selected Scheme is "Priority"

Field Name	Description
Answer with Result Code	Specifies the result code, returned in the answer response, when request message is rejected as part of overload control. Users can select any of the following values from the drop-down menu: • DIAMETER_TOO_BUSY • DIAMETER_UNABLE_TO_COMPLY • DIAMETER_UNABLE_TO_DELIVER • EXPERIMENTAL_RESULT_CODE Note: When the EXPERIMENTAL_RESULT_CODE value is selected, the following two fields are populated on the page: • Result Code: Enter a custom result code. • Vendor ID: Enter a valid value to specify vendor ID.

Table 4-63 Load Shedding Rules Configurations When the Selected Scheme is "Priority and Percentage"

Field Name	Description
Level	Specifies the name of the level. The name specified in this parameter must match the level name in Ingress Gateway's ocdiscardpolicies. Select any of the following values using the drop-down list: • L1 • L2 • L3 Note: If the load levels are not configured, then level transitions will not happen and it will stay at the same level. Also, any existing L4 level data will be removed, as L4 is not supported.
Discard Priority Percentage	Configure the discard priority percentage as explained in the following step.

- c. To configure Discard Priority Percentage, perform the following steps:
- Under **Discard Priority Percentage**, click  **Add** .
This opens the **Add Discard Priority Percentage** dialog box.
Enter values for the available input fields, described in the following table:

Table 4-64 Adding Discard Priority Percentage

Field Name	Description
Priority Range	Specifies the discard priority range for the load shedding rule. The discard priority value can be a number from 0 to 15. Any request message with equal or greater priority will be rejected based on the percentage discard configured.
Discard Percentage	Specifies the discard percentage for the specified priority range.
Answer with Result Code	Specifies the result code, returned in the answer response, when request message is rejected as part of congestion control. Users can select any of the following values from the drop-down menu: • DIAMETER_TOO_BUSY • DIAMETER_UNABLE_TO_COMPLY • DIAMETER_UNABLE_TO_DELIVER • EXPERIMENTAL_RESULT_CODE Note: When the EXPERIMENTAL_RESULT_CODE value is selected, the following two fields are populated on the page: • Result Code: Enter a custom result code. • Vendor ID: Enter a valid value to specify vendor ID.

- ii. Click **Save** to save the discard priority percentage.
- d. Click **Save** to save the load shedding rule.

 Note

You can not edit or delete the load shedding rule for an existing load shedding profile. As a workaround, you must create a new load shedding profile with the updated load shedding rule and delete the old load shedding profile.

4. Click **Save** to save the load shedding profile.
To discard the changes, click **Cancel**

The value gets listed on the Load Shedding Profiles page. Use  or  available under the **Actions** column to update or delete the profile.

Importing Load Shedding Profiles

To import load shedding profiles, perform the following steps:

1. Click  **Import** .
The page opens the **File Upload** dialog box.
2. Upload the file in JSON format by using the **Drag and Drop** button.
3. Click **Import**.

Exporting Load Shedding Profiles

To export load shedding profiles, click **Export**. A json file is saved to your device.

4.9.1.2 Message Priority Profiles

This procedure provides information about creating and managing message priority profiles in Diameter Configurations.

The Message Priority Profiles page allows you to create new and manage existing message priority profiles. The page displays the list of defined profiles and provides the options to import and export data.

To configure Message Priority profiles, perform the following steps:

1. From the navigation menu, under **BSF**, click **Overload Control Configurations**, select **Diameter**, and then select **Message Priority Profiles**.
This opens the Message Priority Profiles page.

2. Click  **Add** .

This opens the Create Message Priority Profiles page.

3. Enter values for the available input fields, described in the following table:

Table 4-65 Message Priority Profiles Configurations

Field Name	Description
Name	Unique name of the message priority profile.

To add message priority rules for the profile, perform the following steps:

- a. Under **Message Priority Rules**, click  **Add** .
This opens the **Add Message Priority Rules** dialog box.
- b. Enter values for the available input fields, described in the following table:

Table 4-66 Message Priority Rules Configurations

Field Name	Description
Name	Specifies the unique name of the message priority rule.
Message Priority	Specifies the priority assigned to the message. It can be a number from 0 to 15.
Rule Priority	Specifies the priority assigned to the message priority rule.
Use DRMP Priority	When this switch is enabled, the priority for the message rule is assigned from DRMP AVP.

Table 4-67 Message Priority Rules Configurations - Conditions

Field Name	Description
Application	Specifies the type of application. Users can select the following value from the drop-down: Rx
Message	Specifies the type of message for the selected application. The supported message values for each application type are as follow: For Rx application, choose a value from AAR, STR, RAR, and ASR.

- c. To add pre-defined AVP conditions, click  **Add** under **Add Pre Defined AVP Conditions**. On the **Add Pre Defined AVP Conditions** dialog box, select Name and enter values as described in the following table:

Table 4-68 Pre Defined AVP Conditions Configurations

Name	Values
Called-Station-Id	This AVP can be used only when the application type is specified a Rx. Users can enter multiple comma-separated values. Note: BSF supports wildcard format for this AVP.
Rx-Request-Type	This AVP can be used for Rx application with message specified as AAR. You can select any of the following valid values from the drop-down list: - INITIAL_REQUEST - UPDATE_REQUEST - PCSCF_RESTORE_REQUEST
Service-URN	This AVP indicates that an AF session is used for emergency traffic. It is of type OctetString. Examples: "sos", "sos.fire", "sos.police" and "sos.ambulance". Note: BSF supports wildcard format for this AVP.
MPS-Identifier	This AVP indicates that an AF session relates to an MPS session and contains the national variant for MPS service name. It is of type OctetString. Example: NGN GETS
MCPTT-Identifier	This AVP includes either one of the namespace values used for MCPTT and may include the name of the MCPTT service provider. It is of type OctetString.
MCVideo-Identifier	This AVP includes the name of the MCVideo service provider. It is of type OctetString.
Reservation-Priority	This AVP is of type Enumerated and is specified in an AA-Request as the main AVP to associate a priority with a resource reservation or modification request. You can specify a value from 0 to 7 for this AVP.

Table 4-68 (Cont.) Pre Defined AVP Conditions Configurations

Name	Values
DRMP	For the Operator, You can select any of the following valid values from the drop-down list: • Equals • Not Equals • Less Than • Less Than Equals • Greater Than • Greater Than Equals For the Values, You can enter any numerical value between 1-15.
OC-Message-Priority	For the Operator, You can select any of the following valid values from the drop-down list: • Equals • Not Equals • Less Than • Less Than Equals • Greater Than • Greater Than Equals For the Values, You can enter any numerical value between 1-31.

Click **Save** to save the pre-defined AVP conditions for the message priority rule.

- d. Click **Save** to save the message priority rule.
4. Click **Save** to save the message priority profile.
To discard the changes, click **Cancel**

The value gets listed on the Message Priority Profiles page. Use  or  available under the **Actions** column to update or delete the profile.

Importing Load Shedding Profiles

To import message priority profiles, perform the following steps:

1. Click  **Import** .
The page opens the **File Upload** dialog box.
2. Upload the file in JSON format by using the **Drag and Drop** button.
3. Click **Import**.

Exporting Message Priority Profiles

To export load shedding profiles, click **Export**. A json file is saved to your device.

4.9.2.1.1 Rate Limiting Policy

This procedure provides information about how to use the Rate Limiting Policy page to manage rate limiting policies for overload control on SBI interface.

To configure rate limiting policy, perform the following steps:

1. From the navigation menu, under **BSF**, click **Overload Control Configurations**, select **SBI**, then select **Rate Limiting**, and then select **Rate Limiting Policy**.

This opens the Rate Limiting Policy page.

- Click **Edit**.
This opens the Edit Rate Limiting Policy page.
- Enter values for the available input fields as described in the following table:

Table 4-69 Rate Limiting Policy Configurations

Field Name	Description
Enable Rate Limiting	Specifies whether to enable or disable rate limiting.
Sampling Period (in milliseconds)	Specifies the time frame for each cycle of rate limiting per service. Its default value is 200 ms.

- Under **Rate Limit Policy**, click  **Add** .
This opens the Add Rate Limit Policy dialog box.
- Enter values for the available input fields as described in the following table:

Table 4-70 Rate Limit Policy Configurations

Field Name	Description
Name	Specifies the name of the rate limit policy that is further used to determine a mapping between route and discard policy name per route.
Discard Priority	Specifies the discard priority for the rate limiting policy. Any request with message priority higher in value than the discard priority is rejected.
Action	Specifies the action taken when requests are discarded. Currently, the only supported value is <code>RejectWithErrorCode</code> .
Scheme	Specifies the scheme for applying rate limiting. Currently, the only supported value is <code>PriorityBased</code> .
Error Code Profile	Specifies the list of error code profiles configured on the Error Code Profiles page.

- Click **Save** to save the rate limit policy. The value gets listed under the **Rate Limit Policy** group.
To discard the changes, click **Cancel**.

Use  or  available under the **Actions** column to update or delete any given policy.
- Click **Save** to save the rate limiting policy.
To discard the changes, click **Cancel**.

4.9.2.1.2 Route Level Mapping

This procedure provides information about how to use the Route Level Mapping page to manage route level mapping for overload control on SBI interface.

To configure route level mapping, perform the following steps:

- From the navigation menu, under **BSF**, click **Overload Control Configurations**, select **SBI**, then select **Route Level Mapping**.

- This opens the Route Level Mapping page.
- Click **Edit**.
This opens the Edit Route Level Mapping page.
 - Under **Route Configuration**, click  **Add** .
This opens the Add Route Configuration dialog box.
 - Enter values for the available input fields as described in the following table:

Table 4-71 Add Route Configuration

Field Name	Description
Id	Specifies the list of route IDs available for BSF. Choose any value from the drop-down list: - BSF Management Register - BSF Management Deregister - BSF Management Discovery
Failure Request Count Profile	Specifies the mapping between error series and failure count. This field allows to select one of the Profiles from the Error Code Series List created on IGW page under Service Configurations .

For more details about the mapping between error series and failure count, see *Error Code Series List at Ingress gateway* sections in *Oracle Communications Cloud Native Core, BSF REST Specification Guide*.

- Under **Rate Limiting**, click  **Add** .
This opens the Add Method dialog box.
- Enter values for the available input fields as described in the following table:

Table 4-72 Add Method Configurations

Field Name	Description
Http Method	Specifies the HTTP method. Depending on the value select for Id , you can select any of the following values from the drop-down list: - POST - PUT - GET - DELETE - PATCH
Message Rate (per sampling period)	Specifies the message rate per sampling period for a given method.
Rate Limit Policy	Select a rate limit policy from the drop-down list. It displays the list of rate limit policies configured using the Rate Limiting Policy page.

- Click **Save** to save the method. The value gets listed under the **Rate Limiting** group. To discard the changes, click **Cancel**.

Use  or  available under the **Actions** column to update or delete any given policy.

8. Click **Save** to save the route configuration.
To discard the changes, click **Cancel**.
9. Click **Save** to save the route level mapping.
To discard the changes, click **Cancel**.

4.9.2.2.1 Discard Policy Mapping

This procedure provides information about how to use the Discard Policy Mapping page to manage discard policy mapping in Overload Control Configurations for SBI interface.

To configure Discard Policy Mapping, perform the following steps:

1. From the navigation menu, under **BSF**, click **Overload Control Configurations**, select **SBI**, then select **Overload Control**, and then select **Discard Policy Mapping**. This opens the Discard Policy Mapping page.
2. Click Edit.
This opens the Edit Discard Policy Mapping page.
3. Enter values for the available input fields as described in the following table:

Table 4-73 Discard Policy Mapping Configurations

Field Name	Description
Enable Overload Control	Specifies whether to enable or disable overload control.
Sampling Period (in milliseconds)	Specifies the time frame for each cycle of overload control per service. Its default value is 200 ms.

4. Under **Mappings**, click  **Add** .
This opens the Add Mappings dialog box.
5. Enter values for the available input fields as described in the following table:

Table 4-74 Mappings Configurations

Field Name	Description
Service Name	Specifies the name of the microservice that is further used to determine a mapping between service and discard policy name per service. BSF Management is the only supported value for this field.
Policy Name	Specifies the name of the discard policy that is used to determine a mapping between service and discard policy name per service. The drop-down list shows the policies configured using the Discard Policy page.

6. Click **Save** to save the mappings.
To discard the changes, click **Cancel**

The value gets listed under the **Mappings** group on the Discard Policy Mapping page. Use  or  available under the **Actions** column to update or delete the mappings.
7. Click **Save** to save the discard policy mapping.

To discard the changes, click **Cancel**.

4.9.2.2.2 Discard Policy

This procedure provides information about how to use the Discard Policy page to manage discard policies for overload control for SBI interface.

To configure discard policy, perform the following steps:

1. From the navigation menu, under **BSF**, click **Overload Control Configurations**, select **SBI**, then select **Overload Control**, and then select **Discard Policy**. This opens the Overload Control Discard Policy page.
2. Click Edit.
This opens the Edit Overload Control Discard Policy page.
3. Click  **Add** .
This opens the Add Discard Policies page.
4. Enter values for the available input fields as described in the following table:

Table 4-75 Discard Policy Configurations

Field Name	Description
Name	Specifies the unique name of the discard policy.
Scheme	Specifies the criteria of dropping requests for a microservice. It could be either priority based or percentage based. If you select the value as Priority Based, configure the values of the parameters under Priority Based Policies. For Percentage Based scheme, configure the parameters under Percentage Based Policies.

To add priority based policies, perform the following steps:

- a. Under **Priority Based Policies**, click  **Add** .
This opens the **Add Priority Based Policies** dialog box.
- b. Enter values for the available input fields as described in the following table:

Table 4-76 Priority Based Policies Configurations

Field Name	Description
Level	Specifies the name of the level. The name specified in this parameter must match the level name in Ingress Gateway's ocdiscardpolicies. Select any of the following values using the drop-down list: • L1 (Load Level 1) • L2 (Load Level 2) • L3 (Load Level 3)
Discard Priority	Specify the discard priority for the discard policy rule. Any request message with equal or lower message priority is rejected. Note: 1 is considered as the highest message priority.

Table 4-76 (Cont.) Priority Based Policies Configurations

Field Name	Description
Error Code Profile	Select an error code profile from the drop-down list. It displays the list of error profiles configured using the Error Code Profile page.
Action	Specifies the action taken when selected requests are rejected. Currently, it only supports the action to reject requests based on error code.

OR

Table 4-77 Percentage Based Policies Configurations

Field Name	Description
Level	Specifies the name of the level. The name specified in this parameter must match the level name in Ingress Gateway's ocdiscardpolicies. Select any of the following values using the drop-down list: • L1 (Load Level 1) • L2 (Load Level 2) • L3 (Load Level 3)
Discard Percentage	Specify the discard percentage for the policy rule. The specified percentage of the calculated rate for service in previous sampling period is discarded in current sampling period.
Error Code Profile	Select an error code profile from the drop-down list. It displays the list of error profiles configured using the Error Code Profile page.
Action	Specifies the action taken when selected requests are rejected. Currently, it only supports the action to reject requests based on error code.

- c. Click **Save** to save the discard policy.
To discard the changes, click **Cancel**.
5. Click **Save** to save the overload control discard policy.
To discard the changes, click **Cancel**

The value gets listed on the Overload Control Discard Policy page. Use  or  available under the **Actions** column to update or delete any given policy.

4.9.3 Congestion Control

4.9.3.1 Settings

The **Settings** page displays the general configurations related to the Congestion Control. The page allows you to edit the configurations.

To edit settings:

1. From the navigation menu under **BSF**, click **Overload and Congestion Control** and select **Congestion Control** and then select **Settings**.
This opens the **Settings** page. The page displays the existing configurations for different BSF services.
2. Click  **Edit**.
This opens the **Edit Settings** page.
3. Enter values for the available input fields as described in the following table:

Table 4-78 Edit Configurations

Field Name	Description
Service Name	Specifies the BSF service name. You can check the default configuration for the all fields and edit as necessary.
Enable	Specifies whether to enable or disable Congestion Control for the selected BSF service. Default value: False
State Change Sample Count	This specifies after how many continuous intervals, the pod state can be changed. This count can range from 1 to 2147483647. Default Value: 5
State Calculation Interval (in milliseconds)	This specifies that after this time duration or interval, the pod congestion state will be re-verified. This interval in milliseconds can range from 50 to 2147483647. Default Value: 200

4. Perform the following steps to configure **Advanced Settings** for configuring algorithm that should be used for calculating the CPU and Queue usage for Congestion Control feature:
 - a. Click the  **Add**.
The page opens the **Add Advanced Settings** dialog box.
 - b. In the dialog box, enter the following **keys** and respective **values**:
The following table describes the keys and values:

Table 4-79 Add Advanced Settings Configurations

Key	Value
CPU_CONGESTION_RESOURCE_CALC_ALGORITHM	This enables the congestion control CPU resource calculation to use the algorithm provided. If invalid value is provided then default method is used to calculate the CPU usage. Allowed Value: EMA
CPU_CONGESTION_RESOURCE_CALC_INTERVAL_MS	This configuration sets the time interval in milliseconds for CPU congestion resource to be calculated by the configured algorithm. Default Value: 50

Table 4-79 (Cont.) Add Advanced Settings Configurations

Key	Value
QUEUE_CONGESTION_RESOURCE_CALC_ALGORITHM	This enables the congestion control Queue resource calculation to use the algorithm provided. If invalid value is provided then default method is used to calculate the Queue usage. Allowed Value: EMA
QUEUE_CONGESTION_RESOURCE_CALC_INTERVAL_MS	This configuration sets the time interval in milliseconds for Queue congestion resource to be calculated by the configured algorithm. Default Value: 70
EMA_CPU_CONGESTION_RESOURCE_CALC_RATIO	This configuration will be used to set the EMA ratio in which congestion resource(s) calculated using EMA algorithm to be considered for calculation. Default Value: 70:30
EMA_QUEUE_CONGESTION_RESOURCE_CALC_RATIO	This configuration will be used to set the EMA ratio in which congestion resource(s) calculated using EMA algorithm to be considered for calculation. Default Value: 70:30


- Click **Save** to save the congestion control configuration. To discard the changes, click **Cancel**.

The value gets listed on the Settings page.

Diameter Gateway Migration of Congestion Control Configurations

After upgrade to BSF 25.1.200, perform migration of Congestion Control configurations using CNC Console, follow these steps:

- From the navigation menu under **BSF**, click **Overload and Congestion Control** and select **Congestion Control** and then either selects **Thresholds** or **Load Shedding Rules**, then the following warning message is displayed on all the respective pages: The following picture illustrates the Diameter Gateway congestion control data migration warning message:
The following picture illustrates the Diameter Gateway congestion control data migration warning message:

Figure 4-19 Diameter Gateway Data Migration Warning

- Click on **Migrate**. This opens the Migration Dialog box, to migrate the Load Shedding Profiles, Message Priority Profiles, and Threshold data.

Note

In the configuration data, if the load shedding profiles are not configured or available for the Diameter Gateway then only Migrate Threshold Data page will be displayed.

3. The Migrate Load Shedding Profiles Data page, displays all the load shedding profiles created using the previous version.
 - a. In the dialog box select values for the available input fields as described in the following table:

Table 4-80 Migrate Load Shedding Profiles Data, with scheme as Priority

Field Name	Description
Load Shedding Profile	Choose a profile from the drop down lists of previous versions load shedding profiles.
Discard Priority	Under Discard Priority Click  to edit DANGER_OF_CONGESTION, CONGESTION_L1, CONGESTION_L2, and CONGESTED states, and configure their respective Discard Priority values.

- b. Click **Next** to continue with the data migration process. This opens the Migrate Load Shedding Profiles Response Code Data page.
 - c. Click **Cancel** to discontinue the data migration process.
4. The Migrate Load Shedding Profiles Data page, displays all the load shedding profiles created using the previous version.
 - a. In the dialog box select values for the available input fields as described in the following table:

Table 4-81 Migrate Load Shedding Profiles Data, with scheme as PriorityAndPercentage

Field Name	Description
Load Shedding Profile	Choose a profile from the drop down lists of load shedding profiles.
Discard Priority	Under Discard Priority Click  to edit CONGESTION_L1 and CONGESTION_L2 states and configure the respective Discard Priority values. Click  to view DANGER_OF_CONGESTION and CONGESTED states respective Discard Priority values.

- b. Click **Next** to continue with the data migration process. This opens the Migrate Load Shedding Profiles Response Code Data page.
 - c. Click **Cancel** to discontinue the data migration process.

- Configure values for the **Congestion Response Codes** input fields as described in the following table:

Table 4-82 Congestion Response Codes

Field Name	Result Code	Vendor Id
DIAMETER_TOO_BUSY	3004	0
DIAMETER_UNABLE_TO_COMPLY	5012	0
DIAMETER_UNABLE_TO_DELETE	3002	0
EXPERIMENTAL_RESULT_CODE	3006	6

- Click **Next** to continue with the data migration process. This opens the Migrate Message Priority Profiles Data page.
- Configure values for the **Migrate Message Priority Profiles Data** input fields as described in the following table:

Table 4-83 Migrate Message Priority Profiles Data

Field Name	Description
Message Priority Profile	Choose a profile from the drop down lists of message priority profiles. This profile will be used for both Congestion Control and Overload Control features.

- Click **Cancel** to discontinue the data migration process.
- Click **Back** to the previous page.
- Click **Next** to continue with the data migration process. This opens the Migrate Message Threshold Data page.
- Configure values for the **Migrate Message Threshold Data** input fields as described in the following table:


Table 4-84 Migrate Message Threshold Data

Field Name	Description
Name of the Profile	Choose a profile from the drop down lists of message threshold profiles.
Threshold Data	Under Threshold Data Click  to edit DANGER_OF_CONGESTION, CONGESTION_L1, CONGESTION_L2, and CONGESTED states and configure their respective Queue and CPU values.

- Click **Migrate** to perform the data migration. This will migrate all the configuration data for Congestion Control to the 25.1.200 versions, and the feature is not enabled for Diameter Gateway services.
- Click **Migrate and Enable** to perform the data migration and enable Congestion Control for Diameter Gateway services.
- Click **Cancel** to discontinue the data migration process.


15. Click **Back** to the previous page and perform any updates or changes.
16. Data migration is successful, the message *migration of old congestion data of diameter-gateway is successful* is shown.

Figure 4-20 Migration of old Diameter Gateway congestion data success status



17. On data migration failure, the message *migration of old congestion data of diameter-gateway failed* is shown.

Figure 4-21 Migration of old congestion data of diameter-gateway failure status



18. To verify if the Diameter Gateway Congestion Control configurations migration is successful,
 - a. From the navigation menu under **BSF**, click **Overload and Congestion Control** and select **Settings**.
This opens the **Settings** page. The page displays the default Congestion Control configurations for the selected services along with other BSF services in the drop-down list, and now the migration message is not displayed.

4.9.3.2 Thresholds

BSF Console allows you to configure the congestion control threshold values for a BSF service using a threshold profile. You can either use a default that is set during BSF deployment or create a new profile using CNC Console. You must activate one of the profiles to use the threshold values. At any time, you can activate only one profile. To open **Congestion Control Thresholds** page:

1. From the navigation menu under **BSF** navigate to **Overload and Congestion Control**, select **Congestion Control** and select **Thresholds**.

This opens the **Congestion Control Thresholds** page. The page shows the existing configurations for a specific BSF service. You can add a new threshold profile configuration or copy the existing profile and modify it.

2. Select from the drop down list fields, as described in the following table:

Table 4-85 Thresholds Configuration

Field Name	Description
Service Name	Provides a drop down lists of BSF services. You can select the service for which the congestion control thresholds are to be configured.

Table 4-85 (Cont.) Thresholds Configuration

Field Name	Description
Threshold Profiles	Provides a drop down lists of all the configured threshold profiles. A default congestion control threshold profile is provided during BSF deployment. You cannot either edit or delete this profile. If the selected profile is already active, you can see Active button in green color next to the profile name.

3. Click **Go To Active** button. The current active profile from the **Threshold Profile** drop down list is selected.
4. Click  **Add** .
This opens the **Create Profile** window and allows you to create a new profile.
 - a. Enter values for the available input fields, described in the following table:

Table 4-86 Create Profile Configurations

Field Name	Description
Service Name	Specifies the selected BSF service for which a new threshold profile is to be created.
Name of the profile	Provide a unique name for the threshold profile to be created. The name can only contain the characters A–Z, a–z, 0–9, period (.), hyphen (-), and underline (_). The maximum length is 255 characters.

- b. Click **Create** to create the threshold profile.
To discard the changes, click **Cancel**.
- c. The value gets listed on the Congestion Control Thresholds page.
By default, the newly created profile has the same values as the default profile configuration values.
- d. Click  to view the Queue and CPU values for the various congested states.
- e. Click  to edit each of the various congested states **Queue** and **CPU** values. This opens the **Edit <Congestion States>** page, where <Congestion States> will be different for different services.
The BSF Management service pod's can be in following Congestion states:

Table 4-87 BSF Management Service Congestion States with Default CPU and Queue values

Congestion States	CPU Percentage	Queue (Pending Requests)
DANGER_OF_CONGESTION (DOC)	75	175
CONGESTION_L1	80	200
CONGESTION_L2	85	225

Table 4-87 (Cont.) BSF Management Service Congestion States with Default CPU and Queue values

Congestion States	CPU Percentage	Queue (Pending Requests)
CONGESTED	90	250

The Diameter Gateway service pod's can be in following Congestion states:

Table 4-88 Diameter Gateway Service Congestion States with Default CPU and Queue values

Congestion States	CPU Percentage	Queue (Pending Requests)
DANGER_OF_CONGESTION (DOC)	60	6300
CONGESTION_L1	70	7200
CONGESTION_L2	75	8100
CONGESTED	80	9000

- f. Enter the following information:

Table 4-89 Edit <Congestion State> Fields

Field Name	Description
Service Name	Specifies the selected service for which threshold profile configurations are to be edited.
Profile Name	Specifies the selected threshold profile for which DOC type queue and CPU values are to be edited.
Queue	Specifies the Queue's pending requests for BSF services across all the congestion levels. The range is between 1 to 2147483647.
CPU	Specifies the CPU percentage across all the congestion levels. The number can be between 1 to 100.

- g. Click **Save** to save the configurations.
To discard the changes, click **Cancel**.

The updated value gets listed on the Congestion Control Thresholds page.

5. Click . This opens the **Copy Profile** window to copy an existing profile. While using the copy profile option, whichever profile is selected in the **Threshold Profile** drop down, the values are copied from that profile.
6. Click . This deletes the selected threshold profile. User cannot delete the default and active threshold profiles.

4.9.3.3 Load Shedding Rules

The Congestion Load Shedding Rules page allows you to create new and manage existing congestion load shedding rules.

To configure Congestion Load Shedding rules, perform the following steps:

1. From the navigation menu under **BSF** navigate to **Overload and Congestion Control**, select **Congestion Control** and then select **Load Shedding Rules**.
This opens the Congestion Load Shedding Rules page. The page shows the existing configurations for a specific Policy service. You can add a new congestion control load shedding rule or copy the existing rule and modify it.
2. Select from the drop down list fields, as described in the following table:

Table 4-90 Congestion Load Shedding Rules

Field Name	Description
Service Name	Provides a drop down lists of BSF services. You can select the service for which the congestion load shedding rules are to be configured.
Scheme Name	Select the schemes to configure load shedding rules.
Load Shedding Rules	Provides a drop down lists of all the configured load shedding rules. A default load shedding rule is provided during BSF deployment. You cannot edit or delete this rule. If the selected profile is already active, you can see Active button in green color next to the profile name.

3. Click **Go To Active** button. The current ACTIVE rule from the **Load Shedding Rules** drop down list is selected.
4. Click  **Add** .
This opens the **Create rule** window and allows you to create a new rule.
 - a. Enter values for the available input fields, described in the following table:

Table 4-91 Create Rule Configurations

Field Name	Description
Service Name	Specifies the selected BSF service for which a new load shedding rule is to be created.
Name of the rule	Provide a unique name for the load shedding rule to be created. The name can only contain the characters A–Z, a–z, 0–9, period (.), hyphen (-), and underline (_). The maximum length is 255 characters.

- b. Click **Create** to create the load shedding rule.
To discard the changes, click **Cancel**.
- c. The value gets listed on the Congestion Load Shedding Rules page.
By default, the newly created rule has the same values as the default rule configuration values.
- d. Click  to view the Discard Priority values for the various congested states.
The BSF Management service pod's can be in following Congestion states:

Table 4-92 BSF Management Service Congestion States with Default Discard Priority

Congestion States	Discard Priority
DANGER_OF_CONGESTION (DOC)	30
CONGESTION_L1	24
CONGESTION_L2	18
CONGESTED	14

The Diameter Gateway service pod's can be in following Congestion states:

Table 4-93 Diameter Gateway Service Congestion States with Default Discard Priority

Congestion States	Discard Priority
DANGER_OF_CONGESTION (DOC)	27
CONGESTION_L1	19
CONGESTION_L2	17
CONGESTED	15

- e. Click  to edit each of the various congested states **Discard Priority** value. This opens the **Edit <Congestion States>** page, where <Congestion States> will be different for different services.
- f. Enter the following information:

Table 4-94 Edit <Congestion State> Fields

Field Name	Description
Service Name	Specifies the selected BSF service for which the load shedding rule is edited.
Rule Name	Specifies the selected load shedding rule for the chosen Congestion state.
Discard Priority	Specifies the discard priority that can be configured for the selected Congestion state. This number can be between 0 to 31.

- g. Click **Save** to save the configurations.
To discard the changes, click **Cancel**.

The updated value gets listed on the Congestion Load Shedding Rules page.

5. Click  . This opens the **Copy Rule** window to copy an existing profile. While using the copy rule option, whichever rule is selected in the **Load Shedding Rules** drop down, the values are copied from that rule.
6. Click  . This deletes the selected load shedding rule. User cannot delete the default and active load shedding rules.

4.9.4 Overload Control Threshold

To open Overload Control Threshold page:

1. From the navigation menu under **BSF**, navigate to **Overload Control Configurations**, and select **Overload Control Threshold**.

Overload Control Threshold page is displayed. The page shows the existing configurations.

2. BSF allows you to configure the threshold values using a profile. You can either use a default or custom profile or create a new profile. You must activate one of the profiles to use the values. At a time, you can activate only one profile.

Note

If there are no profiles activated, Policy calculates the load level based on `overloadLevelThreshold` configured in `custom-values.yaml` file.

If you are upgrading from an older version of BSF to 23.1.0 or later, the following message appears at the top of the page:

System looks to be using old configuration (Not profile based). Request you to please migrate the data or activate one of the profile

- a. Click **Migrate**.
Migrate Data window opens.
- b. Enter the name of the profile.
- c. Either click **Migrate** to migrate the data from the previous version of BSF or click **Migrate and Activate** to migrate the data and activate the profile.

Note

Migrate button will only migrate the existing threshold data to profile. It does not activate the profile. Until one of the profile is not active, the UI will continue to show the above message to migrate the data.

In case of fault recovery, the active profile details are recovered and used from the database. No need to migrate again.

Table 4-95 Create and Manage Threshold Profiles

Button/icon	Action
 Add	Opens the Create Profile window and allows you to create a new profile. By default, the values of the default profile are associated with the new profile. Click  next to each of the services to edit the values.
	Deletes the selected threshold profile. Note: You cannot delete the default system provided profile.
	Opens the Copy Profile window to copy an existing profile. While using the copy profile option, whichever profile is selected in the Threshold Profile dropdown, the values are copied from that profile.

Table 4-95 (Cont.) Create and Manage Threshold Profiles

Button/icon	Action
Threshold Profile	Lists all the threshold profiles. You can select any of the profiles from the list to activate. If the selected profile is already active, you can see (ACTIVE) button in green color next to the profile name.
Activate	Activates the profile selected from Threshold Profile list. At a time, you can have only one active profile.
Go To Active	Go To Active button selects the current active profile from the Threshold Profile drop down list. Note: If none of the profiles are active, the Go To Active button does not appear.

Table 4-96 Configure Threshold Values

Configuration	Description
Service Name	You can configure the threshold values for BSF Management Services:
CPU	Click  to view the abatement and onset values for each of the three levels (L1, L2, and L3). The onset and abatement values for CPU are calculated in percentage (%) and the range is from 1 to 100. Click  to edit the abatement and onset values for each of the levels. Note Make sure that the onset value of L1 is less than the abatement value of L2 and the onset value of L2 is less than the abatement value of L3. You can click  to delete the CPU values for a service. The application prompts you to confirm before deleting the values. Note: If the above mentioned CPU validation fails for any of the threshold levels, importing of the Overload Threshold Profile will be rejected.

Table 4-96 (Cont.) Configure Threshold Values

Configuration	Description
Pending Message Count	Click  Add to add the pending message count for each of the services. Pending message count accepts an integer value between 1 to 1000000.
Failure Count	Click  Add to add the failure count for each of the services. The failure count accepts an integer value between 1 to 1000000.
Memory	Click  Add to add the abatement and onset memory values for each of the three levels (L1, L2, and L3). Memory details are calculated in Percentage (%) and ranges between 1 to 100.

4.10 NF Scoring Configurations

You can configure the NF Scoring feature using the CNC Console. To navigate to NF Scoring, click NF Scoring, under BSF. It shows Settings and Calculated Score, which are described as follows:

Table 4-97 NF Scoring

Field Name	Description
Enable NF Scoring	Specifies whether to enable or disable the NF Scoring
TPS	Traffic
Enable	Enables the TPS.
Max Score	Specifies the maximum score of the TPS.
Max TPS	Specifies the maximum TPS.
Service Health	Specifies the service health of a site.
Enable	Enables the Service Health.
Max Score	Specifies the maximum score of the Service Health.
Signaling Connections	Specifies the Signaling Connections of a site.
Enable	Enables the Signaling Connections.
Max Score	Specifies the maximum score of the Signaling Connections.
Max Connections	Specifies the maximum connections. Range: 1 to 10000
Replication Health	Specifies the Replication Health of a site.
Enable	Enables the Replication Health.
Max Score	Specifies the maximum score of the Replication Health.
Locality/Site Preference	Specifies the Locality or Site Preference.


Table 4-97 (Cont.) NF Scoring

Field Name	Description
Enable	Enables the Locality or Site Preference.
Score	Specifies the score of the Locality or Site Preference.
Active Alert	Specifies the Active Alerts of a site.
Enable	Enables the Active Alert.
Critical Alert Weightage	The site with more critical alerts is unhealthy.
Major Alert Weightage	The site with more major alerts is unhealthy.
Minor Alert Weightage	The site with more minor alerts is unhealthy.


Calculated Score:

To check the calculated Score of a site:

From the navigation menu, under BSF, click NF Scoring and then select Calculated Score. This opens the page displaying the Calculated Score.

Figure 4-22 NF Scoring Disabled

If the NF Scoring feature is enabled. The calculated Screen page displays as:

Figure 4-23 NF Scoring Enabled

Note

If app-info pod is down, you will not be getting the NF Score. You will get an error message that "Data can't be fetched due to internal server error".

Calculated Score shows the total score along with the Instance ID. The total score is shown in either Green or Orange color. If the NF Score is shown in green color there are no failed factors. And, if the NF Score is shown in Orange color there are one or more failed factors. You can click on **Detailed information** to view different criteria and their Max Score, Evaluated Score, and More Info. The criteria in the detailed information tab show the evaluated score. The failed factors are shown with a warning symbol under the evaluated score.

On the top-right of the screen, the *Last Refresh Time* information is available. Moreover, a *Refresh* button is given to refresh the NF Score of a site.

4.11 Viewing cnDBTier APIs in CNC Console

Perform the following procedure to view the cnDBTier version, status of cnDBTier clusters, and georeplication status on the CNC Console.

Note

The following cnDBTier APIs are read only and is available only through CNC Console.

4.11.1 Backup List

Perform the following procedure to view the list of completed backups:

1. From the navigation menu, click **BSF**, and then click **DB Tier**.
2. Click **Backup List** to view the list of completed backups along with Backup ID, Backup size, and Creation Timestamp.
The **Backup List** page is displayed.

Table 4-98 Backup List

Fields	Description
Backup Details	This field displays information such as backup Id, backup size, and backup creation timestamp.
Site Name	This field displays the name of the current site to which BSF is connected.
Backup Id	This field displays the ID of the stored backup.
Backup Size (bytes)	This field displays the size of the stored backup.
Creation TimeStamp	This field displays the time recorded when the backup was stored.

4.11.2 cnDBTier Backup Status

1. From the navigation menu, click **BSF**, and then click **DB Tier**.

- Click **cnDBTier Backup Status** to view the cnDBTier backup status. The **cnDBTier Backup Status** page is displayed. On the **cnDBTier Backup Status** page, click **Refresh** to view the latest cnDBTier backup status as described in the following table:

Table 4-99 cnDBTier Backup Status

Attribute	Description
Current Timestamp	This field displays the current system time.
Is Backup In Progress	This field displays whether any backup is in progress or not.
Next Backup Scheduled Timestamp	This field displays the time for the next scheduled backup.

4.11.3 Database Statistics Report

Perform the following procedure to view the available database:

- From the navigation menu, click **BSF**, and then click **DB Tier**.
- Click the **Database Statistics Report** to view the available database.

Table 4-100 Database Statistics Report

Fields	Description
Database Count	This field displays the number of available database.
Database Tables Count	This field displays the available database names and their table count.
Database Name	This field displays the database name.
Table Count	This field displays the table count for each database.
Database Table Rows Count	This field displays the table rows present in each table.

- Click the **View** icon available next to the database name to view the **View Database Table Rows Count** screen.

Table 4-101 View Database Table Rows Count

Fields	Description
Database Name	This field displays the database name.
Tables	This field displays the table names and the corresponding rows in each table.
Table Name	This field displays the table name.
Row Count	This field displays the table rows present in each table.

4.11.4 Georeplication Recovery

Perform the following procedure to mark cnDBTier cluster as failed, execute georeplication recovery, and monitor their status:

- From the navigation menu, click **BSF**, and then click **DB Tier**.
- Click **Georeplication Recovery** to access the Georeplication Recovery Status of the cnDBTier cluster. This includes options such as Update Cluster As Failed, Start Georeplication Recovery, and Georeplication Recovery Status.

- Click **Update Cluster As Failed** to mark the cluster as FAILED.
The **Update Cluster As Failed** page is displayed.

Table 4-102 Update Cluster As Failed

Fields	Description
Cluster Names	This field lists the names of cnDBTier clusters. It allows you to select a cluster from the available list to mark as failed.
Failed Cluster Names	This field displays the list of names of cnDBTier clusters that are marked as failed.
Update Cluster	Click Update Cluster to mark the cluster as FAILED.
Cancel	Click Cancel to cancel the process.

- Click **Start Georeplication Recovery** to start the georeplication recovery process for a failed site.
The **Start Georeplication Recovery** page is displayed.

Table 4-103 Start Georeplication Recovery

Attribute	Description
Failed Cluster Name	This field displays a list of cnDBTier clusters that have been marked as failed.
Backup Cluster Name (Optional)	This field displays the list of active cnDBTier clusters designated for georeplication recovery.
Start Georeplication Recovery	Click Start Georeplication Recovery to start the georeplication recovery process for a failed site.
Cancel	Click Cancel to cancel the process.

- Click the **Georeplication Recovery Status** to view the status of georeplication recovery for cnDBTier clusters.
The **Georeplication Recovery Status** page is displayed.

Table 4-104 Georeplication Recovery Status

Attribute	Description
Local Cluster Name	This field displays the name of the local cluster.
Georeplication Recover Status Details	This field displays the details of the georeplication recovery status for cnDBTier clusters.
Cluster Name	This field displays the names of all the clusters.
Georeplication Recovery Status	This field displays the current status of georeplication recovery for a cluster.
Refresh	Click Refresh to view the most current data.

The following are the states of Georeplication Recovery:

Table 4-105 Georeplication Recovery States

Georeplication Recovery State	Description
ACTIVE	The cluster is in a healthy state, and replication is up and running with its respective mate cluster.
REINSTALLED	The cluster enters this state during fatal error recovery when the end user reinstalls the cluster.

Table 4-105 (Cont.) Georeplication Recovery States

Georeplication Recovery State	Description
STARTDRRESTORE	When Georeplication recovery is started, the cluster will transition into this state.
INITIATEBACKUP	Once Georeplication recovery is started, the cluster will identify a healthy cluster for backup initiation and transition into this state.
CHECKBACKUP	Once the backup is initiated, the georeplication recovery cluster will monitor the progress of the backup until its completion. If the backup fails, the cluster will restart the backup.
COPY_BACKUP	Upon completion of the backup, the georeplication recovery cluster will request the transfer of the backup from the healthy cluster to the georeplication recovery cluster.
CHECK_BACKUP_COPY	Once backup copy is started georeplication recovery cluster will monitor for the backup transfer progress till it's completion and if it's fails the cluster will re-initiates the backup transfer.
BACKUPCOPIED	Once the backup copy is started, the georeplication recovery cluster will monitor the progress of the backup transfer until its completion. If the transfer fails, the cluster will restart the backup transfer.
BACKUPEXTRACTED	This state indicates that the backup has been successfully extracted at the georeplication recovery cluster, allowing the restoration of the backup to start.
FAILED	This state is used by end user to mark specific cluster as failed and hence georeplication recovery is essential to recover the cluster. This state can also indicates that georeplication recovery started and the database is restored using the healthy cluster backup.
UNKNOWN	This state is used by the end user to mark a specific cluster as failed, necessitating georeplication recovery for cluster recovery. Additionally, this state can indicate that georeplication recovery has started and the database has been restored using the backup from the healthy cluster.
RECONNECTSQLNODES	This state is used to instruct SQL nodes to be offline during backup restoration to prevent any records from entering the binlog of the georeplication recovery cluster.
BACKUPRESTORE	This state indicates that the backup, successfully copied from the healthy cluster, is currently being used to restore the georeplication recovery cluster.
RESTORED	Once the backup is successfully restored in the georeplication recovery cluster, the cluster will enter this state to start the reestablishment of replication channels.
BINLOGINITIALIZED	This state indicates the start of binlogs for the restoration of replication channels, necessary to start the restore process
RECONFIGURE	Once the binlog is restarted, the georeplication recovery cluster will reestablish the replication channels with respect to all its mate clusters.

4.11.5 Georeplication Status

Perform the following procedure to view the local site and remote site name to which BSF is connected.

1. From the navigation menu, click **BSF**, and then click **DB Tier**.
2. Click **Georeplication Status** to view the local site and remote site name to which BSF is connected.

Table 4-106 GeoReplication Status

Fields	Description
Local Site Name	This field displays the local site name to which BSF is connected.
Remote Site Name	This field displays the remote site name.
Replication Status	This field displays the replication status with corresponding sites.
Seconds Behind Remote Site	This field displays the number of seconds that the last record read by the local site is behind the latest record written by the remote site for all the replication groups.

- a. Click the **View** icon in the **Actions** menu to view the **View Georeplication Status** screen.

Table 4-107 Georeplication Status

Fields	Description
Replication Group Delay	This field displays the number of seconds that the last record read by the local site is behind the latest record written by the remote site for individual replication groups.
Replication Channel Group Id	This field displays the ID of the replication channel group.

- b. Click the **View** icon to view the **Replication Group Delay** attributes.

Table 4-108 View Replication Group Delay

Fields	Description
Channel Details	This field displays the channel details such as Remote Replication IP and Role.
Remote Replication IP	This field displays the IP of the remote replication channel.
Role	This field displays the role of the replication channel IP.

4.11.6 Georeplication Status Across All Sites

Perform the following procedure to retrieve the georeplication status across all configured sites, offering a consolidated view of the overall replication health:

1. From the navigation menu, click **BSF**, and then click **DB Tier**.
2. Click **Georeplication Status Across All Sites** to access the georeplication status across all configured sites in the cnDBTier cluster.
The **Georeplication Status Across All Sites** page is displayed.

Table 4-109 Georeplication Status Across All Sites

Fields	Description
Site Name	This field displays the primary site for which the georeplication status is being reported.
Remote Site Name	This field displays the remote site involved in replication with the local site.
Site To Mate Site Replication Status	This field displays the overall georeplication status with all the groups.
Replication Details Per Group ID	This field displays the list of replication-group-specific details with the site name and the remote site name.

3. Click the  option to view the **View Georeplication Status Across All Sites** attributes.
The **View Georeplication Status Across All Sites** page is displayed.

Table 4-110 View Replication Details Per Group ID

Attribute	Description
Replication Group ID	This field displays the unique identifier assigned to the replication group.
Replication Channel Group Status	This field displays the replication status for the specific replication group.
All SQL Status Details	This field displays the detailed SQL and replication state information for all SQL nodes for the specified replication group.

4. Click the  option to view the **View All SQL Status Details** attributes.
The **View All SQL Status Details** page is displayed.

Table 4-111 View All SQL Status Details

Attribute	Description
SQL Node	This field displays the SQL node whose replication metrics are being reported.
Source Host	This field displays the remote host from which the SQL node pulls replication data.
External IP	This field displays the external endpoint through which the SQL node is accessible within the site.
Replica I/O Running Status	This field displays whether the replication I/O thread is currently active or not. Possible values: Yes or No.
Replica SQL Running Status	This field displays whether the SQL thread responsible for applying replicated changes is running or not. Possible values: Yes or No.
Server ID	This field displays the unique server identifier used in the replication topology.
Role	This field displays the role of the SQL node within the replication setup (for example, ACTIVE or STANDBY).
Last Error Code	This field displays the most recent SQL thread error code encountered during replication.

Table 4-111 (Cont.) View All SQL Status Details

Attribute	Description
Last Error Message	This field displays the most recent error message reported by the SQL thread.
Last I/O Error Code	This field displays the most recent I/O thread error code encountered during replication.
Last I/O Error Message	This field displays the most recent error message reported by the I/O thread.
Last SQL Error Code	This field displays the most recent SQL execution error code encountered while applying replicated transactions.
Last SQL Error Message	This field displays the last SQL execution error message reported.
SQL Remaining Delay	This field displays the replication lag or delay (if any) between the source site and this SQL node. Unit: seconds.

4.11.7 cnDBTier Health

1. From the navigation menu, click **BSF**, and then click **DB Tier**.
2. Click **cnDBTier Health** to view the health status of the microservices like replication, backup manager, monitor services, and NDB services.
 - Click **Backup Manager Health Status** to view the health status of the backup manager.
The **Backup Manager Health Status** page is displayed.

Note

The parameters listed in the below table are read-only.

Table 4-112 Backup Manager Health Status

Fields	Description
Service Name	This attribute displays the service name of the backup manager microservice.
Service Status	This attribute displays the service status of the backup manager microservice. Possible values are UP and DOWN.
DB Connection Status	This attribute displays the database connection status of the backup manager microservice. Possible values are UP and DOWN.
Overall Backup Manager Service Health	This attribute displays the overall health status of the backup manager microservice. Possible values are UP and DOWN.

Table 4-112 (Cont.) Backup Manager Health Status

Fields	Description
Backup Executor Health Status	This attribute displays the following information like node id and DB connection status of the backup executor. - Node Id: This attribute displays the id of the node. - DB Connection Status: This attribute displays the backup executor database connection status with the nodes. Possible values are UP and DOWN.

- Click **Monitor Health Status** to view the health status of the services. The **Monitor Health Status** page is displayed.

Note

The parameters listed in the below table are read-only.

Table 4-113 Monitor Health Status details

Attribute	Description
Service Name	This attribute displays the service name of the monitor microservice.
DB Connection Status	This attribute displays the database connection status of the monitor microservice. Possible values are UP and DOWN.
Metric Scrape Status	This attribute displays the status of the metric scrape, that is if the metrics are fetched or not. If the metrics are fetched then the service is up and vice versa. Possible values are UP and DOWN.
Overall Monitor Service Health	This attribute displays the overall health status of the monitor microservice. Possible values are UP and DOWN.

- Click **NDB Health Status** to view the health status of the network database. The **NDB Health Status** page is displayed.

Note

The parameters listed in the below table are read-only.

Table 4-114 NDB Health Status details

Attribute	Description
Local Site Name	This attribute displays the name of the current site. For example, site 1, site 2.

Table 4-114 (Cont.) NDB Health Status details

Attribute	Description
NDB Health Status Details	This attribute displays the health status of the network database like name of the NDB service, status of the service, health status of PVC.
Service Name	This attribute displays the service name. For example, ndbmcmd-0, ndbmtid-0, ndbmyappsqld-1, ndbmysqld-2.
Service Status	This attribute displays the status of the service. Possible values are UP and DOWN.

- Click **Replication Health Status** to view the health status of the replication sites. The **Replication Health Status** page is displayed.

 Note

The parameters listed in the below table are read-only.

Table 4-115 Replication Health Status details

Attribute	Description
Local Site Name	This attribute displays the name of the current site (site 1, site 2).
Health Status Details	This attribute displays the health status details of the local site like replication service name, replication service status, database connection status of the replication service, and the overall health status of the replication microservices. The number of rows in this table varies depending on the type of deployment (for example, two-site, three-site deployments).
Service Name	This attribute displays the name of the available replication service.
Service Status	This attribute displays the status of the available replication service. Possible values are UP and DOWN.
DB Connection Status	This attribute displays the database connection status of the replication microservice. Possible values are UP and DOWN.
Overall Replication Service Health	This attribute displays the overall health status of the replication microservice. Possible values are UP and DOWN.

4.11.8 Heartbeat Status

Perform the following procedure to view the connectivity between local site and remote site name to which BSF is connected.

1. From the navigation menu, click **BSF**, and then click **DB Tier**.
2. Click **HeartBeat Status** to view the connectivity between local site and remote site to which BSF is connected.

Table 4-116 HeartBeat Status Details

Fields	Description
Site Name	This field displays the name of the current site to which BSF is connected.
HeartBeat Details	This field displays information such as the remote site name, heartbeat status, heartbeat lag, and replication channel group id.
Remote Site Name	This field displays the remote site name.
Heartbeat Status	This field displays the connectivity status with corresponding sites.
Heartbeat Lag	This field displays the lag or latency in seconds it took to synchronize between sites.
Replication Channel Group Id	This field displays the ID of the replication channel group.

4.11.9 IP Family Preference

1. From the left navigation menu, navigate to **DBTier** and then click **IP Family Preference**. The **IP Family Reference** page is displayed.
2. Click **Actions** to view the IP preference set for the sites.
 - a. The **Edit IP Preference** page is displayed.

Table 4-117 IP Family Preference

Fields	Description
Local Site Name	This attribute displays the name of the current site to which BSF is connected.
Remote Site Name	This attribute displays the remote site name.
Preferred Family	This attribute displays the current site's preferred IP family type. IPV4- Indicates the IP Type selected is IPV4. IPV6- Indicates the IP Type selected is IPV6.
Status	This attribute displays the status of the IP family switchover. Values are RUNNING/IN-PROGRESS.

4.11.10 Local Cluster Status

Perform the following procedure to view the local cluster status for the current site.

1. From the navigation menu, click **BSF**, and then click **DB Tier**.
2. Click **Local Cluster Status** to view the local cluster status for the current site:

Table 4-118 Local Cluster Status

Fields	Description
Site Name	This field displays the name of the current site to which BSF is connected.
Cluster Status	This field displays the local cluster status for the current site.

4.11.11 On Demand Backup

Perform the following procedure to backup the database on demand.

1. From the navigation menu, click **BSF**, and then click **DB Tier**.
2. Click **On Demand Backup** to create a new backup and view the status of initiated on-demand backups.

Note

On Demand Backup can be initiated on both single site and multisite cnDBTier cluster and can be used to restore the first standalone site. DB Backup will not be initiated if sites are not properly configured.

Table 4-119 On Demand Backup Details

Fields	Description
Site Name	This field displays the name of the current site to which BSF is connected.
DR Status	This field displays the status of DR.
Backup Id	This field displays the ID of the stored backup.
Backup Status	This field displays the status of backup.
Remote Transfer Status	The field displays the status of remote transfer.
Initiate Backup	The field displays whether the backup is initiated or not.

- a. Click the **Edit** icon.
The **Edit** On Demand Backup screen appears.

Note

The **Edit** mode is available only for Initiate Backup.

- b. To enable the Initiate Backup option, click **Save**.
A confirmation message "Save successfully" appears.
- c. Click **Cancel** to navigate back to the On Demand Backup screen.
- d. Click **Refresh** to reload the On Demand Backup screen.

4.11.12 Version

Perform the following procedure to view the cnDBTier version:

1. From the navigation menu, click **BSF**, and then click **DB Tier**.

2. Click **cnDBTier Version** to view the version.

Table 4-120 cnDBTier Version Attributes

Fields	Description
cnDBTier Version	This field displays the cnDBTier version.
NDB Version	This field displays the network database (NDB) version.

4.11.13 Gracefully Stop and Start Georeplication

Stop Georeplication

1. From the left navigation menu, navigate to **DBTier** and then click **Gracefully Stop and Start Georeplication**. The **Gracefully Stop and Start Georeplication** page is displayed.
2. Click **Stop Georeplication** from the left navigation menu to stop the georeplication between the sites. The **Stop Georeplication** page is displayed.

Table 4-121 Stop Georeplication

Fields	Description
Site Name	This attribute displays the name of the current BSF site.
Remote Site Name	This attribute displays the list of the remote BSF site names.
Replication Channel Group Ids	This attribute displays the list of replication group ids of the current site.
Georeplication Details	This attribute displays the details of the georeplication sites and their status.

Table 4-122 Georeplication Details

Fields	Description
Site Name	This attribute displays the name of the current BSF site.
Remote Site Names	This attribute displays the name of the remote BSF site.
Replication Channel Group Id	This attribute displays the replication group id.
Replication Status	This attribute displays the status of the georeplication between sites. • UP- Indicates the georeplication between the sites is up. • DOWN- Indicates the georeplication between the sites is down.
Switch Over Status	This attribute displays the switch-over status of the georeplication between sites. • SWITCH_OVER_ENABLED- Indicates the switch-over status of the georeplication between the sites is up. • SWITCH_OVER_DISABLED- Indicates the switch-over status of the georeplication between the sites is disabled.

3. Click **Stop Georeplication** to stop the replication between the sites.

Start Georeplication

1. Click **Gracefully Stop and Start Georeplication**. The **Gracefully Stop and Start Georeplication** page is displayed.
2. Click **Start Georeplication** from the left navigation menu to start the georeplication between the sites.
The **Start Georeplication** page is displayed.

Table 4-123 Start Georeplication

Fields	Description
Site Name	This attribute displays the name of the current site to which BSF is connected.
Remote Site Name	This attribute displays the list of the remote site name.
Replication Channel Group Ids	This attribute displays the list of replication group ids of the current site.
Georeplication Details	This attribute displays the details of the georeplication sites and their status.

Table 4-124 Georeplication Details

Fields	Description
Site Name	This attribute displays the name of the current site to which BSF is connected.
Remote Site Names	This attribute displays the name of the remote site.
Replication Channel Group Id	This attribute displays the replication group id.
Replication Status	This attribute displays the status of the georeplication between sites. • UP- Indicates the georeplication between the sites is up. • DOWN- Indicates the georeplication between the sites is down.
Switch Over Status	This attribute displays the switch-over status of the georeplication between sites. • SWITCH_OVER_ENABLED- Indicates the switch-over status of the georeplication between the sites is up. • SWITCH_OVER_DISABLED- Indicates the switch-over status of the georeplication between the sites is disabled.

3. Click **Start Georeplication** to start the replication between the sites.

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 .

5.1.24 PCF_BINDING_ERROR_RATE_ABOVE_CRITICAL_THRESHOLD

Table 5-25 PCF_BINDING_ERROR_RATE_ABOVE_CRITICAL_THRESHOLD

Field	Details
Description	PCF Binding Error Rate above 10 Percent in {{ \$labels.microservice }} in {{ \$labels.namespace }}
Summary	PCF Binding Error Rate in {{ \$labels.kubernetes_node }} (current value is: {{ \$value }})
Severity	Critical
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 >= 10
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 10% 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 5-26 INGRESS_CREATE_ERROR_RATE_ABOVE_MINOR_THRESHOLD

Field	Details
Description	BSF Ingress Create Error Rate above 1 Percent in {{ \$labels.microservice }} in {{ \$labels.namespace }}
Summary	Transaction Create Error Rate in {{ \$labels.kubernetes_node }} (current value is: {{ \$value }})
Severity	Minor

Table 5-26 (Cont.) INGRESS_CREATE_ERROR_RATE_ABOVE_MINOR_THRESHOLD

Field	Details
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 >= 1
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 5-27 INGRESS_CREATE_ERROR_RATE_ABOVE_CRITICAL_THRESHOLD

Field	Details
Description	BSF Ingress Create Error Rate above 10 Percent in {{labels.microservice}} in {{labels.namespace}}
Summary	Transaction Create Error Rate in {{labels.kubernetes_node}} (current value is: {{ \$value }})
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 5-28 INGRESS_CREATE_ERROR_RATE_ABOVE_MAJOR_THRESHOLD

Field	Details
Description	BSF Ingress Create Error Rate above 5 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	Major
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 >= 5</pre>
OID	1.3.6.1.4.1.323.5.3.37.1.2.4
Metric Used	<code>http_server_requests_seconds_count</code>
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 5-29 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

Table 5-29 (Cont.) INGRESS_DELETE_ERROR_RATE_ABOVE_MINOR_THRESHOLD

Field	Details
Expression	$\frac{\text{sum by(namespace) (rate(http\_server\_requests\_seconds\_count\{microservice='bsf-management-service', status!= '204', method='DELETE', uri='/nbsf-management/v1/pcfBindings/{bindingId}'}\{24h\}))}}{\text{sum by(namespace) (rate(http\_server\_requests\_seconds\_count\{microservice='bsf-management-service', method='DELETE', uri='/nbsf-management/v1/pcfBindings/{bindingId}'}\{24h\}))}} * 100 \geq 1$
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 5-30 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
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

Table 5-30 (Cont.) INGRESS_DELETE_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 DELETE method. For any assistance, contact My Oracle Support.

5.1.30

INGRESS_DELETE_ERROR_RATE_ABOVE_CRITICAL_THRESHOLD

Table 5-31 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
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 5-32 DB_TIER_DOWN_ALERT

Field	Details
Description	DB cannot be reachable!
Summary	DB cannot be reachable!

Table 5-32 (Cont.) DB_TIER_DOWN_ALERT

Field	Details
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 5-33 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
Metric Used	cgroup_cpu_usage

Table 5-33 (Cont.) CPU_USAGE_PER_SERVICE_ABOVE_MINOR_THRESHOLD

Field	Details
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 5-34 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 5-34 (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 5-35 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 5-35 (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 5-36 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 5-36 (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 5-37 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 5-38 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 5-39 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 5-39 (Cont.) NRF_COMMUNICATION_FAILURE

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

5.1.39 NRF_SERVICE_REQUEST_FAILURE

Table 5-40 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 5-41 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 5-41 (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 5-42 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 5-42 (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 5-43 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 5-43 (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 5-44 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 5-45 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 5-46 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 5-46 (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 5-47 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 5-48 POD_PENDING_REQUEST_DOC

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

Table 5-48 (Cont.) POD_PENDING_REQUEST_DOC

Field	Details
Summary	Pod Resource Congestion status of <code>{{labels.microservice}}</code> service is DoC for PendingRequest type
Severity	Major
Expression	<code>ocbsf_pod_resource_congestion_state{type="queue"} == 1</code>
OID	1.3.6.1.4.1.323.5.3.37.1.2.27
Metric Used	<code>ocbsf_pod_resource_congestion_state{type="queue"}</code>
Recommended Actions	Cause: The pending request queue is in the DANGER_OF_CONGESTION state. Diagnostic Information: Confirm that <code>pod_resource_congestion_state{resourceType="queue"} == 1</code> and review queue-related <code>pod_resource_stress</code>. 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 5-49 POD_PENDING_REQUEST_CONGESTED

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

Table 5-49 (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 5-50 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 5-51 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 5-52 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 5-52 (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 5-53 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 5-53 (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 5-54 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 5-54 (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.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 5-55 (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.

Table 5-56 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 5-56 (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 5-57 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 5-57 (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 5-58 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 5-59 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 5-60 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 5-60 (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 5-61 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 5-62 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 5-63 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 5-63 (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 5-64 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 5-65 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 5-66 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 5-66 (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 5-67 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 5-68 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 5-69 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 5-69 (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 5-69 (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 5-70 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 5-70 (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 5-70 (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 5-71 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 5-71 (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 5-72 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 5-72 (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 5-73 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 5-73 (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(occpn_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>occpn_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 5-74 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 5-74 (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 5-75 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 5-75 (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 5-76 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 5-77 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 5-78 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 5-78 (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 5-79 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 5-80 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 5-81 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 5-81 (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 5-82 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 5-83 EGRESS_GATEWAY_DD_UNREACHABLE_MAJOR

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

Table 5-83 (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 5-84 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 5-85 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 5-86 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 5-86 (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) (occpn_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 5-87 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 5-87 (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 5-88 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 5-89 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 5-89 (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 5-90 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 5-91 POD_PENDING_REQUEST_CONGESTION_L1

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

Table 5-91 (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 5-92 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 5-93 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 5-94 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 5-95 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 5-95 (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 5-96 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 5-97 CERTIFICATE_EXPIRY_MAJOR

Field	Details
Description	Certificate expiry in less than 3 months for {{labels.namespace}}
Summary	Certificate expiry in less than 3 months.
Severity	Major
Expression	security_cert_x509_expiration_seconds - time() <= 7862400
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 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 5-98 CERTIFICATE_EXPIRY_CRITICAL

Field	Details
Description	Certificate expiry in less than a month for {{labels.namespace}}
Summary	Certificate expiry in less than a month.
Severity	Critical
Expression	security_cert_x509_expiration_seconds - time() <= 2592000
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 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 5-99 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 5-100 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 5-100 (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 5-101 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 5-102 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 5-103 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 5-103 (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 5-104 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	<pre>(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</pre>
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 5-105 STA_RX_FAIL_COUNT_EXCEEDS_CRITICAL_THRESHOLD

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

Table 5-105 (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 5-106 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 5-106 (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 5-107 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 5-107 (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 5-108 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{applId=" 16777236", msgType="ASA", responseCode! ~"2.*",direction="in"}[5m])) / sum by (namespace) (rate(ocbsf_diam_response_network_total{applId=" 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 5-109 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 5-110 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 5-110 (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 5-111 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	(increase(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 5-112 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 5-113 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 5-114 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 5-115 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 5-115 (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.

6

BSF KPIs

This section provides information about Key Performance Indicators (KPIs) used for Oracle Communications Cloud Native Core Binding Support Function (BSF).

Note

Sample BSF dashboard for Grafana is delivered to the customer through CNC BSF Custom Templates. The metrics and functions used to achieve KPIs are covered in CNC BSF Custom Templates as well.

KPIs

The following table lists the KPIs used for various operations related to BSF:

Table 6-1 BSF KPIs

KPI Details	Metric used for KPI	Service Operation	Response Code
The BSF Management Service Pod Count KPI measures the number of BSF pods that are available in the system.	count(container_memory_usage_bytes{container='bsf-management-service',namespace=\"\$namespace\"})	Not Applicable	Not Applicable
The Ingress-Gateway Pod Count KPI measures the number of Ingress Gateway pods that are available in the system.	count(container_memory_usage_bytes{container='ocbsf-ingress-gateway',namespace=\"\$namespace\"})	Not Applicable	Not Applicable
The Egress-Gateway Pod Count KPI measures the number of Egress Gateway pods that are available in the system.	count(container_memory_usage_bytes{container='ocbsf-egress-gateway',namespace=\"\$namespace\"})	Not Applicable	Not Applicable
The Total-TPS KPI measures the rate of Ingress Gateway requests at BSF.	sum(rate(http_server_requests_seconds_count{kubernetes_namespace=\"\$namespace\",microservice=\"\$ocbsf-ocbsf-ingress-gateway\"}[5m]))	All	Not Applicable
This Memory-Usage KPI measures the current memory usage in bytes.	sum(container_memory_usage_bytes{image!=\"\",namespace=\"\$namespace\"})	Not Applicable	Not Applicable
This Memory-Usage KPI measures the memory usage (in bytes) for the top 16 memory users by each container.	topk(16,sum(container_memory_usage_bytes{image!=\"\",namespace=\"\$namespace\"}) by (container_name))	Not Applicable	Not Applicable
This CPU-Usage KPI measures the number of cores that are being used by each container.	sum(rate(container_cpu_usage_seconds_total{image!=\"\",namespace=\"\$namespace\",container_name!=\"POD\"}[2m])) by (container_name)	Not Applicable	Not Applicable

Table 6-1 (Cont.) BSF KPIs

KPI Details	Metric used for KPI	Service Operation	Response Code
This Binding-Registration KPI measures the rate of successful binding registration requests at BSF.	sum (rate(http_server_requests_seconds_count{kubernetes_namespace=\"\$namespace\",microservice=\"bsf-management-service\",method=\"POST\",status=\"201\"}[2m]))	NF Register	201
This Binding-Registration KPI measures the rate of unsuccessful binding registration requests at BSF.	sum (rate(http_server_requests_seconds_count{kubernetes_namespace=\"\$namespace\",microservice=\"bsf-management-service\",method=\"POST\",status=\"400\"}[2m]))	NF Register	400
This Binding-Deregistration KPI measures the rate of successful deregistration requests at BSF.	sum (rate(http_server_requests_seconds_count{kubernetes_namespace=\"\$namespace\",microservice=\"bsf-management-service\",method=\"DELETE\",status=\"204\"}[2m]))	NF Deregister	204
This Binding-Deregistration KPI measures the rate of unsuccessful deregistration requests at BSF.	sum (rate(http_server_requests_seconds_count{kubernetes_namespace=\"\$namespace\",microservice=\"bsf-management-service\",method=\"DELETE\",status=\"404\"}[2m]))	NF Deregister	404
This Binding-Discovery KPI measures the rate of successful discovery requests at BSF.	sum (rate(http_server_requests_seconds_count{kubernetes_namespace=\"\$namespace\",microservice=\"bsf-management-service\",method=\"GET\",status=\"200\"}[2m]))	NF Discovery	200
This Binding-Discovery KPI measures the rate of discovery requests that have been rejected by BSF due to errors that may have occurred at NF consumer.	sum (rate(http_server_requests_seconds_count{kubernetes_namespace=\"\$namespace\",microservice=\"bsf-management-service\",method=\"GET\",status=\"404\"}[2m]))	NF Discovery	404
This Binding-Discovery KPI measures the rate of discovery requests that have been stated invalid by BSF due to errors that may have occurred at NF consumer.	sum (rate(http_server_requests_seconds_count{kubernetes_namespace=\"\$namespace\",microservice=\"bsf-management-service\",method=\"GET\",status=\"400\"}[2m]))	NF Discovery	400
This Diameter_ingress-request-response KPI measures the rate of AAR request messages sent to network NFs like AF or PCF (in case of BSF).	sum (rate(ocbsf_diam_request_network_total{kubernetes_namespace=\"\$namespace\",msgType=\"AAR\"}[2m]))	All	Not Applicable
This Diameter_ingress-request-response KPI measures the rate of AAA messages going out to network where response code in AAA is 2001, that is, Diameter_Success.	sum (rate(ocbsf_diam_request_network_total{kubernetes_namespace=\"\$namespace\",msgType=\"AAA\",responseCode=\"2001\"}[2m]))	All	201

Table 6-1 (Cont.) BSF KPIs

KPI Details	Metric used for KPI	Service Operation	Response Code
This Diameter_ingress-request-response KPI measures the rate of STR request messages sent to network NFs.	sum (rate(ocbsf_diam_request_network_total{kubernetes_namespace=\"\$namespace\",msgType=\"STR\"}[2m]))	All	Not Applicable
This Diameter_ingress-request-response KPI measures the rate of STA messages going out to network where response code in STA is 2001, that is, Diameter_Success.	sum (rate(ocbsf_diam_request_network_total{kubernetes_namespace=\"\$namespace\",msgType=\"STA\",responseCode=\"2001\"}[2m]))	All	201
This Egress-Request-Response KPI measures the rate of Delete requests sent to external NFs through Egress Gateway.	sum(rate(oc_egressgateway_http_requests_total{kubernetes_namespace=\"\$namespace\",Method='DELETE'}[2m]))	Delete	Not Applicable
This Egress-Request-Response KPI measures the rate of responses, to Delete requests (initiated at BSF), by external NFs through Egress Gateway.	sum(rate(oc_egressgateway_http_responses_total{kubernetes_namespace=\"\$namespace\",Method='DELETE'}[2m]))	Delete	Not Applicable
This Egress-Request-Response KPI measures the rate of GET requests sent to external NFs through Egress Gateway.	sum(rate(oc_egressgateway_http_requests_total{kubernetes_namespace=\"\$namespace\",Method='GET'}[2m]))	GET	Not Applicable
This Egress-Request-Response KPI measures the rate of responses, to GET requests (initiated at BSF), by external NFs through Egress Gateway.	sum(rate(oc_egressgateway_http_responses_total{kubernetes_namespace=\"\$namespace\",Method='GET'}[2m]))	GET	Not Applicable
This Egress-Request-Response KPI measures the rate of POST requests sent to external NFs through Egress Gateway.	sum(rate(oc_egressgateway_http_requests_total{kubernetes_namespace=\"\$namespace\",Method='POST'}[2m]))	POST	Not Applicable
This Egress-Request-Response KPI measures the rate of responses, to POST requests (initiated at BSF), by external NFs through Egress Gateway.	sum(rate(oc_egressgateway_http_responses_total{kubernetes_namespace=\"\$namespace\",Method='POST'}[2m]))	POST	Not Applicable
This Diameter-egress-Request-Response KPI measures the rate of Update Notify requests sent to external NFs through Egress Gateway.	sum (rate(ocpm_egress_request_total{kubernetes_namespace=\"\$namespace\",operation_type=\"update_notify\",servicename_3gpp=\"\$rx\"}[2m]))	Update Notify	Not Applicable

Table 6-1 (Cont.) BSF KPIs

KPI Details	Metric used for KPI	Service Operation	Response Code
This Diameter-egress-Request-Response KPI measures the rate of Update Notify requests sent to external NFs through Egress Gateway.	sum (rate(ocpm_egress_response_total{kubernetes_namespace=\"\$namespace\",operation_type=\"update_notify\",servicename_3gpp=\"rx\",response_code=\"2xxx\"}[2m])) Note Different microservices may have different nomenclature. For example, operation_type or operationType.	Update Notify	2xxx
This Diameter Connections KPI measures the total number of connections with network peer.	sum (rate(occpn_diam_conn_network{kubernetes_namespace=\"\$namespace\"}[2m]))	Not Applicable	Not Applicable
This Diameter Connections KPI measures the number of connections with network peer for a given application ID.	sum (rate(occpn_diam_conn_app_network{kubernetes_namespace=\"\$namespace\"}[2m]))	Not Applicable	Not Applicable

BSF Metrics

BSF Metrics

This chapter includes information about Metrics for Oracle Communications Cloud Native Core, Binding Support Function (BSF).

BSF uses the observability tool Grafana to analyze and visualize data. The Grafana Dashboard consists of panels displaying the data as graphs, charts or other visualizations. A dashboard in Grafana is represented by a JSON object, which stores the metadata of its dashboard. Following Grafana dashboards are created for BSF Observability and debugging purposes:

- **Observability Dashboard:**
 - This dashboard metadata is stored in the `BSF_Observability_Dashboard.json` file.
 - This dashboard has the panels that monitors the overall health status of the setup & contains details like resource utilization, kmps, latency, etc.
 - Customers can use this Dashboard on a regular basis for observing the status of the setup and also for assessing the setup performance with a run.
 - Depending upon the applicability of a panel or row, the customers have the flexibility to either remove, update or add content to the Dashboard.
- **Debug Dashboard:**
 - This dashboard metadata is stored in the `BSF_Debug_Dashboard.json` file.
 - This dashboard has the panels that monitor the operational status of the setup and is used for debugging & troubleshooting purposes.
 - Customers can use this Dashboard when any issue occur in the NF and also when reporting a probable issue to Oracle Customer Support from the setup.
 - Depending upon the applicability of a panel or row, the customers have the flexibility to either remove, update or add content to the Dashboard.

The name of the metrics may contain suffix such as total, seconds, max and so on. It gets added by the Micrometer registry if it is not present in the metrics name. The metric name has the following format for suffix:

The metric name is equal to `<Baseline of the metric>_<Suffix>`

Table 7-1 Metrics type and Suffix

Metric Type	Suffix
Counter	_total
Gauge	N/A
TimerGauge	_seconds
MultiGauge	N/A

Table 7-1 (Cont.) Metrics type and Suffix

Metric Type	Suffix
Timer	_seconds_max or _seconds Note: There are two types of suffix used for timer metrics. For example, my_timer_seconds_max gauge and my_timer_seconds summary. In summary type, there will be further addition of suffix such as _count or _sum.
DistributionSummary	N/A or _max Note: There are two types of suffix used for DistributionSummary. For example, my_distribution_ratio histogram and my_distribution_ratio_max gauge. In the histogram type there will be further addition of suffix such as bucket, _count, or sum.
LongTaskTimer	_seconds_max or _seconds Note: There are two types of suffix used for LongTaskTimer. For example, long_task_timer_seconds_max gauge and long_task_timer_seconds summary. In summary type there will be further addition of suffix such as _active_count or _duration_sum.

Table 7-2 Dimension Description

Dimension	Description
operation_type	Type of operation Values: • create • get • put • update • terminate • update_notify • terminate_notify • subscribe • unsubscribe • transfer • resubscribe • retry_reason • binding_id • binding_level
dnn	Data Network Name or Access Point Name
snssai	Single Network Slice Selection Assistance Information

Table 7-2 (Cont.) Dimension Description

Dimension	Description
response_code	Response code HTTP interfaces: • 1xx • 2xx • 3xx • 4xx • 5xx Diameter interfaces: • 2xxx • 3xxx • 4xxx • 5xxx
latency	The total time in between request and response. If latency between request and response is 203, then bucket number is 4. Max bucket set to 10 (0-9), Range 50ms.
nf_instance_id	Unique id of the nf Instance. • ingress: source nfInstanceId • egress: destination nfInstanceId HTTP interfaces: Diameter interfaces: • ingress: Origin-Host AVP • egress: Destination-Host AVP
nf_name	This represents the FQDN corresponding to the NF InstanceID present in the nf_instance_id dimension. HTTP interface: egress: NF FQDN
sbi_priority	Service Based Interface
service_version	Service version Value: [UDR = "v1,v2", CHF = "v1"]
service	The complete name of current service. Value: string
namespace	The namespace of current service. Value: string
category	The category of current service. Value: • database • common • infra • pcf • bsf
destHost	Value of destination Host received or sent in the corresponding request message
destRealm	Value of destination Realm received or sent in the corresponding request message
origHost	Value of origination Host received or sent in the corresponding request message
origRealm	Value of origination Realm received or sent in the corresponding request message
reqDestHost	Value of destination Host in corresponding request message of response message.
reqDestRealm	Value of destination Realm in corresponding request message of response message.
reqOrigHost	Value of origination Host in corresponding request message of response message.

Table 7-2 (Cont.) Dimension Description

Dimension	Description
reqOrigRealm	Value of origination Realm in corresponding request message of response message.
direction	Indicates direction of message flow. • "In" means coming towards POD/micro-service • "Out" means going out from POD/micro-service
applId	Application ID exchanged in CEX messages or used in the respective message of an application.
applicationName	Human readable name of corresponding application ID 16777236 => Rx 16777238 => Gx 16777302 => Sy 16777217 => Sh Note: Sh interface is not supported for Converged Policy mode of deployment. 0xffffffff => Relay
cmdCode	Command code value in the received or sent, request or answer message
msgType	Type of the message, for example CCRT-T, CCR-I etc.
responseCode	result code in diameter message
spendingLimitDataSource	Specifies the source from which PCF fetches policy counters. Value: OCS
retry	Identify message is a retry message. Value: • true • false
retryAnswer	Reason for the retry message. Value: • error code • timeout
level	Indicates the current load level or the level of pod congestion. Value: • 0 = Normal • 1 = DOC • 2 = Congested
type	Resource type Value: • PendingRequest • CPU • Memory
le	le is abbreviated as "Less than equal to". Value of a defined bucket for a Histogram.
sessRuleReports	Indicates that session rule report is received at PCF.
policyDecFailureReports	Indicates that Policy decision failure report is received at PCF.
isLeaderPod	Indicates if the pod calculating the threshold level is a leader pod.
prevLevel	Indicates the previous load level prior to current load level calculation.
levelChangeType	Indicates the level change type. The value of this dimension can be: • None: when load level is same • Increment: when level changes to higher level
servicenameNon3gpp	-
serviceResource	-

Table 7-2 (Cont.) Dimension Description

Dimension	Description
params	Lists the API parameters.
outcome	Shows the outcome of an operation such as SUCCESS, FAILURE, TIMEOUT.
cause	Contains the error cause.
peerHost	Indicates the value of peer host received or sent in corresponding connection request message.
perRealm	Indicates the value of peer realm received or sent in corresponding connection request message.
alt_key	Defines which Alt Lookup Table was used. Possible Values: • ipv4_addr_lookup • ipv6_prefix_lookup • mac_addr_48_lookup
db_operation	Defines the DB Operation that was performed. Possible Values: • insert • delete_by_binding_id_and_ue_addr • delete_by_binding_id • update_audit_time_by_binding_id_and_ue_addr • select_by_ue_addr • select_by_binding_id_and_ue_addr
status	Defines if the operation was successful or failed. A DB operation is considered successful if at least one record was inserted/updated/selected/deleted; otherwise if no records were impacted, or if a DB exception occurred, it is considered a failure. Possible Values: • successful • failure
exception	Defines the exception that occurred while performing the DB Operation. Possible Values: • None, if no exception occurred • The name of the DB Exception (e.g. PessimisticLockException)

7.1 Egress Gateway Metrics for SCP

This section provides details about SCP health monitoring metrics and the respective dimensions.

Table 7-3 ocbsf_oc_egressgateway_peer_health_status

Field	Details
Description	It defines Egress Gateway peer health status. • This metric is set to 1, if a peer is unhealthy. • This metric is reset to 0, when it becomes healthy again. • This metric is set to -1, if peer is removed from peerconfiguration.
Type	Gauge
Dimensions	• peer • vfqdn

Table 7-3 (Cont.) ocbsf_oc_egressgateway_peer_health_status

Field	Details
Examples	- ocbsf_oc_egressgateway_peer_health_status{"peer":"10.75.213.172:8080"} 1.0 - ocbsf_oc_egressgateway_peer_health_status{"peer":"10.75.213.172:8080"} 0.0 - ocbsf_oc_egressgateway_peer_health_status{"vfqdn":"http://abc.com","peer":"10.75.213.172:8080"} 1.0 - ocbsf_oc_egressgateway_peer_health_status{"vfqdn":"http://abc.com","peer":"10.75.213.172:8080"} 0.0

Table 7-4 ocbsf_oc_egressgateway_peer_health_ping_request_total

Field	Details
Description	This metric is incremented every time a health ping is sent toward a peer.
Type	Counter
Dimensions	- peer - vfqdn
Examples	- ocbsf_oc_egressgateway_peer_health_ping_request_total{"peer":"10.75.213.172:8080"} 389.0 - ocbsf_oc_egressgateway_peer_health_ping_request_total{"peer":"10.75.213.172:8080"} 439.0 - ocbsf_oc_egressgateway_peer_health_ping_request_total{"vfqdn":"http://abc.com","peer":"10.75.213.172:8080"} 389.0 - ocbsf_oc_egressgateway_peer_health_ping_request_total{"vfqdn":"http://abc.com","peer":"10.75.213.172:8080"} 439.0

Table 7-5 ocbsf_oc_egressgateway_peer_health_ping_response_total

Field	Details
Description	This metric is incremented every time a health ping response is received from a peer irrespective of success or failure.
Type	Counter
Dimensions	- peer - vfqdn - statusCode - cause
Examples	- ocbsf_oc_egressgateway_peer_health_ping_response_total{"peer":"10.75.213.172:8080","status":"httpstatus","cause":""} 89.0 - ocbsf_oc_egressgateway_peer_health_ping_response_total{"peer":"10.75.213.172:8080","status":"Exception","cause":"exception cause"} 39.0 - ocbsf_oc_egressgateway_peer_health_ping_response_total{"vfqdn":"http://abc.com","status":"httpstatus","cause":""} 89.0 - ocbsf_oc_egressgateway_peer_health_ping_response_total{"vfqdn":"http://abc.com","status":"Exception","cause":"exception cause"} 39.0

Table 7-6 ocbsf_oc_egressgateway_peer_health_status_transitions_total

Field	Details
Description	This metric is incremented every time a peer is transitioned from Available to Unavailable or from Unavailable to Available.
Type	Counter
Dimensions	- peer - vfqdn - from - to
Examples	- ocbsf_oc_egressgateway_peer_health_status_transitions_total{"identifier":"10.75.213.172:8080","from":"available","to":"unavailable"} 14.0 - ocbsf_oc_egressgateway_peer_health_status_transitions_total{"identifier":"10.75.213.172:8080","from":"unavailable","to":"available"} 34.0 - ocbsf_oc_egressgateway_peer_health_status_transitions_total{"vfqdn":"http://abc.com","peer":"10.75.213.172:8080","from":"unavailable","to":"available"} 34.0 - ocbsf_oc_egressgateway_peer_health_status_transitions_total{"vfqdn":"http://abc.com","peer":"10.75.213.172:8080","from":"available","to":"unavailable"} 14.0

Table 7-7 ocbsf_oc_egressgateway_peer_count

Field	Details
Description	This metric is incremented every time for the peer count.
Type	Gauge
Dimensions	peerset
Example	ocbsf_oc_egressgateway_peer_count{"peerset":"set-0"} 3.0

Table 7-8 ocbsf_oc_egressgateway_peer_available_count

Field	Details
Description	This metric is incremented every time for the available peer count.
Type	Gauge
Dimensions	peerset
Example	ocbsf_oc_egressgateway_peer_available_count{"peerset":"set-0"} 4.0

Table 7-9 ocbsf_oc_egressgateway_resolve_fqdn_from_ars_failure_total

Field	Details
Description	This metric is incremented each time the Egress Gateway fails to resolve an SCP peer's FQDN through DNS SRV lookup from ARS.
Type	Counter
Dimensions	- microservice - namespace - vfqdn

Table 7-9 (Cont.) ocbsf_oc_egressgateway_resolve_fqdn_from_ars_failure_total

Field	Details
Example	ocbsf_oc_egressgateway_resolve_fqdn_from_ars_failure_total{ namespace="ocbsf", microservice="bsf-binding-service", vfqdn="scp.5gc.mnc001.mcc001.3gppnetwork.org", pod="bsf-binding-service-7d6f8c9b5d-k2xqv" } 3

7.2 Correlation-Info Header Metrics

For every correlation-info header received or newly generated, a metric will be pegged. Following are the list of metrics:

Below mention table list the metrics that shall be implemented as part of this feature with the following Dimensions:

1. operation_type=
{"create","update","delete","subscribe","unsubscribe","terminate","register","deregister"....}
2. correlation_info_type={"imsi", "msisdn", "imsi,msisdn"}

Table 7-10 ocbsf_correlation_info_header_received

Field	Details
Description	BSF reports total outgoing request that are carried by the correlation-info header.
Type	Counter
Dimensions	- operation_type - correlation_info_type
Example	ocbsf_correlation_info_header_received_total '{correlation_info_type="imsi", operation_type="create"}',2

Table 7-11 ocbsf_correlation_info_header_forwarded

Field	Details
Description	BSF reports total incoming request that are carried by the correlation-info header.
Type	Counter
Dimensions	- operation_type - correlation_info_type
Example	ocbsf_correlation_info_header_received_total '{correlation_info_type="imsi", operation_type="update"}',1

Table 7-12 ocbsf_correlation_info_header_generated

Field	Details
Description	BSF reports the total responses carried by the correlation-info header.
Type	Counter
Dimensions	- operation_type - correlation_info_type

Table 7-12 (Cont.) ocbsf_correlation_info_header_generated

Field	Details
Examples	- ocbsf_correlation_info_header_forwarded_total '{correlation_info_type="imsi", operation_type="subscribe"}',2 - ocbsf_correlation_info_header_generated_total '{correlation_info_type="imsi", operation_type="create"}',2

7.3 Configuration Server Metrics

Table 7-13 topic_version

Field	Details
Description	Configuration service will have this metrics a database value from each topic version. The Services fetching the configurations from Configuration Server, will have its current topic version till which configurations has been fetched successfully.
Type	Gauge
Dimensions	- Service Name - Pod Name
Examples	- topic_version{topicName="common.congestionthreshold.diam-gateway"},} 1.0 - topic_version{topicName="common.auditservice.cfg"},} 1.0 - topic_version{topicName="bsf.sbi.errorcodes"},} 36.0 - topic_version{topicName="common.diamsetting"},} 2.0 - topic_version{topicName="common.public.diammessagepriority"},} 2.0 - topic_version{topicName="bsf.managementservice"},} 1.0 - topic_version{topicName="subscriber.activity.logging.idlist.bsf"},} 1.0 - topic_version{topicName="public.policy.error.handler.config"},} 1.0 - topic_version{topicName="common.public.diampeernode"},} 5.0 - topic_version{topicName="nrfclient.cfg"},} 1.0 - topic_version{topicName="common.public.diamloadshedding"},} 2.0 - topic_version{topicName="diameter.errorcodes"},} 7.0 - topic_version{topicName="common.public.diamroutingtable"},} 1.0 - topic_version{topicName="subscriber.activity.logging.mapping.bsf"},} 1.0 - topic_version{topicName="bsf.global.cfg"},} 1.0 - topic_version{topicName="common.public.diamcontrolledshutdown"},} 1.0 - topic_version{topicName="config.level"},} 53.0 - topic_version{topicName="diameter.customerrorcodes"},} 1.0

7.4 Active Sessions Count Metrics

Table 7-14 oc_db_active_session_count

Field	Details
Description	Active sessions count
Type	Gauge
Dimensions	- Service - Table

Table 7-15 inbound_requests_total

Field	Details
Description	Inbound requests total (number of requests received via CM/REST to get instantaneous Active Session Count)
Type	Counter
Dimensions	- requestType (onDemandActiveSessionCount, register, deregister) - tableName

7.5 AppInfo Metrics

Table 7-16 appinfo_service_running

Field	Details
Description	Provides the status of monitored services.
Type	Gauge
Dimensions	- service - namespace - category
Example	appinfo_service_running{service="xxx",namespace="xxx",category="xxx"} 1

Table 7-17 appinfo_category_running

Field	Details
Description	Provides the status of monitored categories.
Type	Gauge
Dimensions	- namespace - category
Example	- appinfo_category_running{category="xxx",namespace="xxx"} 1 - appinfo_category_good{category="xxx",namespace="xxx"} 1

Table 7-18 appinfo_category_good

Field	Details
Description	Provides the readiness of monitored categories.
Type	Gauge
Dimensions	- namespace - category

Table 7-19 nfscore

Field	Details
Description	<i>factor</i> - Contains one of the values: all, successTPS, serviceHealth, signallingConnections, replicationHealth, localityPreference. - When the factor is set to <i>all</i> that means NF score is calculated for all the factors. <i>status</i> <i>success</i>: when the factor is enabled and its value is fetched successfully. <i>failed</i>: when the factor is enabled and fetching the value fails. <i>notCalculated</i>: when the factor is disabled.
Type	Gauge
Dimensions	- nflInstanceId - factor - Status
Example	nfscore{app="testing-appinfo", app_kubernetes_io_instance="testing", app_kubernetes_io_managed_by="Helm", app_kubernetes_io_name="appinfo", app_kubernetes_io_part_of="ocbsf", app_kubernetes_io_version="1.7.1.0.0", application="ocbsf", calculatedStatus="success", container="appinfo", endpoint="cnc-metrics", engVersion="22.4.0-rc.5", factor="localityPreference", helm_sh_chart="appinfo-22.4.0-rc.5", instance="10.233.117.146:9000", job="occne-infra/occne-nf-cnc-podmonitor", microservice="bsf-app-info", mktgVersion="1.7.1.0.0", namespace="biloxi-ns", nflInstanceId="fe7d992b-0541-4c7d-ab84-c6d70b1b0666", pod="testing-appinfo-78dc65865f-hgrhk", pod_template_hash="78dc65865f", vendor="Oracle"} 5

Table 7-20 nfScoringFactorActualValue

Field	Details
Description	factor tag would contain one of the following values: successTPS, serviceHealth, signallingConnections, replicationHealth, localityPreference
Type	Gauge
Dimensions	- nflInstanceId - factor
Example	nfScoringFactorActualValue{app="testing-appinfo", app_kubernetes_io_instance="testing", app_kubernetes_io_managed_by="Helm", app_kubernetes_io_name="appinfo", app_kubernetes_io_part_of="ocbsf", app_kubernetes_io_version="1.7.1.0.0", application="ocbsf", calculatedStatus="success", container="appinfo", endpoint="cnc-metrics", engVersion="22.4.0-rc.5", factor="localityPreference", helm_sh_chart="appinfo-22.4.0-rc.5", instance="10.233.117.146:9000", job="occne-infra/occne-nf-cnc-podmonitor", microservice="bsf-app-info", mktgVersion="1.7.1.0.0", namespace="biloxi-ns", nflInstanceId="fe7d992b-0541-4c7d-ab84-c6d70b1b0666", pod="testing-appinfo-78dc65865f-hgrhk", pod_template_hash="78dc65865f", vendor="Oracle"} 5

7.6 Audit Service Metrics

The following table describes the Audit Service metrics and respective dimensions:

Table 7-21 audit_recs_stale_total

Field	Details
Description	Number of records detected as stale.
Type	Counter
Dimensions	• ServiceName • TableName
Example	audit_recs_stale_total{ServiceName="bsf-management-service",TableName="pcf_binding",} 1.0

Table 7-22 audit_recs_notif_total

Field	Details
Description	Number of stale record notifications sent, applicable for modes: NOTIFY and DELETE_NOTIFY.
Type	Counter
Dimensions	ServiceName
Example	audit_recs_notif_total{ServiceName="bsf-management-service",} 1.0

Table 7-23 audit_recs_deque_for_notif_total

Field	Details
Description	Number of stale records dequeued to send Notification.
Type	Counter
Dimensions	ServiceName
Example	audit_recs_deque_for_notif_total{ServiceName="bsf-management-service",} 1.0

Table 7-24 audit_recs_enqueue_for_notif_total

Field	Details
Description	Number of stale records enqueued from Database.
Type	Counter
Dimensions	ServiceName
Example	audit_recs_enqueue_for_notif_total{ServiceName="bsf-management-service",} 1.0

Table 7-25 oc_db_active_session_count

Field	Details
Description	Reports the session for a given service.
Type	Gauge
Dimensions	• ServiceName • TableName

Table 7-25 (Cont.) oc_db_active_session_count

Field	Details
Example	oc_db_active_session_count{Service="bsf-management-service",Table="pcfbinding",} 1.0

Table 7-26 topic_version

Field	Details
Description	Describes the current applied version of a given topic (mentioned in dimension topic_name) into the pod.
Type	Gauge
Dimensions	- Service Name - Pod Name

Table 7-27 audit_schedule_current_status

Field	Details
Description	This metric defines current status of audit schedule. This metric is set to 1, if the status of the current audit schedule is "CTRL_SHUTDOWN_INTRPT". This metric is reset to 0, if the status is changed.
Type	Gauge
Dimension	- Service - Table - Status - CTRL_SHUTDOWN_INTRPT - INTERRUPTED - RUNNING - COMPLETED - INITIALISING - DEREGISTERED - QUEUED
Example	audit_schedule_current_status{Service="bsf-management-service",Status="CTRL_SHUTDOWN_INTRPT",Table="pcf_binding",application="pcf_session-state-audit",hostname="bsf-ocpm-audit-service-6d65545bbb-mb77m"} 0.0

Table 7-28 occnp_audit_schedule_status_total

Field	Details
Description	This metric indicates the number of times the audit service has transitioned its status for a specific service and table.
Type	Counter

Table 7-28 (Cont.) occnp_audit_schedule_status_total

Field	Details
Dimension	- Service - Table - Status - CTRL_SHUTDOWN_INTRPT - INTERRUPTED - RUNNING - COMPLETED - INITIALISING - DEREGISTERED - QUEUED
Example	occnp_audit_schedule_status_total{Service="bsf-management-service",Status="CTRL_SHUTDOWN_INTRPT",Table="pcf_binding",application="pcf_session-state-audit",hostname="bsf-ocpm-audit-service-6d65545bbb-mb77m"} 1.0

7.7 BSF Management Service

This section describes the metrics used for BSF Management service.

The following table provides information about metrics related to BSF Management service.

Table 7-29 ocbsf_ingress_response_total

Field	Details
Description	This metric is pegged when BSF sends a response message through Ingress Gateway.
Type	Gauge
Dimensions	operation_type dnn snssai application_error_code response_code binding_level binding_id
Example	ocbsf_ingress_response_total{application="bsf-management-service",application_error_code="",binding_id="setxyz.bsfsset.5gc.mnc015.mcc360,setxyz.bsfsset.5gc.mnc015.mcc350",binding_level="nf-set",dnn="dnn1",operation_type="register",response_code="2xx",snssai="100-D143A5",}

Table 7-30 ocbsf_ingress_request_total

Field	Details
Description	This metric is pegged when BSF receives a request message through Ingress Gateway.

Table 7-30 (Cont.) ocbsf_ingress_request_total

Field	Details
Type	Gauge
Dimensions	operation_type dnn snssai pcf_set_id pcf_id
Example	ocbsf_ingress_request_total{application="bsf-management-service",dnn="dnn1",operation_type="register",pcf_id="somePcfId",pcf_set_id="somePcfSetId",snssai="100-D143A5",} 1.0

Table 7-31 ocbsf_audit_notif_request_count_total

Field	Details
Description	This metric is pegged when BSF receives a notification request from the audit service for stale bindings.
Type	Counter
Dimensions	application
Example	ocbsf_audit_notif_request_count_total{application="bsf-management-service",operation_type="audit_notify",} 1.0

Table 7-32 ocbsf_audit_notif_response_count_total

Field	Details
Description	This metric is pegged when BSF sends a notification response for the request received from the Audit service for stale bindings.
Type	Counter
Dimensions	application
Example	ocbsf_audit_notif_response_count_total{application="bsf-management-service",operation_type="audit_notify",response_code="2xx",} 1.0

Table 7-33 ocbsf_audit_delete_records_count_total

Field	Details
Description	This metric is pegged when BSF successfully deletes stale bindings.
Type	Counter
Dimensions	application
Example	ocbsf_audit_delete_records_count_total{application="bsf-management-service",operation_type="audit_notify",} 1.0

Table 7-34 ocbsf_audit_delete_records_max_ttl_count_total

Field	Details
Description	This metric is pegged when the BSF successfully deletes bindings on receiving maxTTL as 'true'
Type	Counter
Dimensions	application
Example	ocbsf_audit_delete_records_count_total{application="bsf-management-service",operation_type="audit_notify",} 1.0

Table 7-35 ocbsf_diamgw_notification_request_count_total

Field	Details
Description	This metric is pegged when BSF management service receives stale binding notification request from Diameter Gateway.
Type	Counter
Dimensions	application
Example	ocbsf_diamgw_notification_request_count_total{application="bsf-management-service",operation_type="audit_notify",} 1.0

Table 7-36 ocbsf_diamgw_notification_response_count_total

Field	Details
Description	This metric is pegged when BSF management service sends a response for the stale binding notification request from Diameter Gateway.
Type	Counter
Dimensions	application
Example	ocbsf_diamgw_notification_response_count_total{application="bsf-management-service",operation_type="audit_notify",} 1.0

Table 7-37 ocbsf_diamgw_delete_records_count_total

Field	Details
Description	This metric is pegged when BSF management service successfully deletes stale PCF bindings due to Diameter Gateway notification requests.
Type	Counter
Dimensions	application
Example	ocbsf_diamgw_delete_records_count_total{application="bsf-management-service",operation_type="audit_notify",} 1.0

Table 7-38 ocbsf_query_request_count_total

Field	Details
Description	This metric is pegged when BSF management service sends request to PCF to confirm if PcfBinding is stale or not.
Type	Counter

Table 7-38 (Cont.) ocbsf_query_request_count_total

Field	Details
Dimensions	• application • operation_type • nfInstanceId • nfName • binding_Id • binding_level • retry • retry_reason
Example	ocbsf_query_request_count_total{application="bsf-management-service",binding_id="",binding_level="",nf_instance_id="",nf_name="nf1stub.s-galileo.svc",operation_type="query",retry="false",retry_reason=""} 1.0 ocbsf_query_request_count_total{application="bsf-management-service",binding_id="set001.pcfset.5gc.mnc012.mcc345",binding_level="NF_SET(1)",nf_instance_id="fe7d992b-0541-4c7d-ab84-c6d70b1b0124",nf_name="nf11stub.s-galileo.svc",operation_type="query",retry="true",retry_reason="alternate_route"} 1.0

Table 7-39 ocbsf_query_response_count_total

Field	Details
Description	This metric is pegged when BSF management service receives a response for PcfBinding stale confirmation request.
Type	Counter
Dimensions	• application • operation_type • response_code • nfInstanceId • nfName • binding_Id • binding_level • retry • retry_reason
Example	ocbsf_query_response_count_total{application="bsf-management-service",binding_id="",binding_level="",nf_instance_id="",nf_name="nf1stub.s-galileo.svc",operation_type="query",response_code="2xx",retry="false",retry_reason=""} 1.0 ocbsf_query_response_count_total{application="bsf-management-service",binding_id="set001.pcfset.5gc.mnc012.mcc345",binding_level="NF_SET(1)",nf_instance_id="",nf_name="nf1stub.s-galileo.svc",operation_type="query",response_code="404",retry="true",retry_reason="alternate_route"} 1.0 ocbsf_query_response_count_total{application="bsf-management-service",binding_id="set001.pcfset.5gc.mnc012.mcc345",binding_level="NF_SET(1)",nf_instance_id="",nf_name="nf1stub.s-galileo.svc",operation_type="query",response_code="5xx",retry="false",retry_reason=""} 1.0

Table 7-40 ocbsf_http_out_conn_request_total

Field	Details
Description	This metric is pegged when BSF sends notification request using HTTP in Diameter Gateway.
Type	Counter
Dimensions	operation_type (delete) serviceResource servicename3gpp (DiameterGateway)

Table 7-41 ocbsf_http_out_conn_response_total

Field	Details
Description	This metric is pegged when BSF receives response for notification request requests using HTTP in Diameter Gateway.
Type	Counter
Dimensions	operation_type (delete) responseCode serviceResource servicename3gpp (DiameterGateway)

Table 7-42 ocbsf_binding_query_request_total

Field	Details
Description	This metric is pegged when BSF receives a query request for PCF bindings. PCF Binding query-service sends SUPI, GPSI, IPv4 address, IPv6 address, and mac address to BSF.
Type	Counter
Dimensions	operation_type (query) criteria_type (SUPI/GPSI/UE Address)
Example	ocbsf_binding_query_request_total{application="bsf-managementservice",gpsi="false",ipv4Addr="true",ipv6Prefix="false",macAddr48="false",operation_type="query",supi="true",} 1.0

Table 7-43 ocbsf_binding_query_response_total

Field	Details
Description	This metric is pegged when BSF sends a response message for a PCF bindings query request. PCF Binding query-service sends SUPI, GPSI, IPv4 address, IPv6 address, and mac address to BSF.
Type	Counter

Table 7-43 (Cont.) ocbsf_binding_query_response_total

Field	Details
Dimensions	response_code (2xx, 4xx) operation_type (query) criteria_type (SUPI/GPSI/UE Address)
Example	ocbsf_binding_query_response_total{application="bsf-managementservice",application_error_code="",gpsi="false",ipv4Addr="true",ipv6P refix="false",macAddr48="false",operation_type="query",response_code="2xx",sup i="true",} 1.0

Table 7-44 ocbsf_binding_delete_request_total

Field	Details
Description	This metric is pegged when BSF receives a request to delete PCF bindings. PCF binding sends the binding ids to BSF.
Type	Counter
Dimensions	application operation_type (DELETE)
Example	ocbsf_binding_delete_request_total{application="bsf-management- service",operation_type="manual_delete",pcf_binding_delete_count="3",}1.0

Table 7-45 ocbsf_binding_delete_response_total

Field	Details
Description	This metric is pegged when BSF sends a response message to the delete PCF bindings request. PCF binding sends the binding ids to BSF.
Type	Counter
Dimensions	application response_code (2xx,5xx) operation_type (DELETE)
Example	ocbsf_binding_delete_response_total{application="bsf-management- service",application_error_code="",operation_type="manual_delete",pcf_binding_d elete_count="3",response_code="2xx",}1.0

Table 7-46 topic_version

Field	Details
Description	Describes the current applied version of a given topic (mentioned in dimension topic_name) into the pod.
Type	Gauge
Dimensions	- Service Name - Pod Name
Example	-

Table 7-47 http_server_requests_seconds_count

Field	Details
Description	BSF Management service overall processing time.
Type	Counter
Dimensions	• application • exception • method • outcome • status • uri
Example	http_server_requests_seconds_count{application="bsf-management-service",exception="None",method="GET",outcome="SUCCESS",status="200",uri="/v3/api-docs/swagger-config",} 1.0

Table 7-48 http_server_requests_seconds_sum

Field	Details
Description	BSF Management service overall processing time.
Type	Gauge
Dimensions	• application • exception • method • outcome • status • uri
Example	http_server_requests_seconds_sum{application="bsf-management-service",exception="None",method="GET",outcome="SUCCESS",status="200",uri="/v3/api-docs/swagger-config",} 0.008189215

Table 7-49 error_handler_exec_total

Field	Details
Description	This metric is incremented on error handling performed by framework.
Type	Counter
Dimensions	• application • applicationException • error_type • operation • origin • rule_name • source_interface • target_interface • wrapped_exception
Example	ocbsf_error_handler_exec_total{application="bsf-management-service",application_exception="JavaException",error_type="INTERNAL",operation="CREATE",origin="JAVA",rule_name="REJECT_WITH_ENHANCED_DETAIL",source_interface="POLICY",status="400",target_interface="BSF",wrapped_exception="ServiceException",}

Table 7-50 error_handler_in_total

Field	Details
Description	This metric is incremented on initialization of error handling.
Type	Counter
Dimensions	- application - applicationException - wrapped_exception
Example	ocbsf_error_handler_in_total{application="bsf-management-service",application_exception="JavaException",status="500",wrapped_exception="ServiceException",}

Table 7-51 error_handler_out_total

Field	Details
Description	This metric is incremented on completion of error handling.
Type	Counter
Dimensions	- application - applicationException - error_resolved - wrapped_exception
Example	ocbsf_error_handler_out_total{application="bsf-management-service",application_exception="JavaException",error_resolved="true",status="400",wrapped_exception="ServiceException",}

Table 7-52 ocbsf_binding_revalidation_request_total

Field	Details
Description	This metric is pegged everytime BSF receives a request for binding revalidation.
Type	Counter
Dimensions	- DNN - snssai - PCF FQDN
Example	-

Table 7-53 ocbsf_binding_revalidation_response_total

Field	Details
Description	This metric is pegged everytime BSF responds to binding revalidation request.
Type	Counter
Dimensions	- DNN - snssai - PCF FQDN - responseCode
Example	ocbsf_binding_revalidation_response_total{application="bsf-management-service",dnn="Internet",pcfFqdn="oc-diam-gateway",response_code="2xx",snssai="1-000007",} 44.0

Table 7-54 ocbsf_binding_revalidation_pcfBinding_missing_total

Field	Details
Description	This metric is pegged when PCF Binding is detected as missing in BSF while responding to binding revalidation request.
Type	Counter
Dimensions	• DNN • snssai • PCF FQDN • responseCode
Example	ocbsf_binding_revalidation_pcfBinding_missing_total{application="bsf-management-service",dnn="Internet",pcfFqdn="oc-diam-gateway",response_code="2xx",snssai="1-000007",} 1.0

Table 7-55 ocbsf_http_congestion_message_reject_total

Field	Details
Description	Tracks the rejected messages due to congestion.
Type	Counter
Dimensions	• message priority (priority) • response code (response_code) • operation (requestMethod) • rest endpoint (requestUri) • Congestion Level (congestionState)

Table 7-56 ocbsf_pcf_cookie_total

Field	Details
Description	This metric is incremented when: • the value of direction dimension is received, BSF Management Service receives ocnf-service-cookie cookie from SM service during SM Create, and the ENABLE_PCF_COOKIE Advanced Settings key for BSF Management Service is set to true. • the value of direction dimension is sent, Diameter Gateway queries BSF Management Service during AAR-I request and ENABLE_PCF_COOKIE Advanced Settings key for BSF Management Service is set to true.
Type	Counter
Dimensions	• application • direction Possible values: — received — sent
Examples	ocbsf_pcf_cookie_total{application="ocbsf",direction="sent",} 1 ocbsf_pcf_cookie_total{application="ocbsf",direction="received",} 1

Table 7-57 ocbsf_http_message_priority_request_count_total

Field	Details
Description	This metric is pegged when BSF receives an incoming request or there is an outgoing request from BSF during the following scenarios: • register • deregister • discovery • staleSession
Type	Counter
Dimensions	• application • direction - Possible values: – in – out • rcvd_priority - Possible values: any value between 0 and 31 • eval_priority - Possible values: any value between 0 and 31 • interfaceType - Possible values: – Nbsf_Management – Npcf_Binding_Audit • msgType - Possible values: – Register – Deregister – Discovery – Audit_Notify
Example	<pre>ocbsf_http_message_priority_request_count_total{application="bsf-management-service",direction="in",eval_priority="6",interfaceType="Nbsf_Management",msgType="Register",rcvd_priority="15"} 1.0</pre>

Table 7-58 ocbsf_http_message_priority_response_count_total

Field	Details
Description	This metric is pegged when BSF receives an incoming response or when there is an outgoing response from BSF during the following scenarios: • register • deregister • discovery • staleSession
Type	Counter

Table 7-58 (Cont.) ocbsf_http_message_priority_response_count_total

Field	Details
Dimensions	- application - direction Possible values: - in - out - rcvd_priority Possible values: any value between 0 and 31 - eval_priority Possible values: any value between 0 and 31 - interfaceType Possible values: - Nbsf_Management - Npcf_Binding_Audit - msgType Possible values: - Register - Deregister - Discovery - Audit_Notify - response_code Possible values: 1xx 2xx 3xx 4xx 5xx Note: response_code always reflects the specific code received, such as 200, 201, 404, or 500. References to 1xx, 2xx, 3xx, 4xx, or 5xx are used as placeholders for groups of response codes and should not be interpreted as actual output values from the metric.
Example	ocbsf_http_message_priority_response_count_total{application="bsf-management-service",direction="out",eval_priority="23",interfaceType="Nbsf_Management",msgType="Register",rcvd_priority="NA",response_code="201"} 4.0

Table 7-59 data_repository_invocations_seconds_sum

Field	Details
Description	The seconds_sum metric for Data repository invocations is a measure of the total time spent executing repository methods.
Type	Gauge

Table 7-59 (Cont.) data_repository_invocations_seconds_sum

Field	Details
Dimension	• app_kubernetes_io_instance • app_kubernetes_io_managed_by • app_kubernetes_io_name • app_kubernetes_io_part_of • app_kubernetes_io_version • application • container • endpoint • engVersion • exception • helm_sh_chart • instance • io_kompose_service • job • method • microservice • mktgVersion • namespace • pod • pod_template_hash • repository • security_istio_io_tlsMode • service_istio_io_canonical_name • service_istio_io_canonical_revision • state • vendor
Example	data_repository_invocations_seconds_sum{app_kubernetes_io_instance="ocbsf", app_kubernetes_io_managed_by="Helm", app_kubernetes_io_name="audit-service", app_kubernetes_io_part_of="ocbsf", app_kubernetes_io_version="1.0.0.0", application="ocbsf", container="audit-service", endpoint="cnc-metrics", engVersion="24.1.4", exception="None", helm_sh_chart="audit-service-24.1.4", instance="10.233.84.88:9000", io_kompose_service="ocbsf-ocpm-audit-service", job="ocne-infra/ocne-nf-cnc-podmonitor", method="getQueuedTablesToAudit", microservice="ocbsf-audit", mktgVersion="1.0.0.0", namespace="ocbsf", pod="ocbsf-ocpm-audit-service-6b6f85f7b9-nshv9", pod_template_hash="6b6f85f7b9", repository="AuditSchedulerRepository", security_istio_io_tlsMode="istio", service_istio_io_canonical_name="audit-service", service_istio_io_canonical_revision="1.0.0.0", state="SUCCESS", vendor="Oracle"}

Table 7-60 data_repository_invocations_seconds_count

Field	Details
Description	This metric measures the total time spent in Data JPA repository invocations. It is a counter metric, which means that it increments each time a repository method is invoked. The metric is reported in seconds.
Type	Counter

Table 7-60 (Cont.) data_repository_invocations_seconds_count

Field	Details
Dimension	- app_kubernetes_io_instance - app_kubernetes_io_managed_by - app_kubernetes_io_name - app_kubernetes_io_part_of - app_kubernetes_io_version - application - container - endpoint - engVersion - exception - helm_sh_chart - instance - io_kompose_service - job - method - microservice - mktgVersion - namespace - pod - service - pod_template_hash - repository - security_istio_io_tlsMode - service_istio_io_canonical_name - service_istio_io_canonical_revision - state - vendor
Example	data_repository_invocations_seconds_count{app_kubernetes_io_instance="ocbsf", app_kubernetes_io_managed_by="Helm", app_kubernetes_io_name="audit-service", app_kubernetes_io_part_of="ocbsf", app_kubernetes_io_version="1.0.0.0", application="ocbsf", container="audit-service", endpoint="cnc-metrics", engVersion="24.1.4", exception="None", helm_sh_chart="audit-service-24.1.4", instance="10.233.127.211:9000", io_kompose_service="ocbsf-ocpm-audit-service", job="occne-infra/occne-nf-cnc-podmonitor", method="findAll", microservice="ocbsf-audit", mktgVersion="1.0.0.0", namespace="ocbsf", pod="ocbsf-ocpm-audit-service-6b6f85f7b9-nmzr4", pod_template_hash="6b6f85f7b9", repository="AuditRegistrationsRepository", security_istio_io_tlsMode="istio", service_istio_io_canonical_name="audit-service", service_istio_io_canonical_revision="1.0.0.0", state="SUCCESS", vendor="Oracle"}

Table 7-61 data_repository_invocations_seconds_max

Field	Details
Description	This metric is used to track the maximum execution time of Data repository invocations. It is a useful metric for monitoring the performance of data repositories and identifying slow-running queries.
Type	Gauge

Table 7-61 (Cont.) data_repository_invocations_seconds_max

Field	Details
Dimension	• app_kubernetes_io_instance • app_kubernetes_io_managed_by • app_kubernetes_io_name • app_kubernetes_io_part_of • app_kubernetes_io_version • application • container • endpoint • engVersion • exception • helm_sh_chart • instance • io_kompose_service • job • method • microservice • mktgVersion • namespace • pod • pod_template_hash • repository • security_istio_io_tlsMode • service_istio_io_canonical_name • service_istio_io_canonical_revision • state • vendor
Example	data_repository_invocations_seconds_max{app_kubernetes_io_instance="ocbsf", app_kubernetes_io_managed_by="Helm", app_kubernetes_io_name="audit-service", app_kubernetes_io_part_of="ocbsf", app_kubernetes_io_version="1.0.0.0", application="ocbsf", container="audit-service", endpoint="cnc-metrics", engVersion="24.1.4", exception="None", helm_sh_chart="audit-service-24.1.4", instance="10.233.127.211:9000", io_kompose_service="ocbsf-ocpm-audit-service", job="ocne-infra/ocne-nf-cnc-podmonitor", method="getQueuedTablesToAudit", microservice="ocbsf-audit", mktgVersion="1.0.0.0", namespace="ocbsf", pod="ocbsf-ocpm-audit-service-6b6f85f7b9-nmzr4", pod_template_hash="6b6f85f7b9", repository="AuditSchedulerRepository", security_istio_io_tlsMode="istio", service_istio_io_canonical_name="audit-service", service_istio_io_canonical_revision="1.0.0.0", state="SUCCESS", vendor="Oracle"}

Table 7-62 ocbsf_ar_request_total

Field	Details
Description	This metric is used to count the alternate route requests.
Type	Counter
Dimensions	• application • fqdn • scheme

Table 7-62 (Cont.) ocbf_ar_request_total

Field	Details
Example	ocbf_ar_request_total{application="ocbf", fqdn="nf11stub.ir-ns3.svc", scheme="http"} 1.0

Table 7-63 ocbf_ar_response_total

Field	Details
Description	This metric is used to count the alternate route responses.
Type	Counter
Dimensions	• application • fqdn • scheme • responseCode
Example	ocbf_ar_response_total{application="ocbf", fqdn="nf11stub.ir-ns3.svc", responseCode="2xx", scheme="http"} 1.0

Table 7-64 ocbf_ar_failure_total

Field	Details
Description	This metric is used to count the alternate route retry failures.
Type	Counter
Dimensions	• application • fqdn • scheme
Example	ocbf_ar_failure_total{application="ocbf", fqdn="nf11stub.ir-ns3.svc", scheme="http"} 1.0

Table 7-65 ocpm_nfDiscovery_request_total

Field	Details
Description	This metric is used to count the nrf-client requests.
Type	Counter
Dimensions	• application • queryParams • targetNfType
Example	ocpm_nfDiscovery_request_total{application="bsf-management-service", queryParameters="target-nf-set-id,service-names,target-nf-type,requester-nf-type", targetNfType="PCF"} 1.0

Table 7-66 ocpm_nfDiscovery_response_total

Field	Details
Description	This metric is used to count the nrf-client responses.
Type	Counter
Dimensions	• application • queryParams • targetNfType • targetNfType cause • cause • responseCode
Example	ocpm_nfDiscovery_response_total{application="bsf-management-service", cause="", queryParameters="target-nf-set-id, service-names, target-nf-type, requester-nf-type", responseCode="200", targetNfType="PCF"} 1.0

Table 7-67 ocbsf_feature_incompatibility

Field	Details
Description	Indicates if one or more of the configured features is incompatible with the system or with each other.
Type	Gauge
Dimensions	incompatibleFeatures (list of features in alphabetical order)
Example	ocbsf_feature_incompatibility{application="bsf-management-service", incompatibleFeatures="ENABLE_PCF_BINDING_TABLE_MIGRATION=true"} 1.0

Table 7-68 ocbsf_binding_record_lookup_latency

Field	Details
Description	Measures the latency time taken by the lookup attempts on both tables.
Type	Histogram
Dimensions	• tableName • duration (le for bucket)
Example	ocbsf_binding_record_lookup_latency_seconds_bucket{application="bsf-management-service", tableName="pcf_binding", le="0.01"} 3

Table 7-69 ocbsf_binding_table_lookup_total

Field	Details
Description	Pegs the number of lookup attempts performed against a certain table. Note 1: During binding deregistration, this metric is pegged only for the pcf_binding_v2 path. The V2 path performs a lookup before deletion, whereas the pcf_binding path deletes the record directly using the primary key binding_id. Note 2: During a session viewer query, the search criteria contains only secondary keys and may match records in either the pcf_binding table or the pcf_binding_v2 table. Therefore, both tables are looked up and the metric is pegged for both tables. When returning the delete result, BSF performs an additional lookup to confirm whether the entry is not found. As a result, the metric is also pegged for both tables with result="not_found"
Type	Counter
Dimensions	- tableName - pcf_binding - pcf_binding_v2 - result: - FOUND - NOT_FOUND
Example	ocbsf_binding_table_lookup_total{application="bsf-management-service",result="found",tableName="pcf_binding"} 6.0

Table 7-70 ocbsf_binding_query_processing_latency

Field	Details
Description	Measures the latency time taken by a binding query/discovery.
Type	Histogram
Dimensions	- queryCriteriaType (below are the types sorted in order if one query has multiple data types) - supi - gpsi - ipv4Addr - ipv6Prefix - macAddr48 - operation_type - responseCode - duration (le for bucket)
Example	ocbsf_binding_query_processing_latency_seconds_bucket{application="bsf-management-service",operation_type="discover",queryCriteria="ipv4",response_code="200",le="0.01"} 0

Table 7-71 ocbsf_alt_key_lookup_total

Field	Details
Description	Number of times a query has been performed to an Alt Lookup Table.
Type	Counter
Dimensions	- Alt Lookup Table - IPv4 Addr - IPv6 Prefix - Mac Addr 48 - Status - success - failure - Cause - none - not_found - Exception Name - Rows fetched 0 1 2 3+
Example	ocbsf_alt_key_lookup_total{alt_key="ipv4_addr_lookup",application="bsf-management-service",cause="none",rows_fetched="1",status="successful"} 3.0

Table 7-72 ocbsf_index_based_lookup_fallback_total

Field	Details
Description	Number of times a fall-back to Index Based Lookup has been performed after not finding a V2 record through Alt Lookup Tables.
Type	Counter
Dimensions	- Alt Lookup Table - IPv4 Addr - IPv6 Prefix - Mac Addr 48 - Status - found - not_found
Example	ocbsf_index_based_lookup_fallback_total{alt_key="ipv4_addr_lookup",application="bsf-management-service",status="found"} 3.0

Table 7-73 ocbsf_lookup_table_reinsertions_total

Field	Details
Description	Number of times a Lookup Record has been re-inserted for a non-stale PCF Binding.
Type	Counter
Dimensions	- Alt Lookup Table - IPv4 Addr - IPv6 Prefix - Mac Addr 48
Example	ocbsf_lookup_table_reinsertions_total{alt_key="ipv4_addr_lookup",application="bsf-management-service"} 3.0 ocbsf

Table 7-74 ocbsf_lookup_table_stale_deletions_total

Field	Details
Description	Number of times a Stale Lookup Record has been deleted.
Type	Counter
Dimensions	- Alt Lookup Table - IPv4 Addr - IPv6 Prefix - Mac Addr 48
Example	ocbsf_lookup_table_stale_deletions_total{alt_key="ipv4_addr_lookup",application="bsf-management-service"} 3.0

Table 7-75 ocbsf_db_latency_seconds_count

Field	Details
Description	Latency for Alt Lookup Table DB operations (INSERT, SELECT, DELETE, UPDATE) in seconds.
Type	Timer
Dimensions	- Alt Lookup Table - IPv4 Addr - IPv6 Prefix - Mac Addr 48 - Database Operation - insert - select_by_binding_id_and_ue_addr - select_by_ue_addr - update_by_binding_id - delete_by_binding_id - delete_by_binding_id_and_ue_addr

Table 7-75 (Cont.) ocbsf_db_latency_seconds_count

Field	Details
Example	ocbsf_db_latency_seconds_count{alt_key="ipv4_addr_lookup",application="bsf-management-service",db_operation="delete_by_binding_id_and_ue_addr",exception="none",status="failure"} 1

Table 7-76 ocbsf_db_latency_seconds

Field	Details
Description	Determines the maximum amount of seconds of DB Latency a quantile had. It is grouped by quantiles. E.g: If quantile "0.5" has `3.5` ms, it means that half of the processed DB requests had a latency lower or equal to 3.5 ms.
Type	Timer
Dimensions	- Alt Lookup Table - IPv4 Addr - IPv6 Prefix - Mac Addr 48 - Database Operation - insert - select_by_binding_id_and_ue_addr - select_by_ue_addr - update_by_binding_id - delete_by_binding_id - delete_by_binding_id_and_ue_addr
Example	ocbsf_db_latency_seconds_count{alt_key="ipv4_addr_lookup",application="bsf-management-service",db_operation="delete_by_binding_id_and_ue_addr",exception="none",status="successful"} 2

Table 7-77 ocbsf_db_latency_seconds_sum

Field	Details
Description	Sums the total amount of DB Latency.
Type	Timer

Table 7-77 (Cont.) ocbsf_db_latency_seconds_sum

Field	Details
Dimensions	- Alt Lookup Table - IPv4 Addr - IPv6 Prefix - Mac Addr 48 - Database Operation - insert - select_by_binding_id_and_ue_addr - select_by_ue_addr - update_by_binding_id - delete_by_binding_id - delete_by_binding_id_and_ue_addr
Example	ocbsf_db_latency_seconds_count{alt_key="ipv4_addr_lookup",application="bsf-management-service",db_operation="insert",exception="none",status="successful"} 6 ocbsf_db_latency_seconds_count{alt_key="ipv4_addr_lookup",application="bsf-management-service",db_operation="select_by_binding_id_and_ue_addr",exception="none",status="failure"} 3

Table 7-78 ocbsf_db_latency_seconds_max

Field	Details
Description	The max amount of DB latency that has been seen so far.
Type	Timer
Dimensions	- Alt Lookup Table - IPv4 Addr - IPv6 Prefix - Mac Addr 48 - Database Operation - insert - select_by_binding_id_and_ue_addr - select_by_ue_addr - update_by_binding_id - delete_by_binding_id - delete_by_binding_id_and_ue_addr

Table 7-78 (Cont.) ocbsf_db_latency_seconds_max

Field	Details
Example	<pre>ocbsf_db_latency_seconds_count{alt_key= "ipv4_addr_lookup",application="bsf- management- service",db_operation="select_by_ue_add r",exception="none",status="failure"} 9 ocbsf_db_latency_seconds_count{alt_key= "ipv4_addr_lookup",application="bsf- management- service",db_operation="select_by_ue_add r",exception="none",status="successful" } 3</pre>

7.8 Collision Detection Metrics

Table 7-79 ocbsf_collision_detection_bad_request_code

Field	Details
Description	This metric is pegged when a BAD_REQUEST error code is detected.
Metric Type	Counter
Dimensions	• DNN • SNSSAI • APPLICATION_ERROR_CODE • RESPONSE_CODE • BINDING_LEVEL • BINDING_ID

Table 7-80 ocbsf_collision_detection_not_found_code

Field	Details
Description	This metric is pegged when a NOT_FOUND error code is detected.
Type	Counter
Dimensions	• DNN • SNSSAI • APPLICATION_ERROR_CODE • RESPONSE_CODE • BINDING_LEVEL • BINDING_ID

Table 7-81 ocbsf_collision_detection_forbidden_code

Field	Details
Description	This metric is pegged when a FORBIDDEN error code is detected.
Type	Counter

Table 7-81 (Cont.) ocbsf_collision_detection_forbidden_code

Field	Details
Dimensions	• DNN • SNSSAI • APPLICATION_ERROR_CODE • RESPONSE_CODE • BINDING_LEVEL • BINDING_ID

Table 7-82 ocbsf_collision_detection_unauthorized_code

Field	Details
Description	This metric is pegged when an UNAUTHORIZED error code is detected.
Type	Counter
Dimensions	• DNN • SNSSAI • APPLICATION_ERROR_CODE • RESPONSE_CODE • BINDING_LEVEL • BINDING_ID

Table 7-83 ocbsf_collision_detection_not_acceptable_code

Field	Details
Description	This metric is pegged when a NON_ACCEPTABLE error code is detected.
Type	Counter
Dimensions	• DNN • SNSSAI • APPLICATION_ERROR_CODE • RESPONSE_CODE • BINDING_LEVEL • BINDING_ID

Table 7-84 ocbsf_collision_detection_service_unavailable_code

Field	Details
Description	This metric is pegged when an UNAVAILABLE error code is detected.
Type	Counter
Dimensions	• DNN • SNSSAI • APPLICATION_ERROR_CODE • RESPONSE_CODE • BINDING_LEVEL • BINDING_ID

Table 7-85 ocbsf_collision_detection_method_not_allowed_code

Field	Details
Description	This metric is pegged when a NOT_ALLOWED error code is detected.
Type	Counter
Dimensions	• DNN • SNSSAI • APPLICATION_ERROR_CODE • RESPONSE_CODE • BINDING_LEVEL • BINDING_ID

Table 7-86 ocbsf_collision_detection_error_code

Field	Details
Description	This metric is pegged when a GENERAL error code is detected.
Type	Counter
Dimensions	• DNN • SNSSAI • APPLICATION_ERROR_CODE • RESPONSE_CODE • BINDING_LEVEL • BINDING_ID

7.9 CM Service Metrics

The following table describes the CM Service metrics and respective dimensions:

Table 7-87 system_operational_state

Field	Details
Description	This metric indicates the current operational state
Type	Gauge
Dimensions	• application • eng_version • namespace • node • pod • vendor
Example	system_operational_state{application="config-server",eng_version="",microservice="",namespace="",node="",pod="occnp-config-mgmt",vendor="oracle",} 1.0

7.10 Diameter Gateway Metrics

Table 7-88 ocbsf_diam_response_local_total

Field	Details
Description	When the diameter request is timed out from the backend peer (diameter connector), this metric is incremented.
Type	Counter
Dimensions	• applId • cmdCode • msgType • reqDestHost • reqDestRealm • reqOrigHost • responseCode • reqOrigRealm
Example	<pre>ocbsf_diam_response_local_total{applId="16777236",cmdCode="265",direction="in",msgType="AAA",reqDestHost="ocbsf",reqDestRealm="xxx.com",reqOrigHost="",reqOrigRealm="",responseCode="timeout"},} 1.0</pre> <pre>ocbsf_diam_request_network_total{applId="16777236",cmdCode="265",destHost="pcf2.example.com",destRealm="example.com",direction="out",msgType="AAR",origHost="bsf-diam-gateway",origRealm="oracle.com",retry="true",retryReason="timeout",immediatePeerHost="dra2.oracle.com",immediatePeerRealm="oracle.com"} 1.0</pre>

Table 7-89 ocbsf_diam_response_network_total

Field	Details
Description	When the diameter request is timed out from the external NF (AF or PCF), this metric is incremented.
Type	Counter
Dimensions	• applId • cmdCode • direction • msgType • reqDestHost • reqDestRealm • reqOrigHost • responseCode • reqOrigRealm • immediatePeerHost • immediatePeerRealm
Example	<pre>ocbsf_diam_response_network_total{applId="16777236",cmdCode="265",direction="in",msgType="AAA",reqDestHost="oc-diam-gateway",reqDestRealm="xxx.com",reqOrigHost="",reqOrigRealm="",responseCode="timeout",immediatePeerHost="pgw.oracle.com",immediatePeerRealm="oracle.com"} 1.0</pre>

Table 7-90 ocbsf_diam_pending_requests

Field	Details
Description	This tracks the total pending diameter requests.
Type	Gauge
Dimension	- namespace - pod
Examples	ocbsf_diam_pending_requests 1.0

Table 7-91 diam_controlled_shutdown_message_reject_total

Field	Details
Description	Indicates failure count because of forced shutdown feature.
Type	Counter
Dimensions	- state - response_code - destHost - destRealm - origHost - origRealm - appId - cmdCode - msgType - direction
Example	diam_controlled_shutdown_message_reject_total{msgType="Gx_CCR_I",appId="16777238",cmdCode="272",destHost="",destRealm="xxx.com",operationalState="PARTIAL_SHUTDOWN",origHost="pgw.xxx.com",origRealm="test.example.com",responseCode="5012"}, 1.0

Table 7-92 topic_version

Field	Details
Description	Describes the current applied version of a given topic (mentioned in dimension topic_name) into the pod.
Type	Gauge
Dimensions	- Service Name - Pod Name

Table 7-93 ocbsf_diam_request_network_total

Field	Details
Description	Tracks total number of request messages of given command code to or from network.
Type	Counter

Table 7-93 (Cont.) ocbsf_diam_request_network_total

Field	Details
Dimensions	• appld • cmdCode • destHost • destRealm • direction • msgType • origHost • origRealm • retry • retryReason • immediatePeerHost • immediatePeerRealm
Example	ocbsf_diam_request_network_total{appld="16777236",cmdCode="265",direction="out",msgType="AAA",reqDestHost="oc-diam-gateway",reqDestRealm="xxx.com",reqOrigHost="",reqOrigRealm="",responseCode="timeout",immediatePeerHost="pgw.oracle.com",immediatePeerRealm="oracle.com"} 1.0

Table 7-94 ocbsf_diam_request_inter_total

Field	Details
Description	Tracks total number of request messages of given command code to or from host services. It will increase only in the case of "inter-pod routing".
Type	Counter
Dimensions	• appld • cmdCode • destHost • destRealm • direction • msgType • origHost • origRealm • retry • retryReason

Table 7-95 ocbsf_diam_realm_validation_failed_total

Field	Details
Description	Used to count the number of failed Destination-Realm validation at Diameter Gateway for BSF.
Type	Counter

Table 7-95 (Cont.) ocbsf_diam_realm_validation_failed_total

Field	Details
Dimensions	• applId • cmdCode • destHost • destRealm • origHost • origRealm • responseCode • validateRealmFailed • validateBindingRealmFailed
Example	ocbsf_diam_realm_validation_failed_total{applId="16777236",cmdCode="265",destHost="",destRealm="xxx.com",origHost="af.xxx.com",origRealm="xxx.com",responseCode="3003",validateRealmFailed="false", validateBindingRealmFailed="true"} 1.0 ; Type-Counter

Table 7-96 ocbsf_ocnf_service_cookie_total

Field	Details
Description	The metrics with direction "IN" is pegged only when BSF response header has ocnfServiceCookie. The metrics with direction "OUT" is pegged only when OCNF-Service-Cookie customAVP is present while sending the request to destination peer.
Type	Counter
Dimensions	• direction • source
Example	ocbsf_ocnf_service_cookie_total{direction="Received",source="Bsf-Mgt-Service"} 1.0 ocbsf_ocnf_service_cookie_total{direction="sent",source=""} 1.0

Table 7-97 ocbsf_diam_service_overall_processing_time_seconds_bucket

Field	Details
Description	Tracks latency of all incoming requests towards Diameter Gateway. This provides an overall time tracking latency for any request towards Diameter Gateway.
Type	Histogram
Dimensions	• applId • cmdCode • msgType • origHost • destHost • le
Example	-

Table 7-98 diam_repaving_disconnections_total

Field	Details
Description	This metric counts the total number of connections terminated by the "force disconnect" mechanism.
Type	Counter
Dimension	- peer_realm - peer_host
Example	diam_repaving_disconnections_total{origHost="pgw.oracle.com",origRealm="oracle.com"} 1.0

Table 7-99 ocbsf_mesh_map_missing_dgw_peer_connection_total

Field	Details
Description	This metric represents the number of times the DGW_IP_ADDRESS entry in the mesh map is detected as missing.
Type	Counter
Dimension	dgwIP
Example	ocbsf_mesh_map_missing_dgw_peer_connection_total {DGW_IP_ADDRESS="10.233.114.26",} 2.0 ocbsf_mesh_map_missing_dgw_peer_connection_total {DGW_IP_ADDRESS="10.233.97.189",} 2.0

7.11 NRF Management Service

Table 7-100 NRF Instance Status

Field	Details
Description	If the metric has value 0 - NRF is unavailable/unhealthy 1 - NRF is available/healthy
Type	Gauge
Dimensions	NrfUri - URI of the NRF Instance

Table 7-101 NRF Instance Status Count

Field	Details
Description	The apiRoot shall be specified in the following format: <i>'scheme'://'fqdn':port'</i> If health check procedure is disabled, all NRF instances are marked as HEALTHY after successful NfRegistration.
Type	Gauge

Table 7-101 (Cont.) NRF Instance Status Count

Field	Details
Dimensions	NrfUri - URI of the NRF Instance HealthStatus FailureReason - Reason for the status

Table 7-102 NRF Instance Consecutive Healthy Count

Field	Details
Description	The metric shall have a minimum value of 0 and maximum value of healthCheckCount.
Type	Gauge
Dimensions	NrfUri - URI of the NRF Instance

Table 7-103 NRF Instance Consecutive Unhealthy Count

Field	Details
Description	The metric shall have a minimum value of 0 and maximum value of healthCheckCount.
Type	Gauge
Dimensions	NrfUri - URI of the NRF Instance

Table 7-104 DNS lookup requests

Field	Details
Description	The metric shall be pegged only if enableVirtualNrfResolution is set to true.
Type	Gauge
Dimensions	Scheme - http or https VirtualFqdn - Fqdn that shall be used by the alternate service for the DNS lookup.

Table 7-105 DNS lookup responses

Field	Details
Description	The metric shall be pegged only if enableVirtualNrfResolution is set to true.
Type	Gauge
Dimensions	Scheme - http or https VirtualFqdn - Fqdn that shall be used by the alternate service for the DNS lookup. HttpStatusCode - The status code as received in the response.

Table 7-106 DNS setup requests

Field	Details
Description	The metric shall be pegged only if enableVirtualNrfResolution is set to true.
Type	Counter
Dimensions	Scheme - http or https VirtualFqdn - Fqdn that shall be used by the alternate service for the DNS lookup.

Table 7-107 nrfclient_nfUpdate_status

Field	Details
Description	Indicates if the NfRegister/NfUpdate operation is successful with NRF after the NfProfile has been updated using the REST API. 0 - Update is not successful. 1 - Update is successful. NFs may use this metric to configure alerts.
Type	Gauge
Dimensions	- NfInstanceId - The NF's nInstanceId as present in the NfProfile. - NfFqdn - The NF's FQDN as present in the NfProfile.

7.12 Overload Control Metrics

Table 7-108 service_resource_stress

Field	Details
Description	This metric tracks CPU, memory, failure count, and pending requests on the basis of which the overload level of a service is calculated.
Type	Gauge
Dimensions	type, Service

Table 7-109 service_resource_overload_level

Field	Details
Description	This metric tracks an individual resource's overload level that is calculated based on the resource usage and configured threshold.
Type	Gauge
Dimensions	type, Service

Table 7-110 load_level

Field	Details
Description	Provides information about the overload manager load level.
Type	Gauge
Dimensions	Service, namespace

Table 7-111 diam_overload_message_reject_total

Field	Details
Description	This metric tracks the total number of messages rejected due to overload on Diameter interface.
Type	Counter
Dimensions	priority, response_code Other dimensions: destHost, destRealm, origHost, origRealm, applId, cmdCode, msgType, direction

7.13 PerfInfo Metrics

Table 7-112 nf_load_info

Field	Details
Description	Provides information about service load.
Type	Gauge
Dimensions	- service - namespace
Example	nf_load_info{namespace="xxx",service="xxx"} 0.8486912141984638

Table 7-113 jvm_cpu_usage

Field	Details
Description	Per service jvm_cpu_usage.
Type	Gauge
Dimensions	- service - namespace
Example	jvm_cpu_usage{namespace="xxx",service="xxx"} 0.2758240242725142

Table 7-114 jvm_memory

Field	Details
Description	Per service jvm_memory.
Type	Gauge
Dimensions	- service - namespace
Example	jvm_memory{namespace="ttz",service="xxx"} 18.361382484436035

Table 7-115 cgroup_cpu_nanoseconds

Field	Details
Description	Reports the total CPU time (in nanoseconds) on each CPU core for all the tasks in the cgroup.
Type	Gauge
Dimensions	NA

Table 7-115 (Cont.) cgroup_cpu_nanoseconds

Field	Details
Example	cgroup_cpu_nanoseconds 2.1782821080274e+013

Table 7-116 cgroup_memory_bytes

Field	Details
Description	Reports the memory usage.
Type	Gauge
Dimensions	NA
Example	cgroup_memory_bytes 1.31289088e+08

Table 7-117 load_level

Field	Details
Description	This metric provides information about the load level of a service.
Type	Gauge
Dimension	- service - namespace - isLeaderPod
Example	load_level{service="xxx"} L1

Table 7-118 system_overload_threshold_config_mode

Field	Details
Description	Indicates whether the overload level threshold configuration is based on STANDALONE or PROFILE mode.
Type	Gauge
Dimension	- namespace - isLeaderPod
Example	system_overload_threshold_config_mode 1.0

Table 7-119 active_overload_threshold_fetch_failed

Field	Details
Description	Indicates whether the active profile data is fetched successfully or failed to fetch.
Type	Gauge
Dimension	- namespace - isLeaderPod The value of this dimension can be either 0 or 1. Where 0, represents "Successfully fetched the active threshold" and 1 represents "Failure in fetching the active threshold".
Example	active_overload_threshold_fetch_failed 1.0

Table 7-120 load_level_report_total

Field	Details
Description	This metric is used to track: - the number of times load level calculation is performed - the number of times load level changes - how long the particular level was active
Type	Counter
Dimension	- level - service - prevLevel - namespace - levelChangeType - isLeaderPod
Example	load_level_report_total{namespace="hi-riley", service="pcf-occpn-pcrf-core", isLeaderPod="True", level="Normal", levelChangeType="-"} 2.0

Table 7-121 service_resource_stress

Field	Details
Description	This metric tracks CPU, memory, failure count, and pending requests on the basis of which the overload level of a service is calculated.
Type	Gauge
Dimension	- type - service - namespace - isLeaderPod
Example	service_resource_stress{service="xxx", type="xxx"} 10.0

Table 7-122 service_resource_overload_level

Field	Details
Description	This metric tracks an individual resource's overload level that is calculated based on the resource usage and configured threshold.
Type	Gauge
Dimension	- type - service - namespace - isLeaderPod
Example	service_resource_overload_level{service="xxx", type="xxx"}2.0

Table 7-123 service_resource_overload_level_report_total

Field	Details
Description	This metric is used to track: - the number of times load level calculation is performed. - the number of times load level changes. - for how long the particular level was active for each metric type.
Type	Counter

Table 7-123 (Cont.) service_resource_overload_level_report_total

Field	Details
Dimension	- level Possible values: - L1 - L2 - L3 - Normal - service - prevLevel - levelChangeType Possible values: - increment - decrement - isLeaderPod Possible values: - True - False - type Possible values: - cpu - memory - svc_failure_count - svc_pending_count - namespace
Example	service_resource_overload_level_report_total{namespace="hi-riley", service="pcf-occp- pcrf-core", isLeaderPod="True", level="Normal", levelChangeType="-", type="cpu"} 2.0

Table 7-124 http_out_conn_request

Field	Details
Description	This counter metric is used to count the number of http API Egress requests.
Type	Counter
Dimension	- servicenameNon3gpp - serviceResource - serviceVersion - operationType - namespace - params - isLeaderPod
Example	-

Table 7-125 http_out_conn_response

Field	Details
Description	This counter metric is used to count the number of http API Egress responses.
Type	Counter

Table 7-125 (Cont.) http_out_conn_response

Field	Details
Dimension	• servicenameNon3gpp • serviceResource • serviceVersion • outcome • namespace • operationType • responseCode • params • cause • isLeaderPod
Example	-

Table 7-126 overload_manager_enabled

Field	Details
Description	This metric indicates whether overload manager is enabled or disabled.
Type	Gauge
Dimension	• source Possible values: – DIAM_GW – INGRESS_GW – PERF_INFO • namespace
Example	overload_manager_enabled{namespace="hi-riley", source="DIAM_GW"} 1

Table 7-127 leader_pod

Field	Details
Description	This metric is used to know the leader pod.
Type	Gauge
Dimension	• namespace
Example	leader_pod{namespace="hi-riley"} 1

For more information about dimensions, see [BSF Metrics](#).

7.14 Pod Congestion Metrics

Table 7-128 ocbsf_pod_congestion_state

Field	Details
Description	Tracks congestion state of pod.
Type	Gauge

Table 7-128 (Cont.) ocbsf_pod_congestion_state

Field	Details
Dimensions	- podName - namespace - level = 0,1, 2, 3, 4 (0=Normal, 1=DANGER_OF_CONGESTION, 2=CONGESTION_L1, 3=CONGESTION_L2, 4=CONGESTED)

Table 7-129 ocbsf_pod_resource_stress

Field	Details
Description	Tracks CPU, memory, queue usage (percentage) based on which POD is calculating its congestion state.
Type	Gauge
Dimensions	type = "PendingRequest","CPU","Memory"

Table 7-130 ocbsf_pod_resource_congestion_state

Field	Details
Description	Tracks individual resource's congestion state calculated based on the resource usage and configured threshold.
Type	Gauge
Dimensions	type = "PendingRequest","CPU","Memory" level = 0,1,2 (0 = Normal, 1 = DoC, 2 = Congested)

Table 7-131 pod_cong_state_report_total

Field	Details
Description	Tracks the total count of change of congestion state.
Type	Counter
Dimensions	- namespace - pod - old_state - new_state

Table 7-132 pod_resource_congestion_state_report_total

Field	Details
Description	Tracks the total count of change of resource specific congestion state.
Type	Counter
Dimensions	- resource - namespace - pod - old_state - new_state

Table 7-133 congestion_http_pending_request_active_counter

Field	Details
Description	Tracks the current active counter for calculating pending HTTP requests.
Type	Gauge
Dimensions	- namespace - pod

Table 7-134 ocbsf_diam_congestion_message_reject_total

Field	Details
Description	Tracks number of messages rejected due to congestion.
Type	Counter
Dimensions	priority = calculated or received DRMP priority of message being rejected response_code = response code sent with rejected message destHost, destRealm, origHost, origRealm, appId, cmdCode, <i>msgType</i>, direction

7.15 User-Agent Header Metrics

The following metric is introduced for User-Agent Header:

Table 7-135 oc.ingressgateway.http.requests

Field	Details
Description	This metric is pegged for every incoming request. If the User-Agent header is not present, then UNKNOWN will be pegged. This will be independent of the User Agent validation feature at the Ingress Gateway.
Type	Counter
Dimensions	- consumerNfType - consumerInstanceId - ConsumerFqdn

7.16 Query Service Metrics

Table 7-136 queryservice_request

Field	Details
Description	Number of query requests received.
Type	Counter
Dimensions	RequestType
Example	queryservice_response_total{Code="2xx",ReqType="GET",} 195.0

Table 7-137 queryservice_response

Field	Details
Description	Number of responses sent for query requests.
Type	Counter
Dimensions	RequestType
Example	queryservice_request_total{ReqType="GET",} 257.0

Table 7-138 queryservice_sessionDelete_request

Field	Details
Description	Total number of query requests to delete PCF bindings.
Type	Counter
Dimensions	- ResourceType - Result_Code
Example	queryservice_sessionDelete_request{ReqType="DELETE",} 257.0

Table 7-139 queryservice_sessionDelete_response

Field	Details
Description	Total number of responses for delete PCF bindings query requests.
Type	Counter
Dimensions	- ResourceType - Result_Code
Example	queryservice_sessionDelete_response{Code="2xx",ReqType="DELETE",} 195.0

7.17 TLS Metrics

The following table describes the TLS metrics and the respective dimensions:

Table 7-140 oc_ingressgateway_incoming_tls_connections

Field	Details
Description	Number of TLS Connections received on the Ingress Gateway and their negotiated TLS versions. The versions can be TLSv1.3 or TLSv1.2
Type	Gauge
Dimensions	- NegotiatedTLSVersion - direction - instanceIdentifier
Example	-

Table 7-141 oc_egressgateway_incoming_tls_connections

Field	Details
Description	Number of TLS Connections received on the Egress Gateway and their negotiated TLS versions. The versions can be TLSv1.3 or TLSv1.2

Table 7-141 (Cont.) oc_egressgateway_incoming_tls_connections

Field	Details
Type	Gauge
Dimensions	- NegotiatedTLSVersion - direction - instanceIdentifier
Example	-

Table 7-142 security_cert_x509_expiration_seconds

Field	Details
Description	Indicates the time to certificate expiry in epoch seconds.
Type	Histogram
Dimensions	serialNumber
Example	-

Table 7-143 diam_conn_network

Field	Details
Description	Indicates the number of TLS connections per TLS version.
Type	Gauge
Dimensions	tlsversion, peerHost, and peerRealm
Example	-

Table 7-144 diam_failed_conn_network

Field	Details
Description	Indicates the number of failed TLS connections. Note: It is applicable when we configure Initiate Connection to <code>true</code> in peer node configurations in the CNC Console.
Type	Gauge
Dimensions	- tlsversion - reason
Example	diam_failed_conn_network{peerName="dgw",reason="SSL Handshake Exception",} 1.0

Table 7-145 diam_conn_network_responder

Field	Details
Description	Indicates the number of allowed TLS responder connections with or without the peer configuration.
Type	Gauge
Dimensions	- tlsversion - peerconfigvalidated - peerHost - peerRealm

Table 7-145 (Cont.) diam_conn_network_responder

Field	Details
Example	-

Table 7-146 dgw_tls_cert_expiration_seconds

Field	Details
Description	Indicates the number of allowed TLS responder connections with or without the peer configuration.
Type	Gauge
Dimensions	- serialNumber (a number assigned by CA to each certificate) - subject (information about the cert issuer)
Example	dgw_tls_cert_expiration_seconds{serialNumber="122859034516052844067924392230910496431389299229",subject="OU=CGIU, O=ORCL, L=BLR, ST=KA, C=IN",} 1.797882352E9

7.18 NRF Client Metrics

NF status and NF load metrics

Table 7-147 nrclient_perf_info_nf_profile_load

Field	Details
Description	The current Load of the NF.
Type	Gauge
Dimensions	-

Table 7-148 nrclient_current_nf_status

Field	Details
Description	The current operative status of the NF. The gauge shall be indicate the status as below: 0 - REGISTERED 1 - DEREGISTERED 2 - SUSPENDED 3 - UNDISCOVERABLE 4 - UNKNOWN
Type	Gauge
Dimensions	- NfType - The NF's NfType as present in the registered NfProfile. - NfInstanceId - The NF's nfInstanceId as present in the registered NfProfile. - NfFqdn - The NF's FQDN as present in the registered NfProfile.

Table 7-149 nrfclient_nf_status_with_nrf

Field	Details
Description	The operative status of the NF communicated to the NRF. The gauge shall be indicate the status as below: • 0 - REGISTERED • 1 - DEREGISTERED • 2 - SUSPENDED • 3 - UNDISCOVERABLE • 4 - UNKNOWN
Type	Gauge
Dimensions	• NfType - The NF's NfType as present in the registered NfProfile. • NfInstanceId - The NF's nfiInstanceid as present in the registered NfProfile. • NfFqdn - The NF's FQDN as present in the registered NfProfile.

NRF Health Status

Table 7-150 nrfclient_nrf_operative_status

Field	Details
Description	The current operative status of the NRF Instance. If the metric has value • 0 - NRF is unavailable or unhealthy • 1 - NRF is available or healthy
Type	Gauge
Dimensions	NrfUri - URI of the NRF Instance

Table 7-151 nrfclient_nrf_status_total

Field	Details
Description	Total number of times an NRF instance is marked as healthy or unhealthy. The apiRoot is specified in the format ' <i>scheme</i> :// <i>fqdn</i> :' <i>port</i> '. Note: If health check procedure is disabled, all NRF instances are marked as HEALTHY after successful NF registration.
Type	Counter
Dimensions	NrfUri- URI of the NRF InstanceHealthStatus FailureReason - Reason for the status

Table 7-152 nrfclient_nrf_successive_healthy_count

Field	Details
Description	The metric shows the consecutive number of times the NRF is considered as healthy. The metric has a minimum value of 0 and maximum value of <i>healthCheckCount</i> .
Type	Counter
Dimensions	NrfUri- URI of the NRF Instance

Table 7-153 nrfclient_nrf_successive_unhealthy_count

Field	Details
Description	The metric shows the consecutive number of times the NRF is considered as unhealthy
Type	Counter
Dimensions	NrfUri- URI of the NRF Instance

NF - NRF-Client metrics**Table 7-154 nrfclient_on_demand_conn_in_request_total**

Field	Details
Description	Total number of on-demand requests received from the backend NF to NRF Client.
Type	Counter
Dimensions	• MessageType - The service request Type. • NfType - The NF's NfType as present in the registered NfProfile. • NfInstanceId - The NF's nfInstanceId as present in the registered NfProfile. • NfFqdn - The NF's FQDN as present in the registered NfProfile.

Table 7-155 nrfclient_on_demand_conn_out_response_total

Field	Details
Description	Total number of on-demand responses sent to the backend NF to NRF Client.
Type	Counter
Dimensions	• MessageType - The service request Type. • NfType - The NF's NfType as present in the registered NfProfile. • NfInstanceId - The NF's nfInstanceId as present in the registered NfProfile. • NfFqdn - The NF's FQDN as present in the registered NfProfile. • StatusCode - The HttpStatusCode as received from the NRF or generated by NRF-client.

Table 7-156 nrfclient_on_demand_processing_latency_ms

Field	Details
Description	Total message processing time duration in milliseconds.
Type	Histogram
Dimensions	• MessageType - The service request Type. • NfType - The NF's NfType as present in the registered NfProfile. • NfInstanceId - The NF's nfInstanceId as present in the registered NfProfile. • NfFqdn - The NF's FQDN as present in the registered NfProfile.

Table 7-157 ocpm_nrf_tracing_request_timeout_total

Field	Details
Description	Total number of requests timeout sent to the backend NF from NRF Client.
Type	Counter

Table 7-157 (Cont.) `ocpm_nrf_tracing_request_timeout_total`

Field	Details
Dimensions	- MessageType - The service request Type. - NfType - The NF's NfType as present in the registered NfProfile. - NfInstanceId - The NF's nfInstanceId as present in the registered NfProfile. - NfFqdn - The NF's FQDN as present in the registered NfProfile.

NRF-Client - NRF metrics

Table 7-158 `nrfclient_nw_conn_out_request_total`

Field	Details
Description	Total number of times NRF-client has sent a request to NRF. This includes autonomous requests as well as on-demand requests.
Type	Counter
Dimensions	- MessageType - The service request Type. - NfType - The NF's NfType as present in the registered NfProfile. - NfInstanceId - The NF's nfInstanceId as present in the registered NfProfile. - NfFqdn - The NF's FQDN as present in the registered NfProfile. - outgoingHeader - priorityValue

Table 7-159 `nrfclient_nw_conn_in_response_total`

Field	Details
Description	Total number of times NRF-client has received a response from NRF.
Type	Counter
Dimensions	- MessageType - The service request Type. - NfType - The NF's NfType as present in the registered NfProfile. - NfInstanceId - The NF's nfInstanceId as present in the registered NfProfile. - NfFqdn - The NF's FQDN as present in the registered NfProfile. - StatusCode - The HttpStatusCode as received from the NRF or generated by NRF-client.

Table 7-160 `nrfclient_nw_conn_in_notify_request_total`

Field	Details
Description	Total number of nfStatusNotify requests received from NRF.
Type	Counter
Dimensions	- EventType - The EventType for which the notification is triggered. - NfInstanceId - The NfInstanceId for which the notification is triggered.

Table 7-161 `nrfclient_nw_conn_out_notify_response_total`

Field	Details
Description	Total number of <i>nfStatusNotify</i> responses sent to NRF.
Type	Counter

Table 7-161 (Cont.) nrfclient_nw_conn_out_notify_response_total

Field	Details
Dimensions	- EventType - The EventType for which the notification is triggered. - NfInstanceID - The NfInstanceID for which the notification is triggered. - HttpStatusCode - The HttpStatusCode sent by NRF-Client.

Table 7-162 nrfclient_network_message_processing_latency

Field	Details
Description	Total message processing time duration.
Type	Histogram
Dimensions	- MessageType - The service request Type. - NfType - The NF's NfType as present in the registered NfProfile. - NfInstanceID - The NF's nfInstanceID as present in the registered NfProfile. - NfFqdn - The NF's FQDN as present in the registered NfProfile.

Table 7-163 nrfclient_nw_conn_out_request_total

Field	Details
Description	This metric provides the total number of requests sent out through NRF Client to the NRF (network) (inclusive of autonomous and on-demand).
Type	Counter
Dimensions	- MessageType - NfType - NfInstanceID - NfFqdn - outgoingHeader - priorityValue - NrfURI - TargetNfType - outgoingHeader - priorityValue
Example	ocbsf_nrfclient_nw_conn_out_request_total{MessageType="AutonomousNfRegistration",NfFqdn="ocbsf1-2-api-gateway.bs1-2.svc.atlantic.morrisville.us.lab.oracle.com",NfInstanceID="fe7d992b-0541-4c7d-ab84-c6d70b1b0666",NfType="BSF",NrfURI="nf1stub.hi-ns14.svc:8080",OutgoingPriorityHeader="3gpp-sbi-message-priority",Priority="16",TargetNfType="UNKNOWN"} 1.0

7.19 Metrics for Automated Certificate Lifecycle Management

The following metrics are used to support automated certificate lifecycle management for BSF:

Table 7-164 oc_certificatemanagement_tls_certificate_info

Field	Details
Description	This metric is used to peg status of TLS certificates. This metric is further used for raising alarms.

Table 7-164 (Cont.) oc_certificatemanagement_tls_certificate_info

Field	Details
Type	Gauge
Dimensions	• CertificateName • SecretName • Status (VALID, CORRUPT, MISSING, EXPIRED) • Service (IngressGateway, EgressGateway)
Example	occnp_oc_certificatemanagement_tls_certificate_info{CertificateName="key-password.txt,bsf.cer,pcf.pem,trust-chain-password.txt,trust-chain.cer",SecretName="bsf-igw-tls-ssl-bundle",Service="IngressGateway",Status="VALID",} 1.0

A

Ingress Gateway Metrics

Ingress Gateway Metrics

Table A-1 oc_ingressgateway_http_requests_total

Field	Details
Description	This metric is pegged as soon as the request reaches the Ingress gateway in the first custom filter of the application.
Type	Counter
Dimension	• NFType • NFServiceType • Host • HttpVersion • Scheme • Route_path • InstanceIdentifier • ClientCertIdentity

Table A-2 oc_ingressgateway_http_responses_total

Field	Details
Description	This metric is pegged in the last custom filter of the Ingress gateway while the response is being sent back to the consumer NF.
Type	Counter
Dimension	• Status • Method • Route_path • NFType • NFServiceType • Host • HttpVersion • Scheme • Identifier • ClientCertIdentity

Table A-3 oc_ingressgateway_request_latency_seconds

Field	Details
Description	This metric is pegged in the last custom filter of the Ingress gateway while the response is being sent back to the consumer NF. This metric tracks the amount of time taken for processing the request. It starts as soon the request reaches the first custom filter of the application and lasts till the response is sent back to the consumer NF from the last custom filter of the application.
Type	Timer

Table A-3 (Cont.) oc_ingressgateway_request_latency_seconds

Field	Details
Dimension	- quantile - InstanceIdentifier

Table A-4 oc_ingressgateway_connection_failure_total

Field	Details
Description	This metric is pegged in the customized Jetty Client as soon as it fails to connect to the destination service with direction as ingressOut. Here in case of Ingress gateway, the destination service is a backend microservice of the NF. And TLS connection failure metrics when connecting to ingress with direction as ingress.
Type	Counter
Dimension	- Host - Port - Direction - InstanceIdentifier - error_reason

Table A-5 oc_ingressgateway_global_ratelimit_total

Field	Details
Description	This metric is pegged in the custom filter implemented to check the global rate limit conditions.
Type	Counter
Dimension	- Method - Route_path - Scheme - InstanceIdentifier - Status (Rate limit Status field is different here)

Table A-6 oc_ingressgateway_route_ratelimit_total

Field	Details
Description	This metric is pegged in the custom filter implemented to check the route level rate limit conditions.
Type	Counter
Dimension	- Method - Route_path - Scheme - InstanceIdentifier - Status (Rate limit Status field is different here)

Table A-7 oc_ingressgateway_request_processing_latency_seconds

Field	Details
Description	This metric is pegged in the last custom filter of the Ingress gateway while the response is being sent back to the consumer NF. This metric captures the amount of time taken for processing of the request only within Ingress gateway. It starts as soon the request reaches the first custom filter of the application and lasts till the request is forwarded to the destination.
Type	Timer
Dimension	- quantile - InstanceIdentifier

Table A-8 oc_ingressgateway_jetty_request_stat_metrics_total

Field	Details
Description	This metric is pegged for every event occurred when a request is sent to IGW.
Type	Counter
Dimension	- event - client_type - InstanceIdentifier

Table A-9 oc_ingressgateway_jetty_response_stat_metrics_total

Field	Details
Description	This metric is pegged for every event occurred when a response is received by IGW
Type	Counter
Dimension	- event - client_type - InstanceIdentifier

Table A-10 oc_ingressgateway_jetty_latency_seconds

Field	Details
Description	This metric is pegged in Jetty response listener that captures the amount of time taken for processing of the request by jetty client.
Type	Timer
Dimension	- quantile - InstanceIdentifier

Table A-11 oc_ingressgateway_netty_latency_seconds

Field	Details
Description	This metric is pegged in Netty outbound handler that captures the amount of time taken for processing of the request by netty server.
Type	Timer
Dimension	- quantile - InstanceIdentifier

Table A-12 oc_ingressgateway_request_content_metrics_total

Field	Details
Description	This metric is pegged by default filter RequestContentMetrics. It pegs whether request has request body or not.
Type	Counter
Dimension	• method • content_available • InstanceIdentifier

Table A-13 oc_ingressgateway_xfcc_header_validate_total

Field	Details
Description	This metric is pegged when xfccHeaderValidation is enabled in XfccHeaderValidationFilter. This metric along with the specified dimension captures the successful/ un-successful validation of XFCC header in the incoming request.
Type	Counter
Dimension	• Route_path • Status • Cause • CertsCompared • InstanceIdentifier • ErrorOriginator

Table A-14 oc_configclient_request_total

Field	Details
Description	This metric is pegged whenever config client is polling for configuration update from common configuration server.
Type	Counter
Dimension	• Release version • Config version

Table A-15 oc_configclient_response_total

Field	Details
Description	This metrics is pegged whenever config client receives response from common configuration server.
Type	Counter
Dimension	• Release version • Config version • Updated

Table A-16 oc_ingressgateway_incoming_tls_connections

Field	Details
Description	Number of TLS Connections received on the Ingress Gateway and their negotiated TLS versions. The versions can be TLSv1.3 or TLSv1.2.
Type	Gauge
Dimension	- NegotiatedTLSVersion - Direction - InstanceIdentifier

Table A-17 security_cert_x509_expiration_seconds

Field	Details
Description	Indicates the time to certificate expiry in epoch seconds.
Type	Histogram
Dimension	- NegotiatedTLSVersion - Host - Direction - InstanceIdentifier

Ingress Metrics Common Tags

Tags	Description	Possible Values
Method	Http method	GET, PUT, POST, DELETE, PATCH
NFType	Name of the NF Type.	For example: Path is / nxxx-yyy /vz/..... Where XXX(Upper Case) is NFType UNKNOWN if unable to extract NFType from the path
NFServiceType	Name of the Service within the NF.	For example: Path is / nxxx-yyy /vz/..... Where nxxx-yyy is NFServiceType UNKNOWN if unable to extract NFServiceType from the path
Host	(Ip or fqdn): port of ingress gateway	NA
HttpVersion	Http protocol version	HTTP/1.1, HTTP/2.0
Scheme	Http protocol scheme	HTTP, HTTPS, UNKNOWN
ClientCertIdentity	Certificate Identity of the client	SAN=127.0.0.1,localhost CN=localhost, N/A if data is not available
Route_Path	Path predicate/Header predicate that matched the current request	NA
InstanceIdentifier	Prefix of the pod configured in helm when there are multiple instances in same deployment	Prefix configured in helm otherwise UNKNOWN
ErrorOriginator	This tag captures the ErrorOriginator	ServiceProducer, Nrf, IngresGW, None

Tags	Description	Possible Values
oc_ingressgateway_route_ratelimit_Status oc_ingressgateway_global_ratelimit_Status	Request accepted or dropped	accepted, dropped
oc_ingressgateway_connection_failure_Host	destination ip/fqdn	NA
oc_ingressgateway_connection_failure_Port	destination port	NA
oc_ingressgateway_xfcc_header_validate_Status	Https Status value after performing xfccHeaderValidation at IGW	200 (OK), 400 (BAD_REQUEST)
oc_ingressgateway_xfcc_header_validate_Cause	This tag determines the validation cause for the xfcc header validation metric being pegged	VALIDATION_FAILURE, VALIDATION_SUCCESS, HEADER_NOT_FOUND
oc_ingressgateway_xfcc_header_validate_CertsCompared	This tag captures the total number of certificates compared in XFCC header at IGW during the header validation	Count of the certificates compared (0,1,2..)
oc_configclient_request_total_releaseVersion	This tag indicates the current release version of ingress gateway	Picked from helm chart{{ .Chart.Version }}
oc_configclient_request_total_configVersion	This tag indicates the configuration version that ingress gateway is currently maintaining	Initial value is 0. Incremental value received from config server whenever there is an update from config server (0, 1, 2...)
oc_configclient_response_total_releaseVersion	This tag indicates the current release version of ingress gateway	Picked from helm chart {{ .Chart.Version }}
oc_configclient_response_total_configVersion	This tag indicates the configuration version that ingress gateway is currently maintaining	Value received from config server (1, 2...)
oc_configclient_response_total_updated	This tag indicates whether the configuration was updated or not	true/false

OAuth Metrics

Below are the metrics and their respective tags that are available in OAuth:

Table A-18 oc_oauth_validation_successful_total

Field	Details
Description	This metric is pegged in the OAuth validator implementation if the received OAuth token is validated successfully. The implementation of OAuth validator is used in Ingress Gateway.
Type	Counter
Dimension	• issuer • subject • scope

Table A-19 oc_oauth_validation_failure_total

Field	Details
Description	This metric is pegged in the implementation of OAuth validator if the validation of the received OAuth token fails. The implementation of OAuth validator is used in Ingress Gateway.
Type	Counter
Dimension	• issuer • subject • scope • reason

OAuth Metrics Common Tags

Tags	Description	Possible Values
scope	NF service name(s) of the NF service producer(s), separated by whitespaces	NA
issuer	NF instance id of NRF	NA
subject	NF instance id of service consumer	NA
reason	reason contains the human readable message for OAuth validation failure	NA

Table A-20 oc_ingressgateway_msgcopy_requests_total

Field	Details
Description	This is incremented whenever ingress request message is sent or acknowledged from Kafka.
Type	Counter
Dimension	

Table A-21 oc_ingressgateway_msgcopy_responses_total

Field	Details
Description	This is incremented whenever ingress response message is sent or acknowledged from Kafka.
Type	Counter
Dimension	

B

Egress Gateway Metrics

Egress Gateway Metrics

Table B-1 oc_egressgateway_http_requests_total

Field	Details
Available Tags	• Method • NFType • NFServiceType • Host • HttpVersion • Scheme • Proxy • InstanceIdentifier
Pegging Instance	This metric is pegged as soon as the request reaches the Egress Gateway in the first custom filter of the application.

Table B-2 oc_egressgateway_http_responses_total

Field	Details
Available Tags	• Status • Method • NFType • NFServiceType • Host • HttpVersion • Scheme • InstanceIdentifier • Direction • BlacklistedFqdn
Pegging Instance	This metric will be pegged in the last custom filter of the Egress gateway while the response is being sent back to backend NF microservice with direction as egress. This will also be pegged when the response is fetched in Jetty responseListener with direction as egressOut. BlacklistedFqdn tag will be filled with BlacklistedFqdn when request is sent with blacklisted producer.

Table B-3 oc_egressgateway_request_latency_seconds

Field	Details
Available Tags	- quantile - InstanceIdentifier
Pegging Instance	This metric is pegged in the last custom filter of the Ingress Gateway while the response is being sent back to the consumer NF. This metric tracks the amount of time taken for processing the request. It starts as soon as the request reaches the first custom filter of the application and lasts till, the response is sent back to the consumer NF from the last custom filter of the application.

Table B-4 oc_egressgateway_connection_failure_total

Field	Details
Available Tags	- InstanceIdentifier - Direction
Pegging Instance	This metric will be pegged in the customized Jetty Client as soon as it fails to connect to the destination service. Here in case of Egress gateway, the destination service will be Producer NF. This will also be pegged when the request to Producer NF fails in Jetty request Listener with direction as egressOut

Table B-5 oc_egressgateway_notification_ratelimit_total

Field	Details
Available Tags	- Method - Scheme - InstanceIdentifier
Pegging Instance	This metric is pegged in the custom filter implemented to check the notification rate limit conditions.

Table B-6 oc_egressgateway_request_processing_latency_seconds

Field	Details
Available Tags	- quantile - InstanceIdentifier
Pegging Instance	This metric is pegged in the last custom filter of the Egress Gateway while the response is sent back to the consumer NF. This metric tracks the amount of time taken for processing the request only within Egress Gateway. It starts as soon as the request reaches the first custom filter of the application and lasts till the request is forwarded to the destination.

Table B-7 oc_egressgateway_jetty_request_stat_metrics_total

Field	Details
Available Tags	- event - client_type - InstanceIdentifier
Pegging Instance	This metric is pegged for every event occurred when a request is sent to EGW

Table B-8 oc_egressgateway_jetty_response_stat_metrics_total

Field	Details
Available Tags	- event - client_type - InstanceIdentifier
Pegging Instance	This metric is pegged for every event occurred when a response is received by EGW

Table B-9 oc_egressgateway_jetty_latency_seconds

Field	Details
Available Tags	- quantile - InstanceIdentifier
Pegging Instance	This metric is pegged in Jetty response listener that captures the amount of time taken for processing of the request by jetty client

Table B-10 oc_egressgateway_netty_latency_seconds

Field	Details
Available Tags	- quantile - InstanceIdentifier
Pegging Instance	This metric is pegged in Netty outbound handler that captures the amount of time taken for processing of the request by netty server

Table B-11 oc_egressgateway_request_content_metrics_total

Field	Details
Available Tags	- method - content_available - InstanceIdentifier
Pegging Instance	This metric is pegged by default filter RequestContentMetrics. It pegs whether request has request body or not and the method.

Table B-12 oc_egressgateway_blacklisted_producer_total

Field	Details
Available Tags	- NFType - NFServiceType - InstanceIdentifier - Host - Route_path
Pegging Instance	This metric is a counter. Track number of times producer is blacklisted.

Table B-13 oc_configclient_request_total

Field	Details
Available Tags	- Release version - Config version
Pegging Instance	This metric will be pegged whenever config client is polling for configuration update from common configuration server

Table B-14 oc_configclient_response_total

Field	Details
Available Tags	- Release version - Config version - Updated
Pegging Instance	This metrics will be pegged whenever config client receives response from common configuration server

Table B-15 oc_egressgateway_sbiRouting_http_responses_total

Field	Details
Available Tags	- Sbi_Fqdn - Reroute_Path - Attempt - HttpVersion - Scheme - InstanceIdentifier
Pegging Instance	This metric is pegged in the SBIRoutingFilter only when SBIRouting feature is enabled for a route to which request is sent to Egress Gateway.

Table B-16 oc_egressgateway_sbiRouting_http_requests_total

Field	Details
Available Tags	- Sbi_Fqdn - Reroute_Path - Attempt - HttpVersion - Scheme - InstanceIdentifier - ErrorOriginator
Pegging Instance	This metric is pegged in the SBIRoutingFilter only when SBIRouting feature is enabled for a route to which request is sent to Egress Gateway and when sbiRerouteEnabled is set to true and reroute mechanism is executed.

Table B-17 oc_egressgateway_outgoing_tls_connections

Field	Details
Description	Number of TLS Connections received on the Egress Gateway and their negotiated TLS versions. The versions can be TLSv1.3 or TLSv1.2.
Type	Gauge
Dimension	- NegotiatedTLSVersion - Direction - InstanceIdentifier

Table B-18 security_cert_x509_expiration_seconds

Field	Details
Description	Indicates the time to certificate expiry in epoch seconds.
Type	Histogram
Dimension	- NegotiatedTLSVersion - Host - Direction - InstanceIdentifier

Egress Gateway Metrics Common Tags

Metric Type	Available Tags	Possible Values
Method	Http method	GET, PUT, POST, DELETE, PATCH
NFType	Name of the NF Type	"UNKNOWN" (Updates are available when Ingress is 5G aware)
NFServiceType	Name of the Service within the NF	"UNKNOWN" (Updates are available when Ingress is 5G aware)
Host	(IP or fqdn): port of ingress gateway	Not Applicable

Metric Type	Available Tags	Possible Values
HttpVersion	Http protocol version (http1.1/http2)	HTTP1.1, HTTP2.0
Scheme	Http protocol scheme (http/https)	HTTP, HTTPS, UNKNOWN
Proxy	Value received for "x-custom-egress-proxy-header".	Unknown or value of "x-custom-egress-proxy-header".
oc_egressgateway_connection_failure_Host	destination ip/fqdn	Not Applicable
oc_egressgateway_connection_failure_Port	destination port	Not Applicable
BlacklistedFqdn	Blacklisted Producer Fqdn	Unknown or Blacklisted Producer Fqdn
oc_configclient_request_total_releaseVersion	This tag indicates the current release version of egress gateway	Picked from helm chart{{ .Chart.Version }}
oc_configclient_request_total_configVersion	This tag indicates the configuration version that egress gateway is currently maintaining	Initial value is 0. Incremental value received from config server whenever there is an update from config server (0, 1, 2...)
oc_configclient_response_total_releaseVersion	This tag indicates the current release version of egress gateway	Picked from helm chart{{ .Chart.Version }}
oc_configclient_response_total_configVersion	This tag indicates the configuration version that egress gateway is currently maintaining	Value received from config server (1, 2...)
oc_configclient_response_total_updated	This tag indicates whether the configuration was updated or not	true/false

Oauth Metrics

Table B-19 oc_oauth_nrf_request_total

Field	Details
Available Tags	- ConsumerNFInstanceId - ConsumerNFType - TargetNFType - TargetNFInstanceId - scope - NrfFqdn
Pegging Instance	This metric is pegged in the OAuth client implementation if the request is sent to NRF for requesting the OAuth token. OAuth client implementation will be used in Egress gateway.

Table B-20 oc_oauth_token_cache_total

Field	Details
Available Tags	- ConsumerNFInstanceId - ConsumerNFType - TargetNFType - TargetNFInstanceId - scope
Pegging Instance	This metric is pegged in the OAuth Client Implementation if the OAuth token is found in the cache.

Table B-21 oc_oauth_validation_successful_total

Field	Details
Available Tags	- issuer - subject - scope
Pegging Instance	This metric is pegged in OAuth validator implementation if the received OAuth token is validated successfully. OAuth validator implementation is used in Ingress Gateway.

Table B-22 oc_oauth_validation_failure_total

Field	Details
Available Tags	- issuer - subject - scope - reason
Pegging Instance	This metric is pegged in OAuth validator implementation if the validation of the the received OAuth token is failed. OAuth validator implementation is used in Ingress gateway.

OAuth Metrics common tags

Common Tags	Description	Possible Values
ConsumerNFInstanceId	NF instance id of the NF service consumer	Not Applicable
ConsumerNFType	The NF type of the NF service consumer	NRF, UDM, AMF, SMF, AUSF, NEF, PCF, SMSF, NSSF, UDR, LMF, GMLC, 5G_EIR, SEPP, UPF, N3IWF, AF, UDSF, BSF, CHF, NWDAF
TargetNFType	The NF type of the NF service producer	NRF, UDM, AMF, SMF, AUSF, NEF, PCF, SMSF, NSSF, UDR, LMF, GMLC, 5G_EIR, SEPP, UPF, N3IWF, AF, UDSF, BSF, CHF, NWDAF

Common Tags	Description	Possible Values
TargetNFInstanceId	NF instance id of the NF service producer	Not Applicable
scope	NF service name(s) of the NF service producer(s), separated by whitespaces	Not Applicable
StatusCode	Status code of NRF access token request	Bad Request, Internal Server Error, and so on (HttpStatus.*)
ErrorOriginator	from where the error is originated (nrf or egress)	Nrf, EgressGW
issuer	NF instance id of NRF	Not Applicable
subject	NF instance id of service consumer	Not Applicable
reason	reason contains the human readable msg for OAuth validation failure	Not Applicable
NrfFqdn	NrfFqdn tag determines the corresponding fqdn of NRF where the request has been forwarded to.	Nrf-Fqdn (dynamic value based on Fqdn), NA
NrfClientUrl	This tag determines the url of NRF-Client Mgmt Svc where subscription requests are sent from OAuth Client module in EGW.	URL of NRF-Client Mgmt Svc (Dynamic value)
EgwNotificationUrl	This tag determines the notification URL mapped in OAuth Client module of EGW where NRF-Client Mgmt Svc will send notifications requests.	Notification URL (Dynamic value)
ConfigurationType	This tag determines the type of configuration in place for OAuth Client in Egress Gateway. If nrfClientQueryEnabled Helm parameter in oauthClient Helm configurations at Egress Gateway is false then the ConfigurationType is STATIC, else DYNAMIC.	STATIC, DYNAMIC

Table B-23 oc_egressgateway_msgcopy_requests_total

Field	Details
Description	This is incremented whenever egress request message is sent or acknowledged from Kafka.
Type	Counter
Dimension	

Table B-24 oc_egressgateway_msgcopy_responses_total

Field	Details
Description	This is incremented whenever egress response message is sent or acknowledged from Kafka.
Type	Counter
Dimension	

C

HTTP Status Codes Supported on SBI

Table C-1 HTTP Status Codes Supported on Service Based Interface (SBI)

HTTP Status Code	Description
204	There is no PCF session binding information matching the query parameters.
400	- The HTTP request contains an unsupported API name or API version in the URI. - The HTTP request has an invalid format. - The HTTP request contains an unsupported query parameter in the URI. - The HTTP request contains an invalid value for a mandatory parameter. - The HTTP request contains a semantically incorrect value for an optional parameter. - A mandatory query parameter is missing in the HTTP request. - The HTTP request contains a semantically incorrect value for a mandatory IE. - The HTTP request contains a semantically incorrect value for an optional IE. - A mandatory IE for an HTTP method is not included in the payload body of the request. - The request is rejected due to a malformed request. - The BSF found more than one binding resource. Therefore, the selected PCF cannot be provided.
401	Unauthorized With Header "WWW-Authenticate"
403	Request is forbidden.
404	The request for modification or deletion of a subscription is rejected because the subscription is not found in the NF.
415	The HTTP request contains an unsupported payload type.
429	The request is rejected due to excessive traffic, which if continued over time, may lead to (or may increase) an overload situation.
500	- The request is rejected due to insufficient resources. - The request is rejected due to unspecified reasons at the NF. - The request is rejected due to a generic error condition in the NF.
503	The NF experiences congestion and performs overload control that does not allow the request to be processed.
504	This error code is generated in case of timeout due to inactivity.

D

Error Code Dictionary

Table D-1 Error Code Dictionary for BSF Management

Application Error ID	Reason	Cause	Action
EC-OBSF-BSF_MGMT-REQVLD-EI-05-01-400-00005-NA-02	The GET request to BSF failed with a 400 status code and cause MANDATORY_QUERY_PARAM_INCORRECT.	The error appears whenever the Binding Request Query contains multiple UE Addresses (for example, a MAC-48 Address and an IPv4 Address), or the mandatory query parameters have an incorrect format.	Check the error details. Based on the error details, Verify that only either UE Address type IPV4 address, or IPv6 Prefix, or MAC-48 Address is sent. Ensure that each mandatory parameter is present and formatted correctly. Correct any errors and resend the request.
EC-OBSF-BSF_MGMT-REQVLD-EI-05-01-400-00006-NA-02	The GET request to BSF failed with a 400 status code and cause OPTIONAL_QUERY_PARAM_INCORRECT.	The error appears whenever the Binding Request optional query parameters has an incorrect format or value.	Check the error details. Based on the error details, review the optional query parameters (such as DNN and SNSSAI) for correct format. Update any values that do not comply before resubmitting the discovery request.
EC-OBSF-BSF_MGMT-REQVLD-EI-05-01-400-00007-NA-02	The GET request to BSF failed with a 400 status code and cause MANDATORY_QUERY_PARAM_MISSING.	The error appears whenever the Binding Request Query does not contain a mandatory parameter. Mandatory parameters are the following: • UE Address – IPv4 Address and/or Ipv6 Prefix – MAC-48 Address	Verify that all mandatory parameters are included in the Binding Request: ensure the UE Address is specified, and at least one valid identifier (IPv4 Address, IPv6 Prefix, or MAC-48 Address) is provided.
EC-OBSF-BSF_MGMT-INTRNL-EI-05-02-500-00000-01-02	The POST request to BSF failed with a 500 status code and cause INTERNAL_SERVER_ERROR.	The error appears whenever an internal error occurs on BSF Management Service while trying to register.	Retry the request. If the error persists, collect logs and contact your system administrator or Oracle Support for further assistance.

Table D-1 (Cont.) Error Code Dictionary for BSF Management

Application Error ID	Reason	Cause	Action
EC-OBSF-BSF_MGMT-REQVLD-EI-05-02-400-00003-01-02	The POST request to BSF failed with a 400 status code and cause INVALID_MSG_FORMAT.	The error appears whenever the Binding Data from the Binding Request contains multiple UE Addresses (For example, a MAC-48 Address and an IPv4 Address and/or IPV6 prefix), or the IP Address is missing its IPv4 Address or IPV6 Prefix or ipDomain is present without IPV4 Address	Ensure that the Binding Data includes only one type of UE Address per request (either MAC-48 Address or IP Address, not both), and if using IP Address, the IPV4 Address or IPV6 prefix must be specified. Ensure IPDomain is present only when IPV4 address is available.
EC-OBSF-BSF_MGMT-REQVLD-EI-05-02-400-00008-01-02	The POST register request to BSF failed with a 400 status code and cause MANDATORY_IE_INCORRECT.	The error appears whenever a mandatory IE in the Binding Data from the Binding Request has an incorrect format.	Check all mandatory fields (IPv4, IPv6, MAC-48, SNSSAI, DNN) in the request for correct formatting and valid values. Correct any invalid entries and retry the request.
EC-OBSF-BSF_MGMT-REQVLD-EI-05-02-400-00009-01-02	The POST request to BSF failed with a 400 status code and cause OPTIONAL_IE_INCORRECT	The error appears whenever the Binding Request optional query parameters has an incorrect format or value.	Review all optional fields (SUPI, GPSI, PCF FQDN, PCF Diameter Realm) for correct format according to the specified patterns defined in the cause column. Update any invalid values and resubmit the request.
EC-OBSF-BSF_MGMT-REQVLD-EI-05-02-400-00010-01-02	The POST register request to BSF failed with a 400 status code and cause MANDATORY_IE_MISSING.	The error appears whenever a mandatory IE is missing in the Binding Data from the Binding Request. Mandatory IE are the following: • UE Address – IPv4 Address and/or Ipv6 Prefix – MAC-48 Address • SNSSAI – SST – SD • DNN	Verify that all required mandatory Information Elements (IE) are present in the request: • UE Address • (At least one of) IPv4 Address, IPv6 Prefix, or MAC-48 Address • SNSSAI (SST and SD) • DNN Add any missing values and resend the request.

Table D-1 (Cont.) Error Code Dictionary for BSF Management

Application Error ID	Reason	Cause	Action
EC-OBSF-BSF_MGMT-INTRNL-EI-05-04-500-00056-01-02	The DELETE deregister request to BSF failed with a 500 status code and cause INTERNAL_SERVER_ERROR	The error appears whenever an internal error occurs on BSF Management Service while trying to deregister.	Retry the deregister request. If the issue continues, gather logs and escalate to your system administrator or Oracle Support.
EC-OBSF-BSF_MGMT-SIG-EI-05-04-404-00060-01-02	The DELETE deregister request to BSF failed with a 404 and cause NOT_FOUND	The error appears whenever there is no resource to delete with the provided bindingID.	Validate that the bindingID that is trying to be deleted is in fact correct, and retry. If the 404 is still obtained, it means that the resource is no longer there and no further action is needed.

Table D-2 Error Code Dictionary for Ingress Gateway

Application Error Code ID	Reason	Cause	Action
OBSF-IGW-E002	REQUEST_TIMEOUT This error code indicates that request processing takes more time than the configured request timeout at ingress-gateway.	A successful TCP connection was established but no response was received within the expected time frame.	To resolve this issue, consider increasing the request timeout setting on the ingress gateway or optimizing the backend services to handle the request faster.
OBSF-IGW-E054	INTERNAL_ERROR This error code indicates that some unknown internal error occurred in ingress-gateway while processing the request.	Internal Errors, occur when there's an issue with improper configurations, mismatch in helm charts and versions.	Identify the cause from the error message sent back in the response and refer to the documentation for the specific feature that is being used and re verify all the configurations. Helm Charts and respective versions can also be checked

Table D-2 (Cont.) Error Code Dictionary for Ingress Gateway

Application Error Code ID	Reason	Cause	Action
OBSF-IGW-E102	LATE_ARRIVAL_EXCEPTION This error code indicates that late arrival exception occurred at ingress-gateway as the request has already expired as part of SbiTimer feature.	Timezone mismatch (or) 12H/24H time format mismatch leading to some time in the past in SBI Timer Header for SBI Timer Feature (or) Header of MaxRspTime might have been configured with the wrong value (or) the calculated request timeout value based on SBI Timer Headers is less than the configured requiredTime value at route level	Make sure the timestamp and max response time attached in the header, are of the correct time with respect to the time zone and not a time in the past. Verify if the requiredTime is accommodating to the calculated timeout value from SBI Timer Header values.
OBSF-IGW-E122	GRACEFUL_SHUTDOWN This error code indicates that requests received at ingress-gateway rejected gracefully during the ingress-gateway shutdown.	This indicates that the service/pod is shutting down and responses sent include this Error Response	(Re) Deploy IGW, or prevent shutdown
OBSF-IGW-E162	XFCC_HEADER_VALIDATION_FAILURE This error code indicates that validation of XFCC header failed at ingress-gateway.	Indicates that the XFCC Header in the request is in the correct format but is not validated as it is the wrong value for the request.	Verify if the right value is updated in the XFCC Header based on the IGW Deployment configuration.
OBSF-IGW-E163	XFCC_HEADER_NOT_PRESENT_OR_EMPTY This error code indicates that request received at ingress gateway with no XFCC header or empty XFCC header.	XFCC Header is not present	Include the XFCC Header in the Request
OBSF-IGW-E164	XFCC_MATCHCERTCOUNT_GREATER_THAN_CERTS_IN_HEADER This error code indicates that XFCC header validation failed due to XFCC certificate count miss match in ingress-gateway.	IGW Configuration requires more certs than the number of certs sent in the request	Check the number of certificates in the request sent and attach the certificates matching the configuration set for IGW
OBSF-IGW-E165	XFCC_HEADER_INVALID_ID This error code indicates that invalid XFCC header received at ingress-gateway.	XFCC Header present in request is not of the right format.	Fix the XFCC Header of the request and send.

Table D-2 (Cont.) Error Code Dictionary for Ingress Gateway

Application Error Code ID	Reason	Cause	Action
OBSF-IGW-E174	OAuth_PRODUCER_SCOPE_NOT_PRESENT This error code indicates that producer scope not present in the OAuth token received at ingress-gateway.	Producer Scope is not present in the token of the request	Check for Producer Scope value in the token used in the request
OBSF-IGW-E175	OAuth_PRODUCER_SCOPE_MISMATCH This error code indicates that producer scope present in the OAuth token miss match with producer scope configured in ingress gateway.	Producer Scope sent through the token either has wrong format, has been used incorrectly, or incorrect value is present.	Verify the Producer Scope Attribute of the token sent through the request
OBSF-IGW-E179	OAuth_TOKEN_INVALID This error code indicates that invalid OAuth token received at ingress-gateway.	Indicates that the OAuth Token received at IGW is not valid, that is., not in the correct format.	Verify if the OAuth Token is in the right format and configured with the right set of values with respect to the configuration at the gateway.
OBSF-IGW-E180	OAuth_TOKEN_ABSENCE This error code indicates that request received at ingress gateway with no OAuth token or empty OAuth token	OAuth Token is missing in the request	Include the OAuth Token in the request
OBSF-IGW-E192	PARTIAL_SHUTDOWN This error code indicates that requests are being discarded partially at ingress-gateway when Control shutdown feature enabled.	Behavior is observed when ControlledShutdown is enabled	Disable ControlledShutdown
OBSF-IGW-E193	COMPLETE_SHUTDOWN This error code indicates that requests are being discarded completely at ingress-gateway when Control shutdown feature enabled.	Behavior is observed when ControlledShutdown is enabled	Disable ControlledShutdown

Table D-2 (Cont.) Error Code Dictionary for Ingress Gateway

Application Error Code ID	Reason	Cause	Action
OBSF-IGW-E201	IGW_OVERLOAD_CONTROL_PERCENTAGE_DISCARD This error code indicates that requests are being rejected when load level matches with configured percentage in discard policy mapping of overload control feature.	The number of requests have reached the configured percentage for OverloadControl. When this value has reached, the respective Error Code is displayed	Change the percentage value to accommodate more requests and not to display the error code.
OBSF-IGW-E202	IGW_OVERLOAD_CONTROL_PRIORITY_DISCARD This error code indicates that requests with sbi-priority header with value greater than configured number of discarded load level matches in discard policy mapping of overload control feature.	The number of requests have reached the configured priority for OverloadControl. When this value has reached, the respective Error Code is displayed	Change the priority value to accommodate more requests and not to display the error code.